AusIMM Conferences

STRATEGIC MINE PLANNING CONFERENCE 2026

PERTH, AUSTRALIA
17–19 NOVEMBER 2026

Conference Proceedings



Strategic Mine Planning Conference 2026

17–19 November 2026
Perth, Australia

The Australasian Institute of Mining and Metallurgy
Publication Series No 6/2026



Published by:
The Australasian Institute of Mining and Metallurgy
Ground Floor, 204 Lygon Street, Carlton Victoria 3053, Australia

© The Australasian Institute of Mining and Metallurgy 2026

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form by any means without permission in writing from the publisher.

All papers published in this volume were peer reviewed before publication.

The AusIMM is not responsible as a body for the facts and opinions advanced in any of its publications.

ISBN 978-1-922395-63-4

Advisory Committee

Roussos Dimitrakopoulos
FAusIMM
Conference Advisory Committee Chair

Marcelo Godoy
FAusIMM(CP)

Ryan Goodfellow

Vaughan Chamberlain
FAusIMM

Orlando Leiva
MAusIMM

Philip Bangerter
FAusIMM

Vanessa Lickfold
FAusIMM

Nigel James
MAusIMM(CP)

AusIMM

Julie Allen
General Manager, Events

Michelle Portelli
Senior Manager, Events

Sasha Pell
Senior Manager, Conference Growth, Events

Elly Richards
Program Manager, Events

James Masters
Senior Manager, Events Marketing and Communications, Events

Reviewers

We would like to thank the following people for their contribution towards enhancing the quality of the papers included in this volume:

Erkan Aslan

Victor Balboa Espinoza

Philip Bangerter

Christian Both

Laura Carelos Andrade

Ivan Carvalho

Joao Pedro De Carvalho

Roussos Dimitrakopoulos

Tarrant Elkington

Matheus Faria

Liam Findlay

Ryan Goodfellow

Julien Herrero

Steven Hoerger

Nigel James

Ashish Kumar

Vanessa Lickfold

Jingyu Luo

Mahlomola Isaac Mabala

Angus Macrae

Luis Montiel Petro

Lidiia Shchichko

Gavin Yeates

Foreword

On behalf of the Organising Committee, our partner institutes and sponsors, we welcome you to explore these proceedings of the AusIMM Strategic Mine Planning Conference 2026 (SMP 2026), which was held in Perth, Western Australia from 17–19 November 2026.

Under the theme Digitalisation and Intelligent Systems Transforming Mineral Value Chains, SMP 2026 brought together leading researchers and industry practitioners from Australia and around the world to examine the evolving challenges and opportunities in strategic mine planning. Building on the legacy of previous symposia held in 2004, 2009 and 2014, the conference provided an important forum for sharing knowledge, challenging established approaches and exploring innovations. These proceedings capture that exchange and offer a resource that will help shape the future of the global mining industry.

Strategic mine planning continues to play a critical role in navigating uncertainty, optimising mineral assets and delivering long-term value. With increasing demand for critical minerals, the energy transition and growing expectations around sustainable development, the industry must embrace new approaches to planning, optimisation and decision-making. Advances in stochastic modelling, digitalisation, machine learning and intelligent systems are creating significant opportunities to improve the resilience, efficiency and sustainability of mineral value chains.

Across its pages, this book showcases international expertise and technical advances in orebody modelling, strategic mine planning optimisation, integrated mining complexes and intelligent decision-making systems. It explores how emerging technologies can bridge the gap between strategic planning and operational execution, while addressing the industry's broader priorities of electrification, decarbonisation and sustainable resource development.

I sincerely thank our sponsors — AngloGold Ashanti, MAPTEK, Rio Tinto, Whittle Consulting, COGNIMINE, KINROSS and Gold Fields — together with our exhibitors, whose support and commitment have helped make possible this event and accompanying proceedings.

I also acknowledge AusIMM and our supporting partners, SAIMM, CIM and SME, for their continued collaboration and their dedication and professionalism in producing such work.

I hope these proceedings provide a valuable platform for technical exchange, meaningful collaboration and new perspectives, and that the knowledge shared within them contributes to advancing strategic mine planning and the future of our industry.

Yours faithfully,

Roussos Dimitrakopoulos
FAusIMM
Strategic Mine Planning Conference 2026 Advisory Committee Chair

Sponsors

Platinum Conference Sponsors









Gold Sponsor



Silver Sponsor



Coffee Cart Sponsor



Contents

Keynotes

Why mine plans fail – value, predictability, flexibility and sustainability 3
M Godoy

Commodity prices have increased significantly in the last two years – have you adjusted your Life Of Asset (LOA) plan yet? Have you adjusted it in the right way? 11
G Whittle

Beyond naive optimisation, revisited 13
G A Yeates

Unlocking value in mining projects – a stochastic mine planning workflow for complex gold-copper leach operations under geological uncertainty 17
A Prawasono and H Camero

From asset review to executable growth – transforming strategic mine planning at Vale Base Metals 21
M Silva, J Todd, B Gadelha and T O'Neill

Holistic optimisation integrating agriculture into mining for sustainable value 33
B King

Planning for uncertainty – the next frontier in strategic mine planning 43
O Leiva

Unlocking managing resource uncertainty in a large mining organisation 45
A Jewbali

AI and machine learning for mine operations

Reinforcement learning-based stochastic optimisation of short-term production scheduling and shovel allocation incorporating truck dispatching policies and environmental constraints 49
V Balboa-Espinoza and R Dimitrakopoulos

Learning interpretable dispatch heuristics for real-time open pit truck dispatching under uncertainty 69
J Luo and R Dimitrakopoulos

LOforSP – Learning-to-optimise for two-stage stochastic programs with continuous recourse 73
L Shchichko, S S Chauhan, E Delage and R Dimitrakopoulos

Closure and sustainable mine planning

Regenerative mine planning – embedding closure, carbon sequestration and circular land use into surface mine design 79
E Chauke and M Mochubele

Digital solutions for sustainable mining – an optimisation methodology for pit-limit and extraction-sequence definition to reduce mining impacts and closure liabilities 97
F Egger and L Fouché

Integrating closure in the mine plan – practical ways to reduce risk and cost, and improve ESG outcomes 107
C Lison, R Irwin, R Ward and J Heyes

Commodity prices and capital allocation

Strategic capital allocation methods for commodity cycle resilience 119
E Deucker and J Chung

Deterministic and stochastic metal price models and implications for strategic mine planning 137
S F Hoerger, T Brady and K Dagdelen

Cut-off grade and economic optimisation

A case study of blending-based cut-off selection in an open pit iron ore mine 155
C Almeida and R Peroni

Leveraging underutilised statistical techniques and low-cost data to drive optimisation and generate value in strategic life-of-mine planning – a case study 173
L Lawrence and K Whitchurch

Cut-off grade as a variable – an emergent property of multi-constraint mine plan optimisation at Evolution Mining's Red Lake operations 177
B Talbot, J Carswell and L Jordan

Fleet, equipment and blasting in mine planning

Diesel hazards to zero-emission safety – techno-economic feasibility of mining fleet electrification at AngloGold Ashanti's Iduapriem gold mine, Ghana 193
A Acheampong, E O Ansah, L K Dadzie and C Abuntori

Incorporating maintenance-driven equipment costs into mine planning 197
D Kaykov and M Kumral

Multi-objective blast optimisation as a decision variable in mine production scheduling 213
R Supriatna, L Rowe and S Raval

Geological modelling, uncertainty and simulation

Effects of non-stationary geostatistical simulation on open pit production scheduling 227
A Aydar and K Dagdelen

Localised conditional simulation in mine planning – bridging deterministic and stochastic workflows 245
T Elkington and R Santos

A kernel mean embedding method for high-order geostatistical simulation in orebody modelling 259
J Herrero and R Dimitrakopoulos

Dynamic structural modelling for near real time inputs for geotechnical risk management and optimisation 279
A Jani

Generative geostatistical simulation using Restricted Boltzmann Machine, comparison with VAEs and GANs 281
Y Mu and F Ferrie

Geometallurgy and Mine to Mill integration

Strategic Mine to Mill as a risk-informed value system 293
F Faramarzi

Assessing the impact of nonlinear combinations from the geometallurgical models and geological uncertainty within the strategic mine planning 311
G Gariba, D Alegre, M Faria, A Mayrink, L C Andrade, R Peroni and J L V Mariz

Debottlenecking underground stoping operations using an integrated Mine to Mill-to-backfill digital twin and optimisation model 321
A Kumar, R Goodfellow and K Dempster

Geometallurgy in mine planning

Quantifying geometallurgical uncertainty for mine planning using machine learning and conditional simulation 327
C Artica, M Noppé and N Fox

A geometallurgical framework using operational mining data for plant performance prediction and stochastic mine production scheduling 333
C Both

Mine scheduling, algorithms and pit optimisation

Multi-objective genetic algorithm for strategic open pit mine planning – generating pareto-optimal schedules for value-feasibility trade-offs 339
J Andre, E Fosdike and M Roberts

Pit optimisation for large models in the presence of geologic and economic uncertainty 343
M Deutsch, J Deutsch and C Deutsch

Strategic conveyor guided mine sequencing for cost stability and value optimisation in nickel open pit operations 347
A R Wilogo

Mining complexes and value chain optimisation

Integrating planning horizons and capital investments for mining complexes with uncertainty using multi-level stochastic optimisation 367
L Findlay and R Dimitrakopoulos

Strategic value optimisation for an integrated copper open pits and underground mining complex 373
B Gadelha, T Câmara, L Montiel and J Todd

Strategic optimisation of the Voisey's Bay mining and processing value chain 379
D Galloway, R Mallett, B Gadelha, L Montiel, T Câmara and J Todd

Assessing strategic production schedules of mining complexes and project evaluation impacts relating to uncertainty 391
W Patton and B Marriott

Production planning under geological uncertainty

Optimisation of selective mining unit (SMU) size through misclassification and economic analysis using sequential Gaussian simulation (SGS) – application to a copper deposit 401
N Bouhlali, A Elghali, A Moradi, Y Taha and M Benzaazoua

Unlocking opportunities through a stochastic production plan optimisation of an iron ore greenfield mining project 405
J Carvalho, A Mayrink, A Amaral, D Barbosa, A Ditlef, G Santos and A Silva

Generating operational mine schedule using a stochastic block-based simultaneous optimisation framework 417
A Mayrink, M Faria, R Senécal, L Silva and A Lioret

Incorporating geological, equipment, and geotechnical uncertainty into mine planning 427
I Minniakhmetov and Z Fletcher

From risk assessment to stochastic optimisation – a systematic approach to de-risk an operating gold mine production schedule 431
A Rispoli, L C Andrade and D Alegre

Strategic mine planning – value, governance and decision-making

Strategic value driven planning through integrated multi-mine optimisation 437
F Damasceno, J Sega and P Silva

Value generation in mining – does mineral resource management still matter? 441
P Knights, H Burgers and C Aji

Closing the gap between strategic and tactical mine planning 453
E Ponce, A Hannett and A Flett

Beyond the pit shell – a framework for strategic mine planning to challenge constraints and support value-driven decisions 457
B Price

Mind the gap – aligning strategic vision with technical evidence in strategic mine planning 465
A Suchak

Strategic planning and optimisation

Incorporating mining direction into strategic mine planning – modelling and implementation 481
I Rudge, S Hill, D Place and C Poblete

Strategic planning in practice – people, projects and decision support

The role and application of techno-economics theory in the development of mining projects 495
M J Fox-Martin and D J de Villiers

Best practices for uncertainty management in mineral projects 519
M Silva, B Gadelha, L Montiel and T Câmara

From file-based workflows to decision support systems – a framework for scalable mine planning in the cloud era 527
L Victor and J Mackenzie

AI-enabled strategic planning and why people are the mining industry's real asset 539
B Yaqub

Underground mine design and method selection

Transforming mine planning – auto-design rules for rapid underground layouts in Deswik and Datamine 547
A Saxena and J R Town

Underground mining method selection for narrow veins at depth at Cuiabá Mine 565
R P S Silva and M M Oliveira

Underground planning – caving and stope optimisation

Integration of machine learning and mixed-integer programming in underground mine planning 579
P Chimunhu, E Topal, M W A Asad, R S Faradonbeh and A D Ajak

Geotechnical simulation for block caving – quantifying rock-mass uncertainty to inform long-term planning 585
S Light, M Musolino, R Barnes, C Chester and I Minniakhmetov

Simultaneous stochastic optimisation of underground mining complexes – production and development scheduling with multiple design alternatives 589
A MacRae and L Petro

A stochastic economic block-model framework for optimising underground stope layouts under grade uncertainty for a platinum-group-element deposit 603
A S Nhleko, M I Mabala and C Musingwini

Dynamic optimisation of integrated mine value chains – a tractable MILP framework for adaptive decision-making 617
J Qu, C Xu and P Dowd

Author index 621

Keynotes

Why mine plans fail – value, predictability, flexibility and sustainability

M Godoy[1]

1. Chief Technology Officer, AngloGold Ashanti, Denver CO 80111, USA.
   Email: mgodoy@aga.gold

ABSTRACT

In the mining sector, mine plan quality has direct capital market consequences. Investors reward predictable delivery and credible growth, and they penalise repeated guidance misses. In practice, many plans fail because deterministic schedules are applied to operating systems shaped by geological, operational, processing, and external uncertainties. This paper argues that effective mine planning should be understood as the design of a robust operating strategy rather than the production of a static schedule. A successful plan must achieve four objectives simultaneously: maximise durable value, support predictable execution, ensure sustainability, and preserve enough flexibility to absorb deviations without damaging long-term performance. To develop a practical framework that makes these plans resilient, this approach emphasises the necessity of credible geological inputs, realistic operating constraints, integrated mine-to-market logic, explicit treatment of uncertainty, and predesigned contingencies that can be activated before a disruption becomes value-destructive.

INTRODUCTION

Mine planning sits at the intersection of physical safety, resource evaluation, operational delivery, and long-term value creation. Yet many plans are still treated primarily as geometric extraction schedules whose purpose is to maximise nominal net present value (NPV) under static assumptions. That view is dangerously incomplete. Mining operates as a non-linear, stochastic system in which geology, rock mass mechanics, equipment performance, process response, labour productivity, water and waste management, permitting, social licence and market conditions interact continuously. Under these conditions, deterministic schedules optimised around a single nominal case are fragile in execution.

The central issue leading to failure is not the existence of uncertainty; it is the practice of treating uncertainty as external to the optimisation engine. Routine deviations in grade, tonnage, rock mass hardness, slope stability, equipment availability, or plant throughput do not merely create short-term noise. They degrade cash flows, compromise long-term ore access, increase energy intensity, and expose personnel to elevated geomechanical hazards.

Accounting for uncertainty by adding arbitrary per cent buffers or running unconstrained Monte Carlo simulations fails to address the pivotal challenge: designing a schedule that is inherently resilient against internal and external disruption. As Rumelt (2011) emphasises, effective strategy requires a coherent response to the core structural challenge.

The thesis of this paper is that a successful mine plan must not be judged solely by headline NPV or schedule compliance. Rather, it must be evaluated as an operating strategy designed to convert a mineral endowment into repeatable cash flow while preserving sustainability and long-term operational optionality. The sections that follow outline why deterministic plans fail, establish an operating strategy framework anchored in four non-negotiable pillars, and detail the governance required to protect value under operating uncertainty.

WHY DETERMINISTIC MINE PLANS FAIL IN PRACTICE

Mine plans are meant to impose structure on a complex operating system, yet they often fail because the assumptions embedded in the plan are treated as stable representations of conditions that are inherently variable. Divergence between plan and realisation is normal rather than exceptional.

Geological variance represents the most fundamental source of instability. No subsurface model can perfectly represent the orebody as it will be mined. Variability in grade, tonnage, hardness, structures, ore-waste boundaries, dilution, and metallurgical response can materially change both

mining sequence and plant performance. Plans built around point estimates without explicit allowance for this variability are vulnerable from the outset.

Beyond the subsurface model, execution shortfalls routinely fracture plan compliance. Deterioration in equipment availability, cycle-time, fragmentation quality, unexpected processing plant downtime, haul-road condition, short- versus long-term compliance, and contractor performance directly affect the mine's ability to deliver scheduled development rates and ore feed. These are not minor disturbances around an otherwise valid plan. They are structural conditions of executability.

Furthermore, a failure to integrate processing constraints into the optimisation function creates a false sense of operational viability. Throughput sensitivity to hardness and fines generation, recovery variability, water and reagent constraints, and tailings limits can materially change the value of a mining sequence. When mine plans are developed without sufficient mine-to-waste integration, an operation may appear compliant against mining targets while underperforming in throughput, recovery, and costs.

External disturbances add further instability. Power reliability, fuel costs, supply-chain disruption, permitting delays, weather events, community-related constraints, and commodity-price movements can all change the feasibility or attractiveness of the planned sequence.

Ultimately, the inherent fragility of deterministic mine plans stems from their rigid reliance on static assumptions within an intrinsically volatile environment. When compounding geological uncertainties, routine operational execution shortfalls, complex processing constraints, and unpredictable external market volatility are treated as stable anomalies rather than structural realities, divergence from the strategic plan becomes inevitable. Because relying on a single estimated resource model and fixed variables is a major cause of technical and financial risk, the mining industry must abandon static forecasting in favour of dynamic, risk-adjusted frameworks that actively embrace and optimise for these inevitable variables.

A STRATEGIC FRAMEWORK FOR PLANNING UNDER UNCERTAINTY

Mine planning should be executed not as a spatial geometry puzzle, but as the design of a robust operating strategy. This distinction aligns with Mintzberg's (1994) separation of analytical strategic planning from true strategic thinking.

A successful mine plan is thus an integrated system for converting a finite mineral endowment into economic value under uncertain operating conditions without violating safety and environmental boundaries. Its quality depends on how well it balances four interdependent pillars: durable economic value, predictable execution, built-in flexibility and sustainability. These objectives are tensioned (Figure 1). Plans that chase maximum nominal value often become brittle, while heavily buffered plans can dilute value. Robust planning depends on conscious trade-offs rather than one-dimensional optimisation. Sustainability serves as the absolute physical boundary; a plan that compromises ground stability or tailings storage limits to chase cash flow is fundamentally invalid.

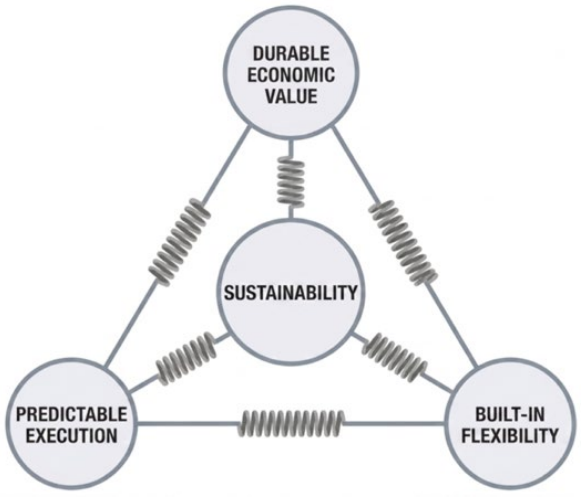


FIG 1 – The four pillars of a robust operating strategy.

Durable economic value

Value maximisation in mine planning means more than delivering tonnes or prioritising near-term high-grade. A plan that improves short-term performance at the expense of stripping, development, blending stability, or future sequencing flexibility can look attractive in the model while destroying value over the planning horizon.

In addition, a mine plan is fundamentally incomplete if it ignores the downstream consequences of upstream decisions. This systemic integration finds its technical roots in the work of Lynch (1977), who pioneered the simulation and optimisation of crushing and grinding circuits as part of a connected system. Mine plans are more robust when extraction sequence and processing performance, both affected by hardness, mineralogy, and blend composition, are modelled together. By assessing value at the processing plant, or more broadly at the mine-to-market level, rather than solely at the mining face, the organisation improves its ability to maximise durable cash flow.

Furthermore, mine planning is a value maximisation exercise that requires complex objective functions to model the intricacies and non-linearities of the full value chain. A significant amount of value has been and continues to be eroded from mining projects and operations because mine planners oversimplify the objective function. Classic examples of that are: not differentiating fixed from variable costs, using deterministic models, ignoring taxes and royalties, assigning a block as ore or waste prematurely, not modelling critical geometallurgical attributes as part of the objective function, not considering a realistic view of stay-in-business capital and the most common, maximising nominal NPV as the sole optimisation criterion. Ultimately, overly simplified economic models yield unrealistic optimisation results.

The true objective is durable economic value, estimated through a distribution of simulated NPVs and cash flows across the full value chain. From such distributions, decision-makers can balance their risk appetite with the probability of achieving desired returns on investment. Practically, this means abandoning a single, fragile financial forecast in favour of stochastic planning results, which allow planners to visualise a series of possible financial outcomes. By applying a risk-adjusted utility function, as introduced by Rendu (2017), to this distribution of outcomes, an executive team can, for example, mathematically penalise scenarios that present a risk of catastrophic financial loss, ensuring that the final mine plan accurately reflects the organisation’s true risk tolerance.

Predictable execution

A high-value schedule has little practical meaning if the operating system cannot deliver it with reasonable consistency. Predictability does not mean eliminating variance. It means building the plan around credible assumptions about grade variability, mining performance, processing plant response, and market dynamics so that routine execution remains within a controllable band.

The goal is not to predict every possible future, but to understand which assumptions have the greatest leverage over value and where the plan becomes fragile. In Taleb’s (2012) terms, the better plan is not one that assumes a stable world, but one that is deliberately structured to withstand volatility and, where possible, benefit from it.

Dimitrakopoulos (2011) advocates the use of stochastic optimisation to move beyond the limitations of deterministic mine planning. This field has achieved a high degree of maturity in the last few years, with many real-world case studies showing the benefits of applying stochastic optimisation to mine planning. Goodfellow and Dimitrakopoulos (2017) introduced a two-stage stochastic integer programming formulation that integrates geological uncertainty into the optimisation of a mining complex value chain. The solution maximises the economic value of the mining sequence while minimising deviations from production targets. This formulation can realistically account for geological uncertainty as well as variability in costs, productivity and commodity prices, simultaneously. The result is a single, optimised physical schedule whose key attributes are presented as a series of possible outcomes, allowing for probabilistic inference.

Predictable delivery is designed, not hoped for. In practice, predictability matters because investors, lenders and boards judge an operation or a project, not only by its headline economics, but by the credibility with which those economics can be delivered.

Built-in flexibility

Because mining systems are continually exposed to operational and market volatility, plans require built-in alternatives that allow management to respond to changing conditions without sacrificing long-term asset value. This flexibility can take the form of alternative ore sources, stockpiles, contingency stopes or benches, alternative access routes, contractor options, or limited surge capacity. Furthermore, creating operational flexibility in a mine is typically linked to ongoing resource conversion and continuous exploration. Flexibility must be designed in advance rather than improvised after a disruption occurs.

Built-in flexibility represents a practical application of real options theory, as developed by Trigeorgis (1996), wherein project value is demonstrably enhanced by managerial adaptability. Brennan and Schwartz (1985) provide the foundational economic logic for this within natural resource investments, demonstrating that the ability to switch operational modes constitutes a significant portion of a project's real-world value.

However, flexibility should not be assumed to be the same as system-wide operational slack, which can be prohibitively expensive and obscure strategic priorities. A superior approach is bounded optionality, which consists of: (a) identifying a limited set of credible scenarios tied to the system's most consequential risks; (b) developing decision-ready alternatives linked to clearly identified trigger points; and (c) conducting a rigorous cost-benefit analysis of these alternatives to support investment decisions.

Sustainability

Geotechnical, social and environmental parameters represent hard physical constraints rather than soft operational targets. When sustainability parameters are divorced from the core financial optimisation framework, they are often relegated to post-optimisation compliance checks. This sequential approach invariably leads to sub-optimal outcomes, exposing operations to catastrophic geomechanical failures, regulatory sanctions, and severe value erosion. For example, optimising schedules without explicit geomechanical constraints risks accelerating extraction in high-stress zones, causing pit wall failures or underground collapses that instantly compromise workers' safety and destroy asset value.

Slope angles, stope hydraulic radii, rock burst energy thresholds, tailings storage capacity, property boundaries, and community relocation timing must bound the spatial extraction sequence. By formulating these physical realities as strict inequalities and dynamic penalty functions within the optimisation engine, mining companies force their strategic plans to reconcile financial ambition with sustainability constraints.

FROM PLAN TO EXECUTION

Even a well-designed plan degrades if the execution system is weak. In an operating mine, the 'execution gap' usually occurs because the organisation lacks a mechanism to manage high-frequency variability. To bridge this gap, the mine must implement Short-Interval Control (SIC) as the mechanism that connects strategic intent to daily delivery. By tracking actual movement, grades, development, and equipment performance in near real-time, the operator can detect deviations early and respond before a local issue becomes a structural failure. In addition, this systematic reconciliation provides management with the controls to monitor predictions and moderate expectations (Shaw *et al,* 2013). This iterative learning loop is essential to enable mining operations to address rapidly changing environments.

Short-interval control as a tactical engine

Short-interval control is the mechanism that connects strategic intent to daily delivery. From the perspective of Goldratt (1984) theory of constraints, SIC acts as the protective capacity of the system. To be effective, SIC must shift from a data collection exercise to a decision-making framework:

- **Rapid intervention**: The goal is faster, better intervention by adjusting blends, resequencing material movement, or reallocating equipment while keeping long-term value intact.

- **The learning loop**: Repeated deviations are data points that reveal which assumptions are wrong. This mirrors Senge (1990) concept of the learning organisation, in which the feedback loop forces a revision of the planning models and assumptions used in the next cycle.

Governance – aligning the organisation

Planning works best when it is treated as an operating system, not as the output of a single department. In that sense, governance ensures that operational metrics reinforce strategic goals rather than drift away from them.

Effective governance establishes clear decision rights over mine planning assumptions, design parameters, and contingency activation. As Pascale and Athos (1981) emphasise, technical strategies fail if they are not supported by organisational alignment and shared values. Proper governance allows the operation to adapt when needed, within explicit rules about who can change the plan and with what economic consequences.

An immediate, high-impact opportunity for mining companies is to breakdown functional silos between the finance team and the technical team that produces the mine plan. The collaboration between these key stakeholders is critical to produce a truly optimised and realistic mine plan.

The social system of mining

Without governance discipline, even a technically sound plan will fragment as it moves through the organisational hierarchy. Predictability and flexibility are not merely engineering outputs; they are the result of a culture that understands the economic consequences of operational deviation. If operator's key performance indicators are strictly tied to pure tonnage rather than blend compliance, or if middle management is penalised for activating contingency plans, the technical strategy inevitably breaks down. Governance must therefore ensure the mine's social system—its people, performance metrics, and incentive structures—is deeply aligned with the technical plan. Achieving this requires transparently communicating the plan's economic drivers to the workforce so that tactical, day-to-day decisions on the ground support the asset's broader strategic intent.

THE HEALTH OF MINE PLAN SCORECARD

To evaluate the likelihood of mine plan failure and diagnose where an operation sits on the spectrum between structural fragility and resilience, a formalised maturity framework is required. This scorecard evaluates the six key factors of strategic mine planning against their practical ability to deliver value, predictability, flexibility and sustainability. These key factors are value, uncertainty, realism, resilience, learning and sustainability. Rather than chasing spurious precision, this tool guides prioritisation. It is designed for executive leadership and operational management to rapidly identify sources of value leakage, exposure to guidance misses, and hidden constraints within the planning ecosystem.

The matrix in Table 1 outlines the three levels of maturity for each factor. Moving from Level 1 to Level 3 represents a shift from reactive, siloed scheduling to a resilient, risk-adjusted operating strategy. The matrix is not rigid; it is meant to be adapted to the realities of the mining asset.

TABLE 1

Health of mine plan – maturity matrix.

Factor\maturity	Level 1 – Deterministic and fragile	Level 2 – Integrated and controllable	Level 3 – Stochastic and resilient
1. Value – Quality of the optimisation's objective function	Focuses on nominal NPV and short-term grade chasing. Objective functions ignore fixed versus variable cost dynamics, taxes, and downstream processing bottlenecks.	Differentiates basic cost structures, integrates high level Mine to Mill constraints, and includes realistic assumptions for productivity and costs.	Solves a full value chain stochastic objective function, producing probabilistic estimates that allow for risk-informed decision-making.
2. Uncertainty – How explicit geologic uncertainty is considered	Relies on single-point deterministic estimates. Uncertainty is treated solely as a resource classification problem.	Applies partial orebody simulations to derive uncertainty measures as statistical ranges.	Deploys full orebody simulation including all relevant geometallurgical variables.
3. Realism – Validity of execution assumptions	Uses idealised, unconstrained operational assumptions like theoretical equipment productivity, processing parameters and costs.	Utilises average historical performance and incorporates baseline operating constraints.	Incorporates simulated conditional distributions of mining equipment productivity, processing plant performance, costs and commodity prices.
4. Resilience – Strength of built-in flexibility	Static schedule with no pre-designed alternatives. Operational disruptions force ad hoc, reactive changes like high-grading.	Features informal flexibility through unmonitored operational slack like stockpiles lacking cost-benefit metrics.	Designs bounded optionality into the plan with cost-optimised alternative ore sources, contingency benches or stopes, and clearly defined trigger points.
5. Learning – Efficacy of reconciliation and feedback loop	Fragmented, historical reconciliation using monthly or quarterly lagging metrics. Planning and operations function in functional silos.	Structured tracking with localised feedback loops. Variance between plan and actuals is noted, but long-term planning assumptions remain static.	Tightly integrates real-time short-interval control with continuous Mine to Mill reconciliation loops, frequently updating all relevant models.
6. Sustainability – Treatment of geotechnical, environmental and social parameters	Treats geotechnical, environmental, and social parameters as soft operational targets that are not explicitly considered as part of the plan optimisation.	Geotechnical parameters are treated as hard physical constraints, but environmental and/or social parameters are soft operational targets.	Sustainability as the absolute physical boundary across geotechnical, environmental and social parameters.

To operationalise the scorecard, management can map their asset's performance onto a six-axis radar diagram (Figure 2). This visual representation highlights the gap between the asset's current

operating state, typically heavily deterministic and fragile, and the target strategic state required to insulate the business from its inherent volatility.

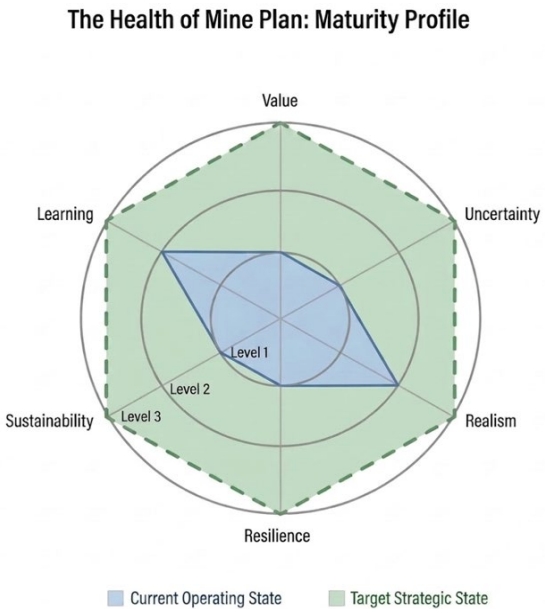


FIG 2 – Diagram representing the gap between an asset's current state and the target.

CONCLUSION – MINE PLANNING AS A ROBUST OPERATING STRATEGY

In practice, the pervasive failure of mine plans is most often a failure of strategic design. When plans are built as deterministic, static schedules, they ignore the inherent variability of the mining system, including geology, equipment performance, processing response, sustainability, and market economics. As this paper has argued, the goal of a modern mine plan is not to eliminate uncertainty, but to create a system that remains resilient in the face of it.

Rather than treating durable economic value, predictable execution, built-in flexibility, and absolute sustainability as isolated targets, successful mining organisations must synthesise these pillars into a unified operating strategy. A robust plan recognises that value is protected only by managing systemic constraints rather than chasing localised efficiencies, and it treats flexibility as a deliberate, engineered investment. Crucially, it respects sustainability parameters as hard physical boundaries rather than negotiable guidelines.

Transitioning to this robust framework requires a fundamental change in both data foundations and organisational governance. Continuous reconciliation of geology, processing, mining productivity and sustainability performance must feed directly into short-interval control loops, narrowing the range of uncertainty for future assumptions. Simultaneously, cross-functional governance must ensure that every level of the organisation, from the mining face to the boardroom, is aligned with the same economic logic and decision rules.

Ultimately, a mine plan is a form of applied strategy. It is the primary tool by which a mining company allocates its scarcest resources, capital and ore reserves, to achieve its strategic goals. As Rumelt (2011) argues, a good strategy identifies the pivotal challenge and provides a coherent response. In mining, that pivotal challenge is uncertainty.

The strongest plans should therefore be understood not as rigid maps, but as robust operating strategies: economically disciplined, operationally credible, sustainability bounded, and flexible enough to sustain value when reality inevitably departs from the plan. The relevant test is not whether a plan eliminates uncertainty, but whether it anticipates plausible deviations and contains decision-ready alternatives long before disruption occurs.

The scorecard presented here is designed to highlight the gaps between an asset's current operating state and the target strategic state required to insulate the business from the uncertainties inherent to the mining value chain. It can be understood as a map that directs operational priorities toward

value maximisation under predictable execution, without ever compromising sustainability and long-term performance.

REFERENCES

Brennan, M J and Schwartz, E S, 1985. Evaluating natural resource investments, *The Journal of Business*, 58(2):135–157.

Dimitrakopoulos, R, 2011. Stochastic optimization for strategic mine planning: A decade of developments, *Journal of Mining Science*, 47(2):138–150.

Goldratt, E M, 1984. *The Goal: A Process of Ongoing Improvement* (North River Press: Great Barrington).

Goodfellow, R and Dimitrakopoulos, R, 2017. Simultaneous Stochastic Optimization of Mining Complexes and Mineral Value Chains, *Mathematical Geosciences,* 49(3):341–360.

Lynch, A J, 1977. *Mineral Crushing and Grinding Circuits: Their Simulation, Optimisation, Design, and Control* (Elsevier Scientific Publishing Company: Amsterdam).

Mintzberg, H, 1994. *The Rise and Fall of Strategic Planning* (Free Press: New York).

Pascale, R T and Athos, A G, 1981. *The Art of Japanese Management: Applications for American Business* (Simon and Schuster: New York).

Rendu, J M, 2017. *Risk Management in Evaluating Mineral Deposits* (Society for Mining, Metallurgy and Exploration: Englewood).

Rumelt, R, 2011. *Good Strategy/Bad Strategy: The Difference and Why It Matter*s (Crown Business: New York).

Senge, P M, 1990. *The Fifth Discipline: The Art and Practice of the Learning Organization* (Doubleday/Currency: New York).

Shaw, W J, Weeks, A, Khosrowshahi, S and Godoy, M, 2013. Reconciliation – delivering on promises, in Proceedings of the 36th APCOM Symposium, Porto Alegre, Brazil.

Taleb, N N, 2012. *Antifragile: Things That Gain from Disorder* (Random House: New York).

Trigeorgis, L, 1996. *Real Options: Managerial Flexibility and Strategy in Resource Allocation* (MIT Press: Cambridge).

Commodity prices have increased significantly in the last two years – have you adjusted your Life Of Asset (LOA) plan yet? Have you adjusted it in the right way?

G Whittle[1]

1. FAusIMM, CEO, Whittle Consulting Pty Ltd, Melbourne Vic 3127.
   Email: gerald@whittleconsulting.com.au

EXTENDED ABSTRACT

Commodity prices (especially gold but also copper and others) have increased dramatically in the last two years – beyond levels mining companies ever contemplated in their previous sensitivity analysis.

If mining companies did nothing other than keep operating the same Life Of Asset (LOA) plan as they had before, they would clearly be making a lot more money at the higher metal prices – they could not help but do so.

Presumably, if a mining company were to update their LOA plans for the new metal prices, they could do even better. Or not!! It all depends on how you do you LOA planning, and which optimisation mechanisms you do or do not apply.

By anecdotal observation (having read hundreds of technical reports) it appears about 80 per cent of the mines worldwide apply what we would call ‘basic’ mine-planning. They optimise the pit, phases and mine-schedule. They use the break-even (marginal) cut-off to decide what material goes to the plant versus the waste dump, with either no stockpiling or just enough stockpiling to balance ore mined to ore processed across the years. Processing calibration and product specifications are fixed.

What happens to the elements of a LOA plan with this style of planning if the metal price goes up?

- Pits get bigger.
- The break-even cut-off grade goes down. Material that was previously regarded as waste is now economic ore.
- Mining rates can come down in the upcoming years to get the same amount of economic material per annum to the plant.
- Life-of-mine (LOM) gets longer due to larger pit and lower cut-off.

Sounds good – or not! Consider the following:

- If you lower the cut-off to the plant, then the average grade processed goes down.
- If you lower the average grade processed, then metal production per annum goes down.

That means you decrease metal production when the metal price goes up and presumably increase metal production when the price goes down. That makes no commercial sense!

We can demonstrate that if metal prices increase it is better to do nothing than to update your LOA plan this way. The trouble is you will not notice this set back as you are so far ahead of the historic plans due to metal price. High metal prices can lead to planning complacency!

A comprehensive approach to LOM planning involves up to 10 Steps (as defined in the Integrated Strategic Planning Seminar):

1. Pit
2. Phases
3. Schedule
4. Cut-off
5. Stockpiling

6. Blending
7. Processing
8. Product Specification
9. Logistics
10. Capital Expenditure

We will show the effect of higher metal prices on all of these mechanisms and prove that the cost of not doing these is far greater at high metal prices than at low.

Beyond naive optimisation, revisited

G A Yeates[1]

1. Principal, Gavin Yeates Consulting, Pty Ltd, San Remo Vic 3925.
   Email: gavin@gavinyeates.com

EXTENDED ABSTRACT

Nearly 20 years ago I co-authored a paper with Peter Monkhouse titled 'Beyond naïve optimisation', which was published in the AusIMM's 2007 *Orebody Modelling and Strategic Mine Planning Spectrum Series, vol 14*. The premise of that paper was that the then general practice of mine planning was 'naïve', and whilst this was often acknowledged, little was being done to address this naivety. In that paper we asserted that the then 'current best practice' in mine planning was to strive to meet the objective function of maximising the net present value (NPV) of a future mining operation. This was achieved by finding the ultimate pit shell, the single mining schedule, shape and sequence of pushbacks, cut-off grades, along with stockpile rules and procedures under a single set of assumptions about orebody tonnes and grade, processing and mining methods and costs, commodity prices, sales volumes in the case of bulk commodities, and discount rates. That paper then argues that this approach is inherently flawed. It points out that the only thing we can be sure of, is that the set of assumptions chosen will be wrong, and that our actions will invariably alter over time as new information becomes available. Yet we continue to naïvely 'optimise' our mine plans assuming this fixed view of the future.

That 2007 paper then goes on to look in detail at two of the major uncertain assumptions, orebody and Price uncertainty. These two key assumptions form the basis on which the 'optimal' mine plans are based. The paper then proposes the use of Real Options as a framework for dealing explicitly with these uncertainties and incorporating these uncertainties into the mine plan such that the resultant mine plan is more sophisticated (less naïve). A simple worked example was presented, to demonstrate how this approach may be used to ultimately generate a more robust plan and valuation for the mine.

This paper revisits the premise of that paper some 20 years later, explores what has changed since 2007 and then answers the question, are we still undertaking naive optimisation?

MINE PLANNING PRACTICE

In the nearly 20 years that have elapsed since this initial paper, the underlying tools and techniques used for mine planning have barely changed. Yes, the software used has been updated, with better interfaces and graphics, and the speed of processing has advanced. The thinking and framing of the mine planning problem has not evolved and as such the underlying algorithms along with the treatment of the assumptions as being fixed, is very much the same as it was in 2007. Industry practice still focuses on final pit limits based on a fixed set of assumptions, then sequencing and scheduling the material within the final pit shell in a manner that maximises NPV given those fixed assumptions and within a set of chosen constraints. The application of these mine planning tools and techniques for everyday use by mine planners either in project studies or for operational plans in both companies and consultancies is very much the same it was in 2007. So yes, we are still undertaking naïve optimisation. Whilst there exists more advanced tools and techniques such as simulation (either conditional or Monte Carlo) or the use of Real Options or range analysis, these remain confined to research or for very niche specialist applications.

KEY SOURCES OF UNCERTAINTY

Whilst the sources of uncertainty remain much the same as they were 20 years ago, there have been developments that tell us more about the magnitude of the uncertainty associated with these key assumptions. Whilst Orebody and Price uncertainty remain the most significant sources of uncertainty, the complexity of mining and processing has risen over the intervening 20 years resulting in increasing levels of uncertainty in the chosen assumptions such as metallurgical recovery and costs, mining costs, recovery and dilution.

Orebody uncertainty

The recent Parker challenge undertaken in conjunction with the AusIMM Orebody Modelling and Resource Estimation conference has provided the industry with quantitative measures of orebody uncertainty, as measured by the between estimator distribution of outcomes. For a given blind data set, nearly 30 estimators undertook their own resource estimate and JORC compliant report. Even though each of these estimates were deemed to be compliant with the JORC Code, the distribution of outcomes that ranged greater than ±200 per cent in contained metal. The generally accepted industry expectation for the uncertainty in an Measured and Indicated resource estimate would be in order of ±25 per cent (Dunham, 2023). The Parker Challenge has demonstrated conclusively that the uncertainty associated with a resource estimate, a key input to a mine plan, is significantly higher than industry expectations (potentially higher by an order of magnitude). It is completely naïve to enter into a mine planning study selecting one and only one of these estimates as 'right' and ignore the distribution or uncertainty associated with that estimate.

Commodity price uncertainty (or volume for bulk commodities)

The last two decades has seen an increase in the volatility of commodity prices. This is evidenced by the approximate doubling of volatility for gold over that time, and we see similar increases for other commodities. This has resulted in mine plans prepared only a few years ago, being shown to have been massively sub-optimal with many of the strategic decisions in hind sight to have been wrong or value destructive. Conventional mine planning practice still uses a single commodity price as input to its mine planning processes. This single price assumption is used in the generation of value within each block, the determination of cut-off, for pit optimisation and scheduling processes, then ultimately in the valuation of the project. The impact of using a single price point is exacerbated by the conventional mine planning process which focuses heavily on optimisation and selection of the final pit shell. The actual mining of this final shell can be decades into the future when price and other assumptions are likely to be vastly different to those chosen today. So a single view of price, based on history and what we know now, is used to make planning decisions that will come into effect decades into the future. Some of these decisions can be irreversible or value destructive and potentially valuable options are often extinguished. It remains completely naïve to continue to use single point estimates for commodity prices that take no account of the uncertainty or variability associated with those assumptions particularly in a world of increasing volatility.

Solutions exist

There are readily available tools that are routinely used in other industries to incorporate uncertainty into planning and operational decision-making. Techniques such as range analysis, probabilistic determination, real options valuation and simulation are all well-established tools to first measure the level of uncertainty associated with an input variable, and secondly to incorporate that uncertainty into planning and decision-making process. Even though these tools exist they are not in routine use in the mining industry.

CONCLUSION

The commonly used industry practice used to plan, evaluate and value mines has barely changed over the last 20 years. The industry is fixated on providing a single point estimate (often to many significant figures) of the 'optimised' NPV for a single set of assumptions. These assumptions include taking a single view of the orebody even when we know that any geological model is uncertain and is based on sparse data and an interpretation that can change. These mine plans are also based on a single view of commodity prices even in a world where the volatility is increasing, as well as a single view of other inputs such as costs, recoveries throughputs etc. To continue to plan and make decisions based on a single set of assumptions for a mine that can have a life measured in decades is naïve. We would all serve the industry better if we moved away from 'optimisation' towards robust or resilient planning. We need to develop plans that are robust and resilient, in a range of future worlds and plans that preserve options for later execution should conditions change. To continue to ignore uncertainty in our planning is naïve.

REFERENCES

Dunham, S, 2023, 27 May. Parker Challenge Thoughts (the first bit) [online], LinkedIn. Available from: <https://www.linkedin.com/pulse/parker-challenge-thoughts-first-bit-scott-dunham/>

Monkhouse, P H L and Yeates, G A, 2007. Beyond naive optimisation, in *Orebody Modelling and Strategic Mine Planning* (ed: R Dimitrakopoulos), Spectrum Series, pp 3–18 (The Australasian Institute of Mining and Metallurgy: Melbourne; and Springer: Cham).

Unlocking value in mining projects – a stochastic mine planning workflow for complex gold-copper leach operations under geological uncertainty

A Prawasono[1] and H Camero[2]

1. Senior Director Mine Planning, Kinross Gold Corporation, Toronto ON, Canada. Email: agung.prawasono@kinross.com
2. Principal Mine Planning Engineer, Kinross Gold Corporation, Denver CO, USA. Email: harold.camero@kinross.com

INTRODUCTION

Complex gold-copper deposits pose non-trivial strategic planning challenges for mine projects: several economic elements and ore types must be optimised at once, cut-off decisions depend on more than one metal and blending and product tonnages respond to reward and cost structures that change over time. Conventional workflows for long-term mine planning attempt to resolve these decisions sequentially from a single estimated block model (Whittle, 1990; Hustrulid, Kuchta and Martin, 2013): pit and phase limits ignore product synergies, cut-off grades rely on metrics such as equivalent grades and deterministic forecasts are exposed to non-compliance because variability is smoothed out and uncertainty is not accounted for (Ravenscroft, 1992; Dimitrakopoulos, Farrelly and Godoy, 2002; Dimitrakopoulos and Ramazan, 2009; Asad and Dimitrakopoulos, 2013).

Shortcomings intensify when considering geological uncertainty, which affects the grade's spatial distribution, grade proportions, and correlations between metals if more than one. Stochastic mine planning replaces the estimated model with geostatistical simulations that reproduce variability and quantify uncertainty of the available data (Goovaerts, 1997; Boucher and Dimitrakopoulos, 2009). Risk analysis then measures, in space and time, the risk attached to a pit geometry and production schedule, surpassing traditional sensitivity studies (Ravenscroft, 1992; Dowd, 1994; Dimitrakopoulos and Ramazan, 2009). The simultaneous stochastic optimisation of mining complexes used in this case study generates the extraction sequence, destinations and material flows concurrently, maximising value and managing risk associated with the production targets of the mining complex (Goodfellow and Dimitrakopoulos, 2016, 2017; Dimitrakopoulos and Lamghari, 2022).

This study presents a complete simultaneous stochastic optimisation workflow applied to a gold-copper heap-leach operation – a multi-mine mining complex (two open pits) with oxide, mixed and sulfide ore and a hard physical limit on leach-pad capacity. The workflow advances through four integrated stages: (1) geostatistical simulations; (2) risk analysis of the deterministic plan; (3) direct block scheduling through the simultaneous stochastic optimisation of mining complexes; and (4) the generation of operationalised mine pushbacks based on stochastic optimisation, closing the gap between an optimal schedule and an executable phase design.

GEOSTATISTICAL SIMULATIONS

A set of conditional simulations of gold grades was generated for each deposit and upscaled to the mine planning block support. Compared with the estimated model, the simulations reproduce the data statistics and expose pronounced smoothing effects. In the present case study, the ordinary kriging model overestimates tonnages above cut-off used while underestimating gold grades above the cut-off.

RISK ANALYSIS OF DETERMINISTIC PLAN

The deterministic schedule (sequence and destinations) was applied to the full set of simulations to build $P_{10}/P_{50}/P_{90}$ risk profiles of mineral production and economic production forecasts. Fixing block destinations from the deterministic mine plan exposed a clear risk of overestimating *in situ* ounces, leach head grades and recoverable gold, concentrated in the late-period sulfides. Fixing the sequence and the gold cut-off grade, allowing destinations to change based on the simulated values, revealed a risk of ore shortage and reduction in produced gold to the leach pad.

Consolidated across the life-of-mine, the deterministic mine plans of the open pits overestimated recoverable gold by 8.8 per cent, translating to a net present value (NPV) overestimation of 24 per cent. The investigation shows that the deterministic plan systematically misclassifies material, with probable waste dilution (waste material going to the ore processing streams) and ore loss (ore material sent to the waste dump).

SIMULTANEOUS STOCHASTIC OPTIMISATION

Using the same simulations of the mineral deposits, a single risk resilient plan was generated through the simultaneous stochastic optimisation of mining complexes, which simultaneously optimises the extraction sequence through direct block scheduling (DBS), together with cut-off policies across all simulated orebody realisations. The model spans 20 mining periods (years) and 22 plant periods to honour the delayed leach recoveries, and respects hard operational and environmental constraints: 15 Mt plant capacity, 35 Mt mining capacity, a maximum of 30 per cent inferred material at the plant, a 10-bench sinking rate and a 100 m minimum mining width.

Relative to the risk-assessed deterministic plan (taken as the baseline), the simultaneous stochastic optimisation unlocked value in material previously classified as economic waste through dynamic cut-off optimisation, delivering more produced gold, more ore with less waste, an additional year of mine life and an expected NPV 35 per cent higher. The optimised cut-off policies vary per period and per material type, reflecting the optimisation ability to adapt to grade uncertainty and blending rather than apply a fixed deterministic threshold at a block scale.

PUSHBACK DESIGN AND FINAL SIMULTANEOUS STOCHASTIC OPTIMISATION WITHIN DESIGNED PHASES

The optimised block schedules of the two mines in the mining complex considered were operationalised into mineable phases for both open pits. The designs incorporated practical access, ramp, switchback and slope criteria. The phases were tied to one another and to the final pit to avoid stranded remnants, keeping the geometry as close as possible to the DBS stochastically optimal solution.

In the final stage, the operationalised pushbacks were re-introduced into the stochastic optimiser to re-schedule the phases under geological uncertainty, producing realistic risk profiles, an operable sequence that is different from the deterministic schedule and a credible NPV. The pushback sequence respected the sinking-rate limits (maximum 10 benches), pushback dependency and advancement (two-bench lag between phases) and minimum mining widths (100 m). Operationalising the phases inevitably reduced the value captured by the unconstrained stochastic DBS solution, especially for deposits located in a hill, which challenging access considerations. The expected NPV of the pushback stochastic optimisation settles at 26 per cent above the deterministic risk-analysis baseline, as expected, yet still a substantial and operationally achievable gain. This quantifies the real cost of operationalisation and confirms that the bulk of the stochastic upside remains after the related transition to an executable plan.

Figure 1 compares the geometry of the deterministic phases and extraction sequence with the phases defined by the DBS simultaneous stochastic optimisation and with the final operational stochastic schedule.

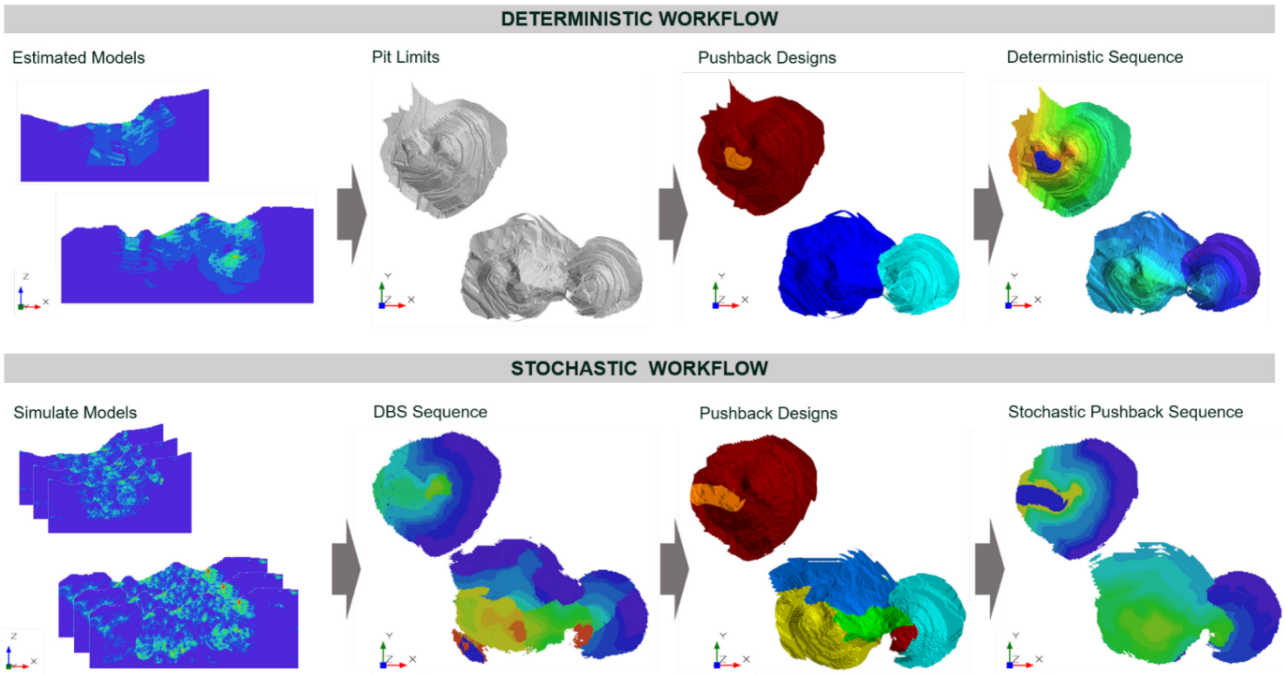


FIG 1 – Deterministic versus stochastic workflow: estimated/simulated models, pit limit/DBS sequence, pushback design and resulting extraction sequences.

Figure 2 tracks recovered gold and NPV along the workflow, both expressed as a percentage of the deterministic risk-analysis P_{50}, showing the value generated by coupling executable phase designs with geostatistical simulations and stochastic optimisation.

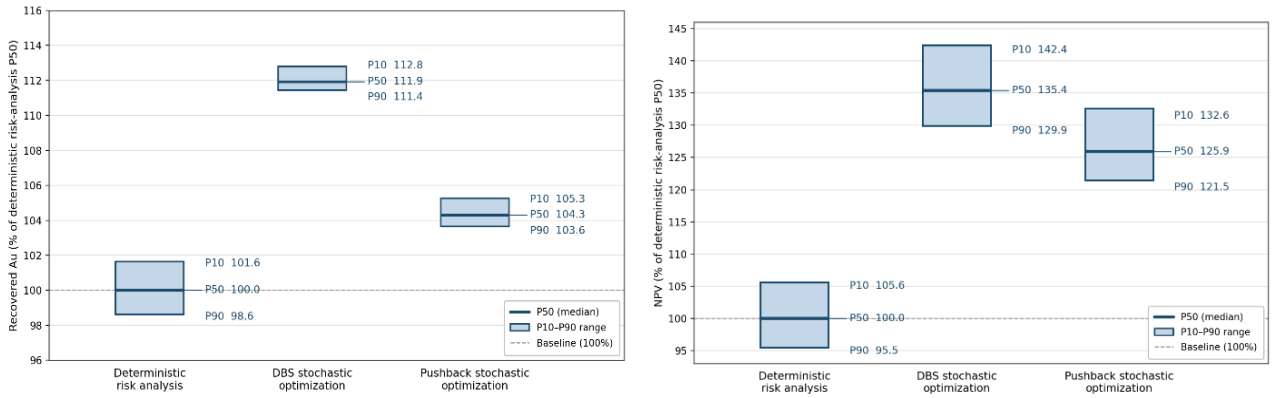


FIG 2 – Evolution of (a) recovered gold; and (b) NPV evolution across the workflow (P_{10}/P_{50}/P_{90}; deterministic risk analysis = 100 per cent). Different vertical scales in (a) and (b).

CONCLUSION

This study demonstrates an end-to-end simultaneous stochastic optimisation workflow for a gold-copper heap-leach mining complex including two open pit mines. Risk analysis shows that the deterministic plan systematically overestimates ore, recoverable and produced gold and overstates NPV by 24 per cent. Simultaneous stochastic optimisation converts this exposure into opportunity, lifting expected NPV by 35 per cent while respecting every hard operational and environmental constraint. Notably, carrying the optimised solution through pushback design and a final pushback stochastic scheduling retains an NPV 26 per cent higher than the deterministic baseline, documenting that the value identified under uncertainty for mining projects is operationally realisable.

REFERENCES

Asad, M W A and Dimitrakopoulos, R, 2013. A heuristic approach to stochastic cut-off grade optimization for open pit mining complexes with multiple processing streams, *Resources Policy*, 38(4):591–597.

Boucher, A and Dimitrakopoulos, R, 2009. Block simulation of multiple correlated variables, *Mathematical Geosciences*, 41(2):215–237.

Dimitrakopoulos, R and Lamghari, A, 2022. Simultaneous stochastic optimization of mining complexes/mineral value chains: An overview of concepts, examples and comparisons, *Int J Mining Reclamation and Environment*, 36:6:443–460. https://doi.org/10.1080/17480930.2022.2065730

Dimitrakopoulos, R and Ramazan, S, 2009. Stochastic integer programming for optimising long term production schedules of open pit mines: methods, application and value of stochastic solutions, *IMM Trans Mining Technology*, 117(4):155–160. https://doi.org/10.1179/174328609X417279

Dimitrakopoulos, R, Farrelly, C T and Godoy, M, 2002. Moving forward from traditional optimization: grade uncertainty and risk effects in open-pit design, *Mining Technology*, 111(1).

Dowd, P, 1994. Risk assessment in reserve estimation and open-pit planning, *International Journal of Rock Mechanics and Mining Sciences and Geomechanics Abstracts,* 32(4):A168.

Goodfellow, R and Dimitrakopoulos, R, 2016. Global optimization of open pit mining complexes with uncertainty, *Applied Soft Computing*, 40:292–304. https://doi.org/10.1016/j.asoc.2015.11.038

Goodfellow, R and Dimitrakopoulos, R, 2017. Simultaneous stochastic optimization of mining complexes and mineral value chains, *Mathematical Geosciences*, 49(3):341–360.

Goovaerts, P, 1997. *Geostatistics for Natural Resources Evaluation* (Oxford University Press: New York).

Hustrulid, W A, Kuchta, M and Martin, R, 2013. *Open Pit Mine Planning and Design* (CRC Press – Taylor and Francis Group: London).

Ravenscroft, P J, 1992. Risk analysis for mine scheduling by conditional simulation, *Transactions of the Institution of Mining and Metallurgy (Section A)*, 101: A104–A108.

Whittle, J, 1990. Open pit optimization, in *Surface Mining* (ed: B A Kennedy), 2nd ed, pp 470–475 (Society for Mining, Metallurgy and Exploration: Littleton).

From asset review to executable growth – transforming strategic mine planning at Vale Base Metals

M Silva[1], J Todd[2], B Gadelha[3] and T O'Neill[4]

1. Technical Director, Mining, Vale Base Metals, Toronto ON M5J 2K2, Canada.
   Email: mario.silva@vale.com
2. Head of Strategic Mine Planning, Vale Base Metals, Toronto ON M5J 2K2, Canada.
   Email: jody.todd@vale.com
3. RDP Manager, Vale Base Metals, Sudbury ON P0M 1N0, Canada.
   Email: bruno.gade@vale.com
4. Non-Executive Director, Woodside, Melbourne Vic 3000, Australia.

ABSTRACT

Copper demand is projected to outrun supply by around 10 Mt/a by 2040, while lead times for new mines have roughly tripled since the 1990s. Vale Base Metals (VBM) responded with a comprehensive review of its global mineral endowment and, critically, of the business planning practices used to convert resource potential into executable options. The review found four gaps: inconsistent valuation across assets, bottom-up capital allocation, no mechanism for comparing fundamentally unlike asset types, and a disconnect between long-range strategic intent and near-term execution. It established a blueprint for change – a move from asset-by-asset evaluation to an integrated, portfolio-level planning system.

Building on that blueprint, VBM rebuilt its Strategic Mine Planning business process around the Resource Development Plan (RDP), an unconstrained, pre-scoping instrument that explores the full potential of the endowment without committing study-stage resources, fully integrated with Life of Asset and Life of Business planning across operating assets, brownfield and greenfield projects, and open pit to underground transitions. At its core is a multi-asset strategic optimisation that evaluates alternative development and sequencing options against a portfolio objective, ranks them on capital intensity and net present value, and screens them for materiality and strategic risk.

This paper presents the strategic rationale, governance model and planning architecture underpinning the transformation. It also sets out what the approach cannot do, including the limitations of capital intensity as a ranking metric and the deliberate choice of deterministic scenarios at the concept stage.

By embedding continuous reassessment, portfolio-level optimisation and structured scenario analysis in core business processes, VBM has strengthened decision quality and capital discipline. The framework underpins a defined pathway to approximately 700 kt per annum (kt/a) of copper by 2035, nearly double 2025 production.

INTRODUCTION

Global copper demand is expected to rise by approximately 50 per cent, from 28 Mt in 2025 to 42 Mt by 2040, driven principally by electrification and the build-out of artificial intelligence data centres, while supply is projected to fall short by around 10 Mt over the same period (S&P Global, 2025; Vale Base Metals, 2026b). The average lead time for new copper mines has reached 17.8 years from discovery to production for those that became operational between 2020 and 2024, nearly three times the lead time of mines commissioned in the 1990s (S&P Global Market Intelligence, 2025). With nickel, cobalt and platinum group metals facing similar dynamics, diversified base metals producers face both an extraordinary opportunity and a complex portfolio planning challenge.

Vale Base Metals (VBM) is a leading global nickel and copper producer with more than 120 years of mining history. It operates 13 mines, five concentrators, two smelters and five refineries across Canada, Brazil, Japan and the United Kingdom, with an operating presence in Indonesia through PT Vale Indonesia Tbk (Vale Base Metals, 2026a). Owned 90 per cent by Vale S.A. and 10 per cent by Manara Minerals Investment Company, VBM's portfolio is anchored by Brazil's Carajás copper district – including Salobo, the world's second-largest iron oxide copper gold (IOCG) deposit – and

by Canadian operations at Sudbury, Thompson and Voisey's Bay, with longer-term growth optionality from the Hu'u porphyry copper gold project in Indonesia.

VBM's copper and nickel endowment supports its ambition to grow copper production from 382 kt in 2025 to approximately 700 kt/a by 2035, while remaining the largest Western Class I nickel supplier and targeting 210–250 kt/a by 2030 (Vale Base Metals, 2026b). Realising this growth requires a strategic mine planning framework capable of evaluating, prioritising and sequencing development options across a diverse, geographically distributed portfolio.

In 2022, Vale S.A. completed the restructuring of its base metals business, establishing VBM as a standalone entity with dedicated governance and capital allocation processes. This structural change created both the opportunity and the imperative for a comprehensive reassessment of strategic planning capability. A series of asset reviews, conducted with the support of experienced industry leaders including T O'Neill, found that VBM's legacy approach – asset-by-asset evaluation, bottom-up capital allocation and limited portfolio-level integration, a pattern long recognised in the strategic planning literature (Smith, 2012) – was insufficient to unlock the full value of the mineral endowment or to support the company's growth ambitions in critical minerals.

In response, VBM introduced the Resource Development Plan (RDP): a portfolio-level instrument providing an unconstrained, full-potential view of the mineral endowment. The concept itself is not new. Several major producers run some form of unconstrained or full-potential portfolio assessment ahead of formal study work, under a variety of names, and the underlying principle – that option generation should be separated from option evaluation, and both from capital commitment – is well established in the strategic planning literature (Smith, 2012; Whittle *et al,* 2018). What is less common, and what this paper argues is the source of the value, is the degree of integration. The RDP at VBM is not a periodic strategic study running alongside the planning calendar; it is a permanent tier of the business planning hierarchy, formally connected to the Life of Asset and Life of Business Plans and, through them, to the annual budget. Together with a refreshed approach to project design, delivery and licensing, it forms the foundation of VBM's trajectory toward being a significant contributor to closing the projected copper supply gap. This paper presents the strategic rationale, governance model and planning architecture underpinning that transformation, and the lessons learned to date.

THE ASSET REVIEW – ESTABLISHING THE STRATEGIC BLUEPRINT

The asset review was initiated as a comprehensive assessment of VBM's global mineral endowment and its readiness to support long-term growth in critical minerals. The review tested both the asset base and the business planning practices used to convert resource potential into executable options. Led with direction from experienced external industry leader T O'Neill, it challenged legacy assumptions, introduced international best practice and helped frame the strategic vision. Its scope covered VBM's operating assets, development projects and exploration targets in Brazil and Canada, together with its international refining assets and strategic joint ventures.

The review identified four critical gaps in the planning approach. First, valuation and prioritisation were inconsistent across assets: different operations used different assumptions, discount rates and planning horizons, which prevented like-for-like comparison. Second, capital followed site requests rather than a structured assessment of where incremental capital created the most value. Third, there was no formal mechanism to weigh value, capital intensity and risk across fundamentally different asset types – large-scale open pit copper operations in Brazil against deep underground polymetallic nickel mines in Canada. Fourth, strategy was disconnected from execution: the Life of Asset Plans (LoAPs) and annual budgets were not systematically informed by, or aligned with, a long-term strategic view. Each of these findings maps directly to a mechanism in the framework that replaced it (Figure 1).

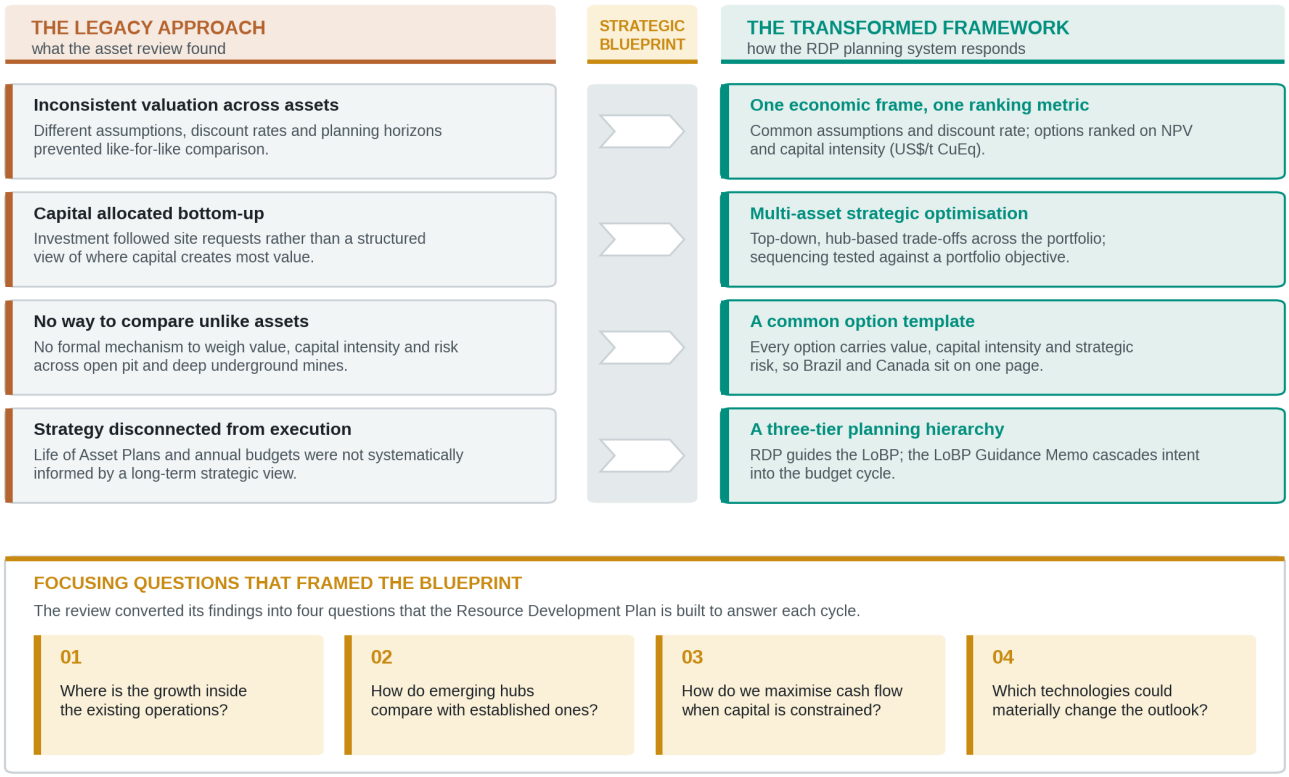


FIG 1 – Findings of the asset review and the corresponding design responses embedded in the Resource Development Plan framework.

The review also produced a set of focusing questions to guide strategic direction: where the key growth opportunities lie within existing operations; how emerging hubs compare with established ones; how to maximise cash flow in a capital-constrained environment; and which technology opportunities could materially change the outlook for specific assets. These questions became the foundation of the strategic blueprint – a directive to move toward an integrated, portfolio-level planning system with the RDP as its central instrument. They are re-asked, and re-answered, in every planning cycle.

THE TRANSFORMATION FRAMEWORK – FROM ASSET-LEVEL TO PORTFOLIO-LEVEL PLANNING

Planning philosophy and hierarchy

The transformation introduced a planning hierarchy of three tiers, designed to create a clear line of sight from long-term strategy to near-term execution. At the apex sits the RDP, which provides strategic direction and portfolio-level guidance. The second tier comprises the LoAPs and Life of Business Plans (LoBPs), which translate strategic guidance into detailed, asset-level and business-unit-level plans. The third tier is the annual budget and execution planning cycle, which operationalises those plans into actionable programmes with defined deliverables, timelines and accountabilities.

Within those three tiers, the business runs six planning stages across distinct horizons, from the 25-year strategic view through to daily execution, each with its own purpose, outputs and cadence (Figure 2). The distinction matters: the tiers describe where authority and integration sit, while the stages describe the work that is actually performed and the artefacts it produces.

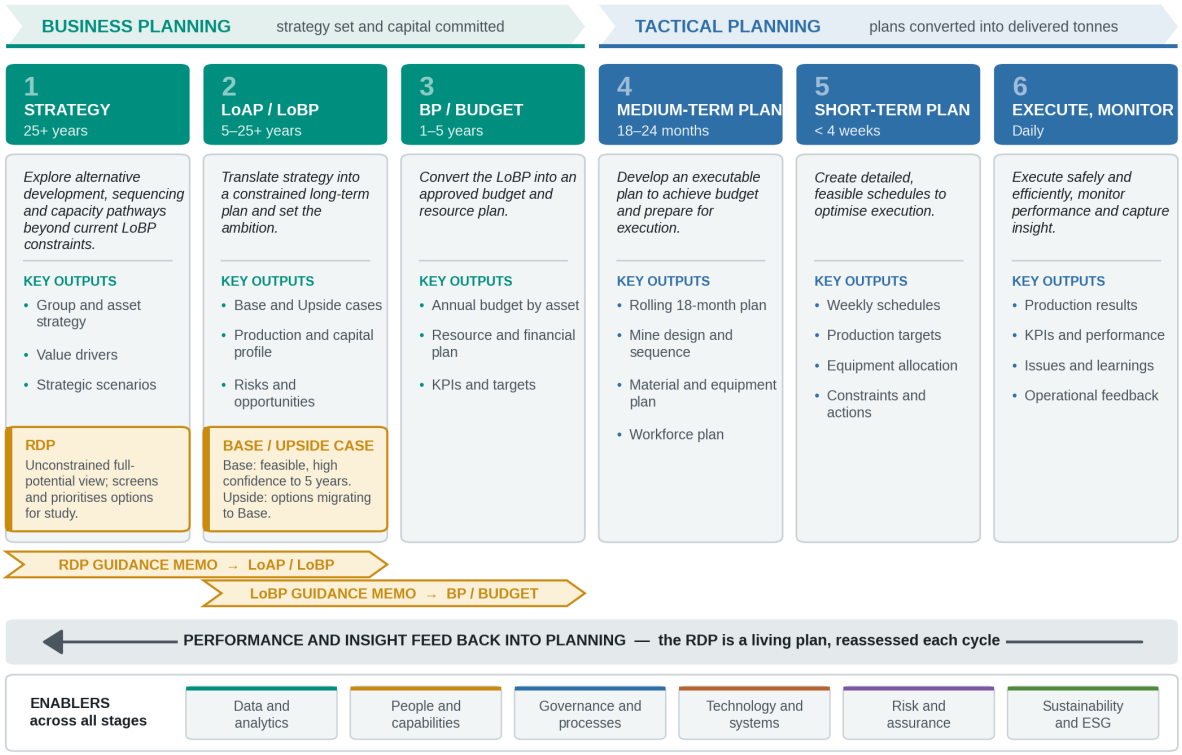


FIG 2 – VBM's integrated planning architecture, linking the 25-year strategic horizon to daily execution through the Resource Development Plan, the Life of Asset and Life of Business Plans, and the annual budget.

A fundamental principle of the framework is that the RDP is a strategic guidance exercise, not a conventional mine planning task. The RDP operates in the pre-scoping and concept stage, where the objective is to explore the full potential of the mineral endowment, identify the most promising development pathways and provide strategic direction, without committing the time and resources required for detailed engineering studies. This allows the organisation to generate, analyse and refine opportunities continuously within a structured framework, accelerating the identification of value-creating options.

The RDP begins not with just a mine plan but with a view of how much endowment there might be. For each deposit, the resource base is expressed as a range rather than a single figure, built up in steps of declining confidence: Mineral Reserves; Measured and Indicated Resources; Inferred Resources; declared exploration targets; and, beyond these, the conceptual endowment implied by the geological footprint, camp-scale analogues and the extent to which mineralisation remains open. Each step carries an explicit basis — drilling density, demonstrated continuity, geophysical response, analogue endowment — so that the confidence attached to any option can be traced back to the material it rests on.

This is the principal reason the RDP can see options that the Life of Business Plan cannot. A plan that must be supportable by Mineral Reserves is, by construction, blind to the part of the endowment that has not yet been drilled to that confidence, and that is precisely the part that determines whether a shaft, a concentrator or a tailings facility has been correctly sized. Infrastructure decisions are irreversible on a scale that mining decisions are not: a pushback can be re-sequenced, a hoisting capacity cannot (not without a lot of effort/impact). The RDP therefore sizes long-lead infrastructure against the endowment range and its associated risk, while committing capital only against the confident end of it.

The resource range input in the RDP – identifies where an incremental tonne of Resource is worth the most: not where the deposit is largest, but where additional drilling would move an option across a materiality threshold, change the sizing of a shared facility, or reverse a sequencing decision between hubs. This converts exploration prioritisation from a geological argument into a portfolio

one, and it is the mechanism by which the RDP directs brownfield drilling programmes toward the outcomes most likely to change a decision.

The framework recognises three mine planning stages, each corresponding to a different level of confidence and detail (Whittle *et al,* 2018): the Unconstrained stage (pre-scoping and concept), where strategic options are explored without operational constraints; the Full Potential stage (FEL-1 or PEA, and FEL-2 or PFS), where promising options are advanced through scoping and pre-feasibility assessment; and the Feasibility stage (FEL-3 or FS), where selected projects are developed to feasibility and approved for execution. This staged approach ensures that planning effort and investment remain proportionate to the maturity and confidence of each opportunity.

The Resource Development Plan

The RDP is designed to fulfil several interconnected purposes:

- provide an unconstrained, full-potential view of every Asset/Hub/Business Unit in the portfolio
- offer a range of strategic growth options for evaluation
- test how resource endowments are incorporated using recent exploration and project data
- identify bottlenecks, constraints and optimal sizing for key infrastructure
- advise on capital allocation and project sequencing
- test innovative technical solutions
- prioritise exploration expenditure
- incorporate environmental, social and governance (ESG) risks as integral planning inputs.

The RDP framework is integrated into VBM's broader business model as a concept horizon for pseudo-projects – opportunities that have not yet entered the formal project pipeline but that warrant strategic evaluation. This creates a structured mechanism for the early identification and assessment of growth opportunities, ensuring that the pipeline of future projects is continuously refreshed and that strategic decisions are informed by the broadest possible view of the endowment.

Every option carried in the RDP is dispositioned as advance, study, retain or defer, and each disposition has a defined destination in the study and delivery pipeline (Figure 3). Options that are retained or deferred are not discarded; they re-enter the following cycle and are re-tested against updated assumptions.

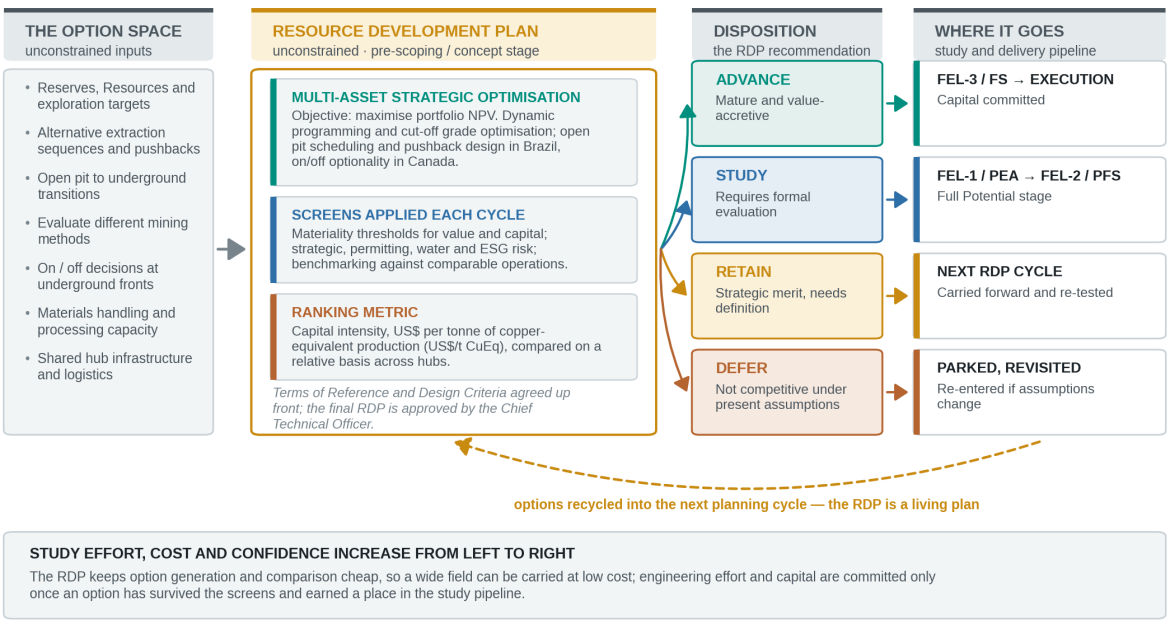


FIG 3 – The Resource Development Plan as an option-generation and screening engine, showing the inputs, the optimisation and screening logic, and the four dispositions applied to each option.

The RDP is intentionally iterative and adaptive. Financial analysis focuses on relative comparisons between options rather than absolute metrics, reflecting the early-stage nature of the exercise. The RDP also identifies strategic risks unique to each proposed option – changes to resource maturity, mining method, geotechnical uncertainty, permitting timelines, water management, community considerations and technology readiness – ensuring that decisions are informed by more than financial metrics alone.

Multi-asset strategic optimisation

At the core of the RDP is a multi-asset strategic optimisation that evaluates alternative development and sequencing scenarios across the portfolio. The objective function is to maximise portfolio net present value (NPV) over a horizon of 25 years and beyond, at a common real discount rate, subject to current and potential hub-level capacity constraints, an indicative capital envelope and known licensing limits. The decision variables are the selection, scale and timing of development options; the resulting outputs are time-phased profiles of metal, operating cost and capital for each option and for the portfolio as a whole.

Making Brazilian open pits and Canadian underground mines commensurate within one problem is the central technical difficulty, and the framework resolves it by deliberately not comparing mine designs. For the Brazilian copper operations, which are predominantly large-scale open pits, the decision set is close to continuous: pushback configuration, extraction sequence, cut-off grade policy, destination and plant feed blend, and the potential transition from open pit to underground. These are addressed with scheduling optimisation and cut-off grade optimisation techniques (Hall, 2019) implemented in optimisation tools available within the industry. For the Canadian nickel and copper operations, which are predominantly underground mines with largely fixed extraction geometry, the decision set collapses to a discrete problem of on/off optionality – when to start, pause, restart or cease production at an individual mining front or operation. Both classes of problem are then expressed in a single common currency: the incremental cash flow and capital signature that each option contributes to the portfolio, evaluated on a common price deck, discount rate, exchange rate and treatment and refining terms. What is compared across the portfolio is therefore not the engineering of the option but its contribution to portfolio value and its claim on capital.

Hub-level constraints enter the problem as shared capacity envelopes rather than as asset-level assumptions. Concentrator capacity, tailings storage, licensed throughput, power supply, and rail and port concentrate logistics are each modelled at the hub node, and options that compete for the same envelope are evaluated as mutually exclusive or as a sequence rather than as independent additions. This is the mechanism through which the RDP identifies infrastructure bottlenecks and optimal sizing, and it is the reason a portfolio-level view can reach conclusions that no asset-level study would reach on its own.

Alongside sequencing, the framework evaluates alternative mining methods – established techniques such as AVOCA, sub level open stoping (SLOS) and newer approaches including sub-level caving or shrinkage and Alimak mining – together with targeted capital investment in materials handling and processing technology, such as in-pit crushing and conveying, trolley assist and coarse particle flotation. Options are optimised jointly within a hub, where shared infrastructure creates genuine interdependence, and compared across hubs at portfolio level. This two-stage decomposition keeps the problem tractable at concept stage while capturing the interactions that matter most to portfolio value.

The consequence of working with relative rather than absolute financial comparisons should be stated plainly. It is what allows the RDP to move quickly across a wide field of options at consistent quality, and it is appropriate to the confidence level of the underlying inputs. It also means that no RDP output is a business case. Absolute value is established through the study stages, and a option's RDP ranking is a claim on study effort, not on capital.

The ranking metric and its limitations

Options are ranked primarily on capital intensity, expressed as United States dollars per tonne of copper-equivalent production (USD/t CuEq). The metric was chosen because it is comparable

across commodities, asset types and jurisdictions, and because capital availability, rather than resource availability, is one constraint on the growth ambition.

Capital intensity is nevertheless a capital efficiency screen, not a measure of value, and ranking on it alone introduces a systematic bias. It favours brownfield over greenfield development, short-life over long-life options, and incremental debottlenecking over the structural expansion that ultimately unlocks a large endowment. Left uncorrected, a portfolio optimised on capital intensity would deliver an efficient decade and an impoverished one after it. The framework therefore applies three lenses to every option rather than one: capital intensity establishes the claim on scarce capital; portfolio NPV establishes the value created and protects long-life, high-capital options that a pure intensity ranking would suppress; and the strategic risk overlay captures what neither number sees – permitting exposure, water and tailings capacity, geotechnical uncertainty, community considerations and technology readiness. An option is advanced only where all three support it, and the reasons for advancing an option that ranks poorly on one lens are recorded explicitly.

Treatment of uncertainty

The RDP is a deterministic, scenario-based process. Uncertainty is addressed through structured scenario analysis and sensitivity testing – commodity prices, discount rates, capital availability and development timing – rather than through stochastic optimisation of the mining complex and its supply chain (Dimitrakopoulos and Lamghari, 2022). This is a deliberate choice, and its rationale is throughput. Simultaneous stochastic optimisation requires geological uncertainty models and a model build effort that is difficult to justify for an option that may be deferred after a single cycle, and the RDP's value lies in carrying a wide field of options cheaply. Stochastic methods are therefore reserved for the Full Potential stage and beyond, where the option set has narrowed and the additional confidence changes a real decision.

The risk this carries is that a deterministic ranking understates the option value embedded in flexibility, particularly for the on/off decisions that dominate the Canadian portfolio, where the ability to pause and restart a mining front is precisely a response to uncertainty. The framework mitigates this by testing the on/off decision set explicitly across scenarios rather than optimising it under a single price path, and by recording flexibility as a qualitative strategic factor where it cannot yet be valued. Extending stochastic treatment further up the planning hierarchy remains an open area of development for the team.

Hub-based portfolio organisation

The portfolio is organised into hub-based structures to facilitate integrated analysis. The Canadian region comprises the Ontario Hub (including Sudbury operations), Thompson and Voisey's Bay. The Brazilian region comprises the Salobo Hub, the Sossego and South Hub, Paulo Afonso, and Alemão. This organisation enables the analysis of shared infrastructure, processing capacity and logistics constraints within each region, while maintaining the ability to compare and optimise across the entire portfolio – particularly as projects mature and exploration advances within a hub.

Optimisation outputs are consolidated into portfolio-level scenarios and visualised through interactive dashboards. These enable rapid scenario comparison and sensitivity analysis, allowing decision-makers to explore the implications of alternative assumptions for commodity prices, discount rates, capital availability and development timing. They also serve as a communication tool that makes complex multi-asset analysis accessible to senior leadership and supports informed, timely decision-making.

Integration with Life of Asset and Life of Business planning

A key design principle is the full integration of the RDP with Life of Asset planning. The LoBP is the primary vehicle through which strategic direction from the RDP is translated into detailed, directional and executable plans, and it now serves as VBM's foundational pillar guiding the budgetary process (Vale Base Metals, 2026b).

LoBPs are constructed with two scenarios. The Base Case represents the most likely development pathway, with high confidence in the first five years, supported by Mineral Reserves and approved or near-approved projects. The Upside Case captures higher-value development options, including

mature initiatives expected to migrate toward the Base Case in subsequent planning cycles, and forms a key focus for corporate teams during the planning cycle. Within the Base Case, one selected scenario is designated the Go-Forward Case, and it is this scenario that becomes the basis of the budget.

The deliverables follow a hierarchical structure: the LoAP at the individual site and business-unit level, and the LoBP at the VBM global level. This ensures consistency and traceability from site-level mine plans through to enterprise-level strategy, while allowing appropriate levels of detail and accountability at each tier. Several business planning principles underpin the integration: timely provision of guidance from the RDP to inform LoBP development; systematic evaluation of multiple scenarios to test robustness; explicit cascading of strategic intent from the RDP through to asset-level plans; optimisation of validation time so that plans are reviewed and approved efficiently; and locking mine plans at the latest practicable date so that the most current information is reflected in the outputs.

GOVERNANCE AND ORGANISATIONAL DESIGN

The transformation required not only new processes and tools but a purposeful redesign of governance structures and organisational roles. The Strategic Mine Planning function sits at the core of the Mines and Geotechnical organisation, reflecting its role as the integrating function that connects geology, geotechnical engineering, mining engineering, water management, processing, infrastructure, finance and corporate strategy into a coherent long-range plan.

The mandate of the Strategic Mine Planning team is fourfold: to define credible, value-optimised long-term plans for each asset and for the portfolio as a whole; to provide transparent decision support to senior leadership; to shape the strategic direction of portfolio growth through the RDP; and to maintain alignment between Mineral Reserves, mine plans, annual budgets and corporate strategy. This last function is particularly important, as it ensures that the Mineral Reserves reported to the market are consistent with the mine plans that underpin them, and that both are consistent with the capital and operating assumptions embedded in the budget.

The RDP is established as the reference view for portfolio prioritisation, infrastructure timing and asset transition decisions. That status is enforced rather than assumed: the RDP is version-controlled and approved by the Chief Technical Officer; the LoBP Guidance Memo derived from it is the mandated input to the budget cycle; and any material divergence between a site plan and the current RDP must be reconciled and documented through the LoBP reconciliation review, rather than resolved locally.

A significant organisational innovation was the establishment of the Copper Growth Forum, a structured, cross-functional platform designed to accelerate the identification, evaluation and advancement of copper growth opportunities within the VBM portfolio including those identified within the Upside Case. The Forum brings together representatives from Technical Services, Projects, Geology, Geotechnical Engineering, Permitting, Finance and Corporate Development. Its objectives are to ensure on-time delivery of the copper growth roadmap, to enable fast and informed decision-making, to avoid duplication of effort across functions, and to share strategic context across a multi-disciplinary team. Its practical effect has been on decision latency: decisions that previously required sequential functional review over many months or planning periods now close within a single monthly forum. The Strategic Mine Planning function plays a central role, framing the strategic questions and providing the analytical foundation for comparative evaluation of growth options.

Process discipline is enforced through the early establishment of Terms of Reference and Design Criteria for each planning cycle. These documents define the scope, assumptions, scenarios and deliverables for the RDP and LoBP processes, and are agreed by all participating functions before detailed work commences.

BRIDGING STRATEGY AND EXECUTION

One of the most persistent challenges in strategic mine planning is bridging the gap between long-term strategy and near-term execution. Plans developed at the strategic level must ultimately be translated into actionable programmes that site-based teams can implement. VBM has addressed

this through a combination of structured information transfer, clear planning definitions and integrated governance.

The primary mechanism is the RDP and LoBP Guidance Memo. Issued at the conclusion of each RDP and LoBP cycle, it communicates the key outcomes of the strategic planning process and translates them into specific guidance for the subsequent LoAP and budget planning cycle. It specifies the agreed Base Case and Upside Case parameters, key milestones and decision points, capital allocation envelopes, and any strategic directives that must be reflected in the budget. It provides sufficient context and direction for site-based teams to develop their annual budgets in alignment with strategic intent, without requiring them to replicate the full strategic analysis.

The Base Case is a feasible plan with high confidence in the first five years, supported by Mineral Reserves and approved or near-approved projects. The Upside Case is a higher-value pathway that incorporates mature initiatives and projects expected to migrate toward the Base Case as they advance through study and approval. Over successive cycles the Upside Case therefore acts as a leading indicator of future Base Case content, creating a natural progression from opportunity identification to committed execution, and defining one of the principal focus areas for the Copper Growth Forum.

VBM's planning process spans a set of interconnected horizons, each with a defined purpose and time frame: Strategy (25 years and beyond, set by the RDP); Life of Asset Plan (0 to 25 years and beyond); Business Plan and Budget (one to five years); Medium-Term Plan (18 to 24 months); and Short-Term Plan (less than four weeks). The explicit definition of these horizons, and the mechanisms that connect them, ensures that strategic decisions made at the RDP level are progressively translated into increasingly detailed and time-specific plans, with feedback loops that allow execution experience to inform future strategic assessment.

OUTCOMES AND VALUE DELIVERED

The first RDP, completed in 2024, was a landmark: for the first time, all Canadian and Brazilian assets were incorporated into a single integrated strategic study, enabling direct comparison and portfolio-level optimisation across geographies, commodities and mining methods – a capability that had not previously existed within the organisation. The 2025 RDP extended the review to encompass all business planning horizons from the RDP through to five-year planning, closing the remaining gaps between long-term strategy and near-term execution and creating a continuous planning system from the 25-year strategic horizon through to the annual budget.

The most consequential outcome is a credible, capital-efficient pathway to nearly double copper production over the next decade. By evaluating, prioritising and sequencing development options within an integrated portfolio-level framework, VBM has translated a substantial mineral endowment into a defined roadmap from 382 kt of copper in 2025 to approximately 700 kt/a by 2035, anchored by a pipeline of low-capital-intensity, high-return projects across the Brazilian and Canadian portfolios (Figure 4). In parallel, the framework has enabled optimisation of the polymetallic Canadian assets to strengthen VBM's position as the largest Western supplier of Class I nickel, with production targeted to reach 210–250 kt/a by 2030.

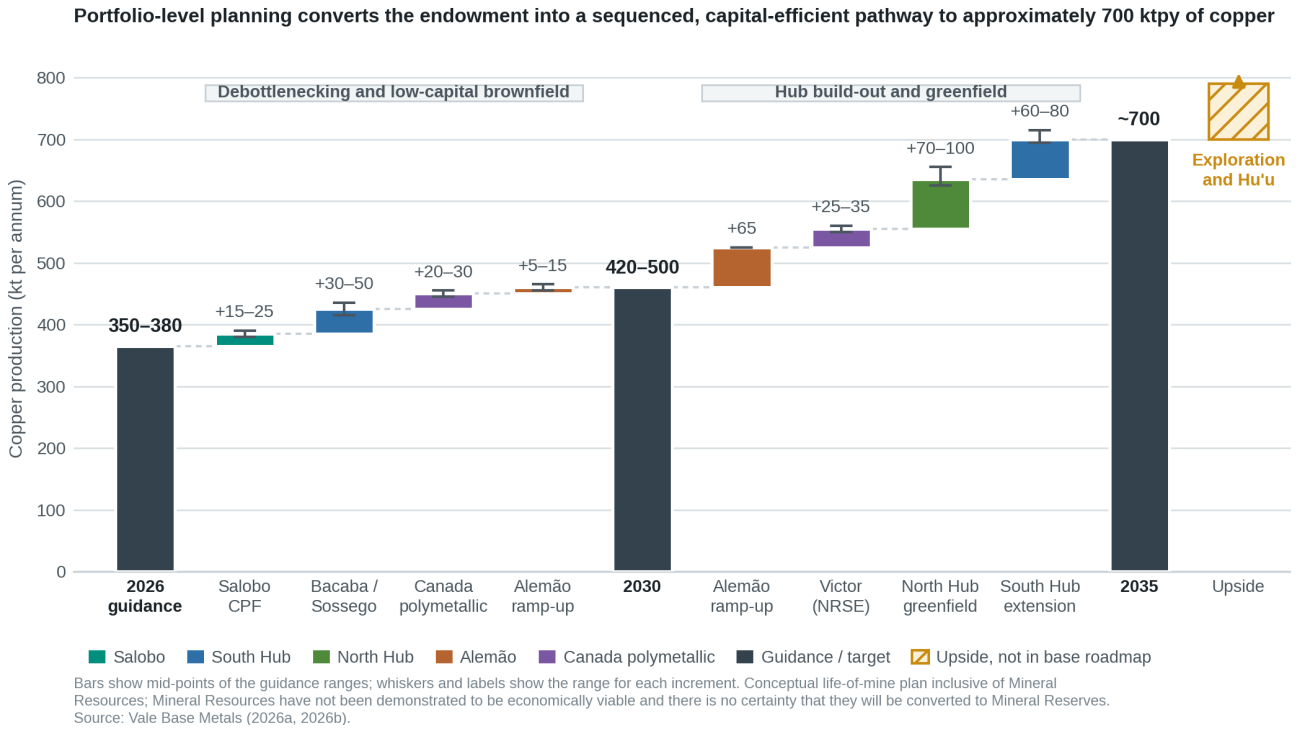


FIG 4 – Copper production bridge from 2026 guidance to approximately 700 kt/a by 2035, showing the contribution of each hub. Increments are guidance mid-points; whiskers show the stated range.

The effect on Mineral Reserves and Resources is also notable, although it should be read with care. As of the 2025 reporting cycle, VBM's copper Mineral Reserves and Measured and Indicated Resources increased by a net 3.7 per cent year-on-year, and nickel Mineral Reserves and Measured and Indicated Resources by a net 9.4 per cent (Vale Base Metals, 2026a). These movements reflect drilling, price and reconciliation effects as well as planning; the contribution attributable to the Strategic Mine Planning team is the portion arising from optimised mine plans and revised economic parameters, and from the prioritisation of brownfield drilling programmes through the RDP toward the opportunities most likely to influence portfolio-level outcomes.

LESSONS LEARNED AND RECOMMENDATIONS

Five lessons from the transformation may be of value to other companies facing similar challenges. The first two are the least obvious and, in the authors' experience, the most important.

First, establish materiality thresholds before generating options, not after. The unconstrained nature of the RDP means that a very large number of options can be generated, and not all of them are material at portfolio level. Clear thresholds for both value uplift and capital expenditure prevent the process becoming unfocused and keep analytical effort directed at the options that can genuinely influence portfolio decisions. VBM set these thresholds in the Terms of Reference for each cycle, agreed across functions before work commenced; the discipline of agreeing them in advance is what makes it possible to decline an interesting but immaterial option without relitigating it.

Second, benchmark relentlessly, because unconstrained analysis invites optimism. An exercise that explores full potential is structurally vulnerable to embedding pessimistic/optimistic assumptions about grades, recoveries, throughputs, costs and schedules, and equally vulnerable to unexamined conservatism inherited from legacy plans. Systematic benchmarking against comparable operations and independent industry data sets – project pipeline and capital intensity data (Wood Mackenzie, 2025) – combined with structured challenge from experienced peers, is what keeps RDP outputs credible when they are tested through subsequent study stages.

Third, start with the strategic question, not the software. There is a natural tendency in the industry to treat optimisation tools, scheduling software and digital platforms as the drivers of improvement. They are important enablers, but the value lies in defining the right strategic questions, establishing a coherent planning framework and building the organisational discipline to execute it consistently.

At VBM the focusing questions from the asset review preceded any tool selection, and they still frame each cycle. The tools follow the framework, not the reverse.

Fourth, portfolio thinking is a cultural change and requires both sponsorship and structured cross-functional integration. Moving from ‘my asset’ to ‘our portfolio’ is a human challenge as much as a technical one. It requires sustained leadership engagement – the VBM transformation was championed by the Mining Director and the Chief Technical Officer, whose active engagement provided the authority, resources and organisational mandate to drive change – and it requires structured opportunities for geology, geotechnical engineering, processing, permitting and finance to contribute, challenge and validate at every stage. Strategic mine plans developed in isolation from those functions do not survive contact with reality.

Fifth, treat the RDP as a living document and manage the balance between speed and rigour explicitly. Static plans lose relevance quickly, so the process must embed continuous reassessment, with mechanisms for updating assumptions, incorporating new information and revisiting recommendations as conditions change. That speed must not come at the expense of analytical credibility: clear standards for the level of detail, the quality of inputs and the documentation of assumptions are what maintain stakeholder trust in a process that deliberately operates at low confidence.

CONCLUSIONS

Vale Base Metals has moved from a traditional, asset-by-asset planning paradigm to an integrated, portfolio-level system centred on the Resource Development Plan. The RDP provides strategic direction, the Life of Asset and Life of Business Plans translate that direction into executable plans, and the budget process operationalises them – with structured governance, cross-functional integration and dedicated forums ensuring the system operates with discipline and accountability. The result is a credible, capital-efficient pathway to nearly double copper production by 2035, and a planning process capable of identifying, well in advance, which decisions actually govern the endowment – including those, such as exploration priorities, tailings capacity and permitting pathways, that are not mining decisions at all.

The principles are not unique to VBM’s circumstances. Portfolio-level planning, structured scenario analysis, integration of strategy with execution and disciplined governance are transferable to any multi-asset, multi-commodity company seeking to optimise the development of its mineral endowment. In an era in which the copper supply gap is projected to reach 10 Mt/a by 2040, the ability to systematically identify, sequence and deliver capital-efficient growth options will be a defining source of competitive advantage.

One question remains open, and the authors would welcome dialogue on it. The RDP’s speed depends on deterministic analysis, yet the decisions it most influences – on/off optionality, the timing of infrastructure, the sequencing of an expansion against a permitting pathway – are precisely those whose value lies in flexibility under uncertainty. How much stochastic treatment can be pushed up the planning hierarchy before it destroys the throughput that makes the unconstrained stage worth having is, in the authors’ view, the most important unresolved question in portfolio-level strategic mine planning.

ACKNOWLEDGEMENTS

The authors acknowledge the contributions of the broader Strategic Mine Planning team at Vale Base Metals and the cross-functional teams who have supported the development and implementation of the planning framework described in this paper. The authors also acknowledge the support of Vale Base Metals’ senior leadership in championing the transformation and providing the mandate and resources necessary for its success, and the approval of Vale Base Metals to publish this work.

REFERENCES

Dimitrakopoulos, R and Lamghari, A, 2022. Simultaneous stochastic optimization of mining complexes and their supply chains, *International Journal of Mining, Reclamation and Environment*, 36(4):247–270.

Hall, B E, 2019. *Cut-off Grades and Optimising the Strategic Mine Plan*, Spectrum Series 20 (The Australasian Institute of Mining and Metallurgy: Melbourne).

S&P Global Market Intelligence, 2025. From 6 years to 18 years: the increasing trend of mine lead times, 11 April 2025 (S&P Global Market Intelligence: New York).

S&P Global, 2025. Copper in the Age of AI: Challenges of Electrification (S&P Global: New York).

Smith, G L, 2012. Strategic long term planning in mining, *Journal of the Southern African Institute of Mining and Metallurgy*, 112(9):761–774.

Vale Base Metals, 2026a. 2025 Asset Handbook (Vale Base Metals: Toronto), 30 March 2026.

Vale Base Metals, 2026b. Vale Day 2026 Presentation (Vale Base Metals: Toronto).

Whittle, D, Whittle, J, Wharton, C and Hall, G, 2018. *Strategic Mine Planning*, 8th edition (Gemcom Software International Inc.).

Wood Mackenzie, 2025. Copper Project Database, Q4 2025 (Wood Mackenzie).

Holistic optimisation integrating agriculture into mining for sustainable value

B King[1]

1. MAusIMM, CEO, COMET Strategy Pty Ltd, Cleveland Qld 4163.
   Email: brett.king@cometstrategy.com

ABSTRACT

This paper presents a proof-of-concept framework for integrating agricultural land use into long-term mine planning. Conventional mine optimisation treats rehabilitation primarily as a closure activity and generally excludes the productive value that land may generate while mining continues. The proposed Holistic Optimisation framework extends net present value to a Holistic Present Value that incorporates mining, agricultural, environmental and social cash flows within a common schedule. Dynamic programming was selected because it can represent nonlinear and discontinuous relationships among material movement, available surface area, soil regeneration, agricultural production and closure outcomes. The demonstration model calculates agricultural area as a waste dump develops, estimates soil depth with a logistic growth function, introduces crops and livestock after minimum soil-depth thresholds are reached, and values continuing agricultural cash flows as a perpetuity. Two schedules were compared: a Base Case without agricultural value and a Soil Case that incorporated grass, laying hens and sheep. The Soil Case accelerated early material movement to make agricultural land available sooner. Agricultural value began after nine years and reached full production in year 31. At full production, annual agricultural cash flow was approximately $7.85 million. The modelled Holistic Present Value increased by $172 million, or 1.8 per cent, relative to the Base Case, and mine life increased by 1.32 years. These results show that incorporating productive post-mining land use can change an economically optimised mining schedule, even when agricultural cash flows are small relative to mining cash flows. The values are illustrative and depend on site-specific geology, climate, soil, water, markets, rehabilitation methods and operating constraints. Further work should validate the agricultural parameters with local expertise and test the framework at an operating mine.

INTRODUCTION

Mining and agriculture are foundational industries that provide essential resources for modern life. Mining supplies raw materials such as coal, iron and copper for equipment, vehicles, electronics and infrastructure, while agriculture supplies food and materials such as cotton and timber. The distinction between materials that are grown and those that are mined illustrates the complementary role of the two industries.

Traditional mine optimisation models often overlook agriculture. Existing software typically addresses geology, logistics, environmental constraints and financial considerations but rarely includes agricultural potential. Integrating agriculture into mine planning could create additional revenue and allow mine planners to evaluate environmental and community outcomes throughout the project life. Mine optimisation algorithms should also reflect environmental, social and governance (ESG) concerns, including sustainable food supplies, access to clean water and soil regeneration.

Mine optimisation software selection

Current mine optimisation software is well suited to managing geological constraints, mining processes, logistics and environmental impacts. However, these systems typically do not account for agricultural potential, limiting their ability to evaluate productive land use during and after mining.

This framework cannot be implemented directly in every mine-planning system. Much of the available software uses linear or mixed-integer linear programming, which may require simplified representations of nonlinear agricultural area and value relationships. Dynamic programming can represent these relationships more flexibly, as described by Bellman (1957), Lane (1988) and King

(1999), and as implemented in the COMET Optimal Scheduler. This approach enables the model to combine agricultural and mine optimisation relationships.

The selected mine optimisation software already integrates geology, mining operations, logistics, marketing, and financial and socio-environmental constraints that evolve over a mining project's life cycle, often 10–100 years. It did not, however, originally include agricultural area and growth functions. The modifications extend the optimisation horizon to 300–500 years or more, allowing a broader, multigenerational view of land use. In the demonstration software, a 500-year project can produce a solution within hours and, depending on project complexity, sometimes within minutes. This additional complexity remains within practical time frames for strategic scenario analysis.

Study objective

This study develops an integrated mine optimisation model that includes agricultural components and can rapidly evaluate numerous scenarios. Net present value (NPV) is traditionally used to evaluate mining and processing operations based on costs and revenues over a project's life. A broader assessment must also consider agricultural and environmental effects, together with regulatory and community constraints.

To assess the full value of a project, this study introduces the Holistic Optimisation (HO) framework, which evaluates Holistic Present Value (HPV). HPV extends NPV by including environmental and social costs and revenues and discounting them to present value. This supports a broader economic assessment consistent with established corporate-finance principles (Brealey *et al,* 2022).

METHOD

Part 1 – Surface area calculation

Mining operations traditionally focus on the mass and volume of material moved because these drive costs. Agriculture is primarily concerned with surface area and the availability of resources such as water and sunlight. A dynamic agricultural function was therefore developed to calculate the surface area available for agriculture as mining progresses. The function adjusts to mine progression and surface topography and can be refined to include more detailed relationships.

Part 2 – Soil quality estimation

Once surface area is calculated, soil quality must be evaluated because it is central to regenerative agriculture. Introductory sources on soil regeneration include *Holistic Management* (Savory and Butterfield, 2016), *A Soil Owner's Manual* (Stika, 2016) and *Dirt to Soil* (Brown, 2018).

The soil-quality functions consider soil composition, contamination and land-rehabilitation strategies. Local experts may be needed to define appropriate parameters for each site and for distinct areas within it. The functions use the mining schedule, including waste-dump and tailings activities, to determine when soil regeneration can begin.

To model the initial change in soil depth, a logistic growth model was used. It represents growth that begins slowly, accelerates and then slows as it approaches a maximum.

Characteristics of logistic growth

1. Growth is initially slow because the starting soil depth is small.
2. Growth accelerates as time progresses.
3. Growth slows as soil depth approaches its maximum.
4. The resulting curve is sigmoidal: slow initially, faster in the middle and level near the upper limit.

The logistic growth function is:

$$S(t) = S_max/[1 + \exp(-k(t - t_0))] \quad (1)$$

where:

S(t) = soil depth at time t (mm)

S_max = maximum soil depth (mm)

k = growth rate

t_0 = time at which the growth rate is highest (years)

The proof of concept assumes an initial soil depth of zero that approaches a plateau of 30 mm after approximately 10 years (S_max = 30, t_0 = 5 and k = 2.1).

Figure 1 illustrates the assumed relationship between soil depth and time.

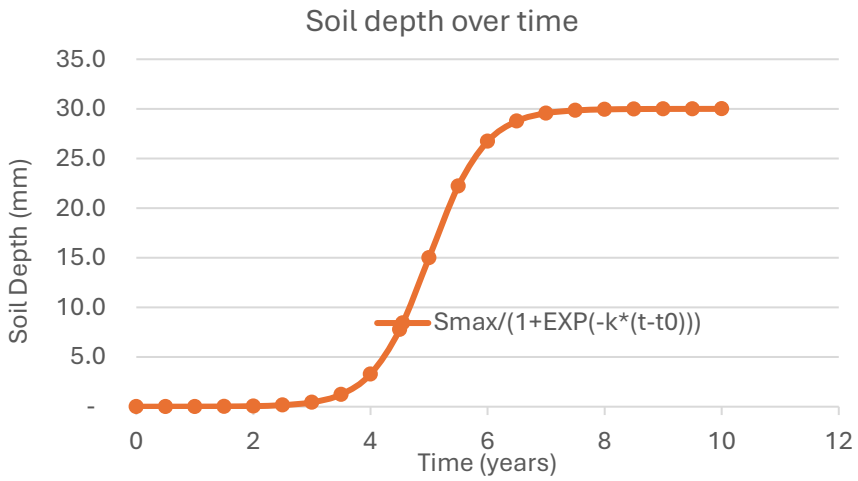


FIG 1 – Soil depth over time.

The formulation should be refined using expertise specific to the site and its earthworks. For this proof of concept, the model approximates soil depth over time. Dynamic programming is well suited to these nonlinear relationships, whereas linear and mixed-integer linear programming formulations may require simplification.

Part 3 – Agricultural value calculation

To integrate agricultural activities into the mine optimisation model, another function was developed to estimate their financial value. The function uses the following components:

- Surface area and soil quality determine the potential scale of agricultural activities.
- Crop yields, livestock prices, market conditions and related metrics determine agricultural value.
- The resulting values allow agriculture to be evaluated as part of the strategic mine plan.

As soil depth increases, plant root growth, nutrient availability, soil biota and water-holding capacity generally improve. Deeper regenerated soil supports crops and other plants, which in turn support livestock and wildlife. Plants and animals both contribute to soil regeneration; vegetation alone may not restore the full ecosystem (Savory and Butterfield, 2016). Soil organisms, including fungi, bacteria, worms and insects, help cycle nutrients and support plant growth. Native plants and animals can further contribute to soil regeneration and a viable long-term ecosystem.

Separate value functions were developed for plant and animal contributions. Both can be represented by the same basic equation:

$$V_A = D \times Y \times P + C \quad (2)$$

where:

V_A = annual value of the agricultural product ($/m^2)

D = soil depth (mm)

Y = yield or quantity of agricultural products (crop yield or animals/m^2)

P = net annual value per product ($/product)

C = annual cost per unit area ($/$m^2$), expressed as a negative value under the COMET convention. It represents the initial cost at zero soil depth and may include crop irrigation and water supplies for animals

Once calculated, the function can be applied to each area designated for agriculture using that area's soil depth. If several products are produced, their values can be combined.

Some agricultural products require a minimum soil-depth threshold before production can begin. For example, larger animals may require a minimum soil depth (D_min) and corresponding vegetation production. This condition can be expressed as:

$$\text{If } D \geq D_min, \text{ then } V_A = D \times Y \times P + C; \text{ otherwise, } V_A = 0 \quad (3)$$

For reporting, both the quantity and value of each product are recorded. Any number of agricultural products can be assigned to areas designated for agriculture, and each area may suit different products. For example, a tailings storage facility may support different plants and animals from a waste dump, and waste materials may differ in properties such as clay content or acid-generation potential. The selected plants and animals may also change as the soil regenerates.

This formulation is a starting point. Crop-production costs may decline as soil depth increases and plants reseed. Animal rotation is beyond the scope of this paper but is important to regenerative management (Savory and Butterfield, 2016). Non-harvested species, including birds, lizards and insects, also contribute to the ecosystem. Native plants and animals should be considered where appropriate.

A benefit of dynamic programming algorithms is that they can represent realistic, complex relationships that may be difficult to express in conventional linear programming models.

Part 4 – Closure impact assessment

This part assesses the long-term impact of agricultural activities on mine closure and subsequent land reclamation.

- The function ideally models agricultural activities continuing indefinitely, even after mining operations cease.
- It assesses the implications for closure plans, land rehabilitation, and post-mining site management.
- Stakeholders gain insights into the long-term environmental and socio-economic sustainability of the project.

The present value of perpetual annual costs and revenues is calculated as:

$$PV = A/r \quad (4)$$

where:

A is the perpetual annual cash flow

r is the discount rate expressed as a decimal (eg 10 per cent = 0.1)

Rather than using a separate custom function, the model calculates this value in Part 3 and adds it to the schedule once the soil is fully developed. This reduces the number of inputs and the risk of omitting a required value.

Traditional closure costs

Modern mine-closure costs include dismantling the processing plant and mine site infrastructure and, traditionally, revegetating the remaining areas. Under the Holistic Optimisation framework, many

land-rehabilitation costs are incurred while mining and processing remain active so that agricultural land becomes available at the optimal time.

Proof-of-concept analysis cases

A simple mining schedule was used to demonstrate how agricultural factors could be integrated and to assess their effects on the mining schedule and project value.

Two cases were analysed:

1. Base Case: a standard mine optimisation model without agricultural elements.
2. Soil Case: a model incorporating the four stages outlined above.

Both models were optimised using the COMET Optimal Scheduler.

Agricultural case parameters

One waste dump was allocated to agriculture. It had a top width of 1000 m, a maximum height of 100 m, a top length of 1000 m and slope angles of 20°. Results are reported in hectares using 0.0001 ha/m^2.

Figure 2 shows the cash flow configuration parameters used for agriculture and mining.

Stream Name	Supply Location	Cut-off	Attribute	Factor	Capacity	Values		Material Flow	Custom Function
			.			Activity	Period Timing		
Mining	Total Period		Mass		100	-1.50			
Mill Feed	Custom COG	LP(Base)	Mass		PV(Mill Capacity)				
HG SP Separation	Fixed COG	4	Mass					Add to HG SP	
Waste Split	Fixed COG	-10	Mass					Add to Waste	
Send to Agricultural Dump	Sent to Waste		Mass	400000					Agriculture1AreaOfRectangularDump(WasteDumpForAgriculture, OverflowForAgriculture, 1000, 100, 1000, 20, 0.000247105)
Genreate Soil	Empty		Mass						Agriculture2SoilDepth(WasteDumpForAgriculture, 0, 30, 10, 0.5)
Grass hay (t) value	Empty		Mass					Add to Grass	Agriculture3LifeValue(WasteDumpForAgriculture, 0, 0.0666666666666667, 0,-0.0001)
Chicken egg value	Empty		Mass					Add to Chicken Eggs	Agriculture3LifeValue(WasteDumpForAgriculture, 5, 1, 0.000408,0)
Sheep value	Empty		Mass					Add to Sheep	Agriculture3LifeValue(WasteDumpForAgriculture, 10, 0.166666666666667, 0.0002,0)
Mining Period Value	Total Period		Mass		1		-55.00		ApplyAboveLimit()
Ore Haul	Sent to Ore		CrusherTrk			-250.00			
SP Haul	Sent to HG SP		DumpTrk			-250.00			
Waste Haul	Sent to Waste		DumpTrk			-250.00			
Ore Processing	Sent to Ore		Mass			-5.00			
Copper Revenue	Sent to Ore		CuRec	2204.62		PV(Copper Price)			
Gold Revenue	Sent to Ore		AuRec			PV(Gold Price)			

FIG 2 – Cash flow configuration parameters for agriculture and mining.

Table 1 summarises the agricultural production parameters used in the proof of concept.

TABLE 1

Agricultural production parameters.

Agricultural system	Minimum soil depth (mm)	Annual stocking or yield at full depth	Valued output	Annual value per unit ($)	Annual area cost ($/ha)
Grass	0	4.94 t/ha	Grass forage	0	-247
Laying hens	5	74.1 hens/ha	Egg production per hen	408	0
Sheep	10	12.4 sheep/ha	Sheep	200	0

To improve realism, livestock and crops were introduced only after the required soil depth was reached. Chickens required 5 mm of soil and sheep required 10 mm. These assumptions demonstrate how agricultural constraints can be integrated with a mining schedule.

Grass was not harvested as hay; instead, it supported chickens and, later, sheep. The model also assigned ongoing grass-related costs, such as irrigation, as a function of area.

Laying hens were used for egg production rather than meat. The model assumed that each hen produced 340 medium-sized eggs per annum and generated annual revenue of $408.

Animal numbers were scaled by soil depth. For example, the full stocking rate was 12.4 sheep/ha at a soil depth of 30 mm. At a depth of 10 mm across 18.2 ha, the model therefore allowed approximately 75 sheep (12.4 × 18.2 × 10/30).

These assumptions will vary between projects and may also vary within a project. This proof of concept is intended to demonstrate how the parameters can be integrated and how they may influence the mining schedule and project value.

RESULTS AND DISCUSSION

Holistic present value overview

The HPV of each scenario depends heavily on the geological resource, climate and other site-specific factors. The results should therefore be interpreted within the assumptions of this proof of concept.

Table 2 compares the Base and Soil cases.

TABLE 2

Summary of the Base and Soil cases.

Case	Holistic present value ($ million)	Mine life (years)	Final value ($ million)
Base	9695	23.68	-10
Soil	9867	25.00	68.5
Change	+172 (1.8%)	+1.32 (+5.6%)	+78.5

The Soil Case increased HPV from $9.695 billion to $9.867 billion, an improvement of $172 million or 1.8 per cent. It generated no agricultural value during the first nine years and did not reach full production until year 31. At the end of the mine life, the $10 million terminal cost was incurred while agricultural products were generating annual cash flows of $7.85 million. At a 10 per cent discount rate, the perpetuity value of those cash flows was $78.5 million. Discounting that terminal value back 31 years produced a present value of $4.1 million. The remaining increase arose from changes to the optimised mining schedule.

Mining schedules

The mining schedules were compared using total material moved and agricultural area available. The Base Case had no incentive to develop the waste dump for agriculture.

Figure 3 compares total material movement and the creation of agricultural area in the Base and Soil cases.

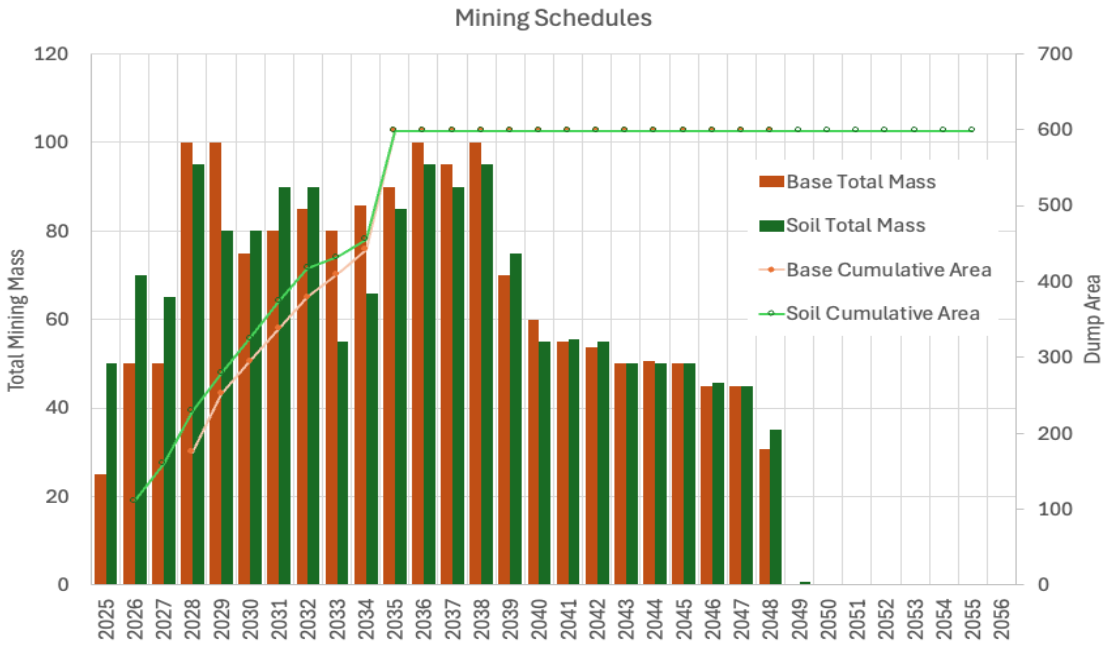


FIG 3 – Total material movement (left axis; Mt) and agricultural area created (right axis; acres).

The Soil Case used faster mining rates in the early years, creating agricultural area sooner. Both schedules finished at a similar time, with slightly lower mining rates after the dump was completed. The optimised cut-off grades were slightly higher during the faster early mining period, while lower-grade material supported slower mining later in the schedule.

The relationship between material movement and available agricultural area is neither linear nor continuous. This makes it difficult to represent accurately in linear programming models but well suited to the dynamic programming approach used here.

Agricultural area availability

Figure 4 shows total material moved and the areas available for grass, laying hens and sheep.

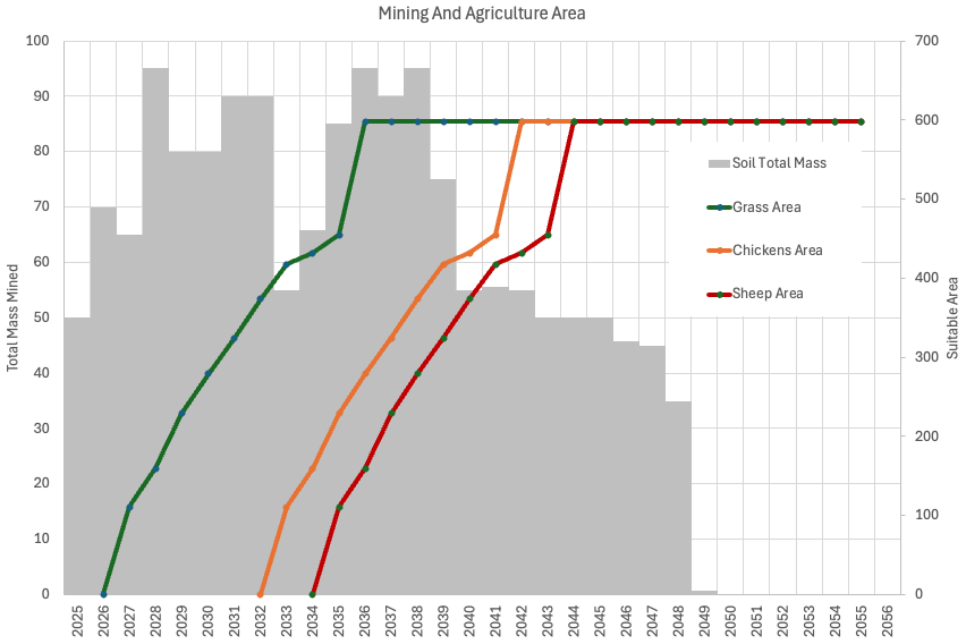


FIG 4 – Total material movement (left axis; Mt) and agricultural area availability (right axis; acres).

The offsets in available area arise because chickens require soil at least 5 mm deep and sheep require 10 mm before they can be introduced.

Mining and agricultural cash flows

Figure 5 shows annual cash flows from mining and the three agricultural products. The source-model chart retains agricultural-series labels ending in ‘Area’; in this figure, those series represent annual cash flow streams rather than physical area.

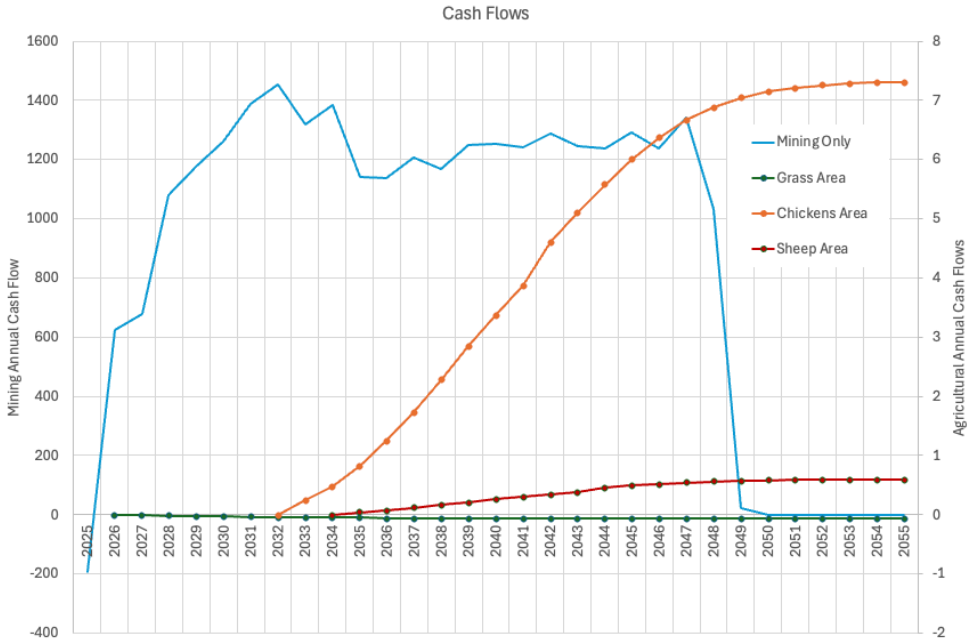


FIG 5 – Mining cash flow (left axis; $M) and agricultural cash flows (right axis; $M).

Several features are evident in the mining and agricultural cash flows:

- Mining cash flows are significantly higher than agricultural cash flows and are therefore plotted on a separate axis. Agricultural values continue indefinitely, while mining cash flows end.
- Grass incurs costs but generates no direct revenue; livestock products, such as chicken eggs and sheep, drive agricultural income.
- Eggs generated the most revenue among the three agricultural products. At full production, the model included approximately 18 000 laying hens across nearly 243 ha. An operation of that scale would require substantial infrastructure to handle and market the eggs. Alternatively, some laying hens could be replaced with meat chickens. These estimates require detailed analysis and validated parameters; the resulting value could be considerably lower than assumed in this case.
- The model produced nearly 3000 sheep for sale each year. The assumed net price of $200 per sheep reflected meat production only and excluded wool revenue.
- The framework assumes that soil regeneration requires a mix of plant and animal species rather than grass or a single animal species alone.
- Native plants and animals were not included in this analysis, but the model could be extended to include them.

CONCLUSIONS

This proof-of-concept study demonstrates the potential for integrating agriculture into mining operations to increase project value and improve sustainability. The modelled value increase was modest at 1.8 per cent, while agriculture also sustained food production and supported land regeneration over the longer term. These results show how a first-principles Holistic Optimisation framework can assess mining and agriculture together. The framework provides a method for incorporating environmental and social considerations into value-based mine planning while evaluating a sustainable agricultural business that could benefit local communities.

FURTHER WORK

Further research should refine the agricultural models and test them in different mining environments, including more complex relationships between agriculture and mining, soil-reclamation techniques and environmental effects. Expanding the Holistic Optimisation framework to include additional social, environmental and economic factors would allow it to be tested across a wider range of mine settings. A practical next step is to identify a suitable area at an operating mine for a case study and publish the findings.

ACKNOWLEDGEMENTS

The author thanks Moises Alvarez (Codelco), Pamela Castello (Escondida) and Orlando Leiva (Anglo American) for their enthusiastic support and encouragement.

DISCLOSURE STATEMENT

The COMET Optimal Scheduler features used in the case study are available as part of the standard licence.

FUNDING DETAILS

COMET Strategy Pty Ltd funded the research.

DATA AVAILABILITY STATEMENT

The underlying data are confidential and are not publicly available.

DATA DEPOSITION

No data were deposited.

REFERENCES

Bellman, R E, 1957. *Dynamic Programming* (Princeton University Press: Princeton).

Brealey, R A, Myers, S C, Allen, F and Edmans, A, 2022. *Principles of Corporate Finance*, 14th edn (McGraw-Hill Education: New York).

Brown, G, 2018. *Dirt to Soil: One Family's Journey into Regenerative Agriculture* (Chelsea Green Publishing: White River Junction).

King, B, 1999. Schedule optimisation of large complex mining operations, in *Proceedings 28th International Symposium on Application of Computers and Operations Research in the Mineral Industry*, pp 749–762 (Colorado School of Mines: Denver).

Lane, K F, 1988. *The Economic Definition of Ore: Cut-off Grades in Theory and Practice* (Mining Journal Books Ltd: London).

Savory, A and Butterfield, J, 2016. *Holistic Management: A Commonsense Revolution to Restore Our Environment*, 3rd edn (Island Press: Washington, DC).

Stika, J, 2016. *A Soil Owner's Manual: How to Restore and Maintain Soil Health* (CreateSpace Independent Publishing Platform).

Planning for uncertainty – the next frontier in strategic mine planning

O Leiva[1]

1. VP, Head of Strategic Planning, Anglo American, London, UK.
   Email: orlando.leiva@angloamerican.com

ABSTRACT

The evolution of strategic mine planning has been shaped by the software tools and data available to planners. It has progressed from simplified representations of mine operations to increasingly integrated, dynamic and computationally intensive approaches for the full value chain. This evolution has created significant value for the industry, enabled through increasingly sophisticated schedule optimisation techniques. However, in an increasingly dynamic and volatile world, static analyses with deterministic inputs are no longer fit for purpose. Uncertainty in inputs must be directly represented, analysed and used to inform strategic decision-making.

The development of modern mine planning can be seen as a series of step changes. The Lerchs-Grossmann algorithm provided a mathematical foundation for pit optimisation in the 1960s, allowing planners to move beyond manual design logic and identify economically optimal pit limits under defined assumptions. In later decades, planning advanced from 2D to 3D modelling, improving the representation of orebody continuity and mining constraints. Commercial optimisation tools such as Whittle 4X, incorporated the work of Kenneth Lane for cut-off grade optimisation, testing staged development options. More recently, integrated strategic optimisation software's such as COMET and Minemax support multi-policy schedule optimisation; supporting sophisticated value-based trade-offs with much greater speed and consistency. Cloud architecture and scalable compute now allow thousands of cases to be generated in hours. However, this introduces a new challenge; more scenarios does not automatically create better decisions. Without a clear framework for ranking options or understanding confidence and variability in inputs, additional cases can create noise rather than insight.

A key limitation of traditional planning approaches is that they remain deterministic at their core. A single 'optimal' result can be found for production capacity, extraction sequence, cut-off policies and cash flow. However, the inherent uncertainty in orebody knowledge, processing performance and future economic conditions is masked, with plans often appearing more precise than the underlying data justifies. By compressing a range of possible futures into one assumed outcome, deterministic plans are fragile and often optimistic; not able to respond to dynamic changes in key inputs. Strategic mine planning must therefore evolve from deterministic scenario comparison to probabilistic decision support.

Recent collaboration between Anglo American and KPI COSMO has explored how stochastic planning methods and confidence frameworks can support capital decision-making and development pathway selection across mine assets. Stochastic planning supported a deeper understanding of orebody uncertainty and production risk to enable targeted data collection to improve confidence. Development pathways apply a structured assessment of value, confidence and risk to explicitly highlight uncertainty in decision-making. In this context, the most valuable plan is not necessarily the one that maximises value in a single deterministic case, but the one that performs well across uncertainty while supporting transparent, risk-informed choices.

The mine plan of the future will combine automation, probabilistic modelling and human judgement in a continuous planning loop. Anglo American has already started exploring this direction through technologies such as AI-assisted core logging, active hyperspectral sensing for real-time ore characterisation, digital twins for plant optimisation and automated process control, alongside faster resource model and planning updates. These capabilities point towards a future where orebody knowledge is captured continuously with real-time model updates and mine plan refreshes, enabling strategic options to be tested as soon as new information is available. Technology provides the tools to improve the scale, speed and complexity of analyses; giving experienced planners the ability to focus on true value-adding strategic optimisation and value-based decision-making.

The industry has progressed from simple 2D representations to optimising integrated life-of-asset strategies. The next frontier is probabilistic optimisation; using uncertainty as a critical input to robust decision-making. In an uncertain world, the best mine plans will not be those that appear most precise. They will be those that are transparent about confidence, resilient across possible futures and capable of adapting as knowledge improves.

Unlocking managing resource uncertainty in a large mining organisation

A Jewbali[1]

1. Global Head, Resource Management, Newmont Corporation, Denver CO 80237, USA.
   Email: arja.jewbali@newmont.com

ABSTRACT

Mining companies make high-value, long-term decisions using resource models constructed from incomplete data, geological interpretation, statistical assumptions, and imperfect knowledge of the orebody. These models form the foundation for reserve estimation, mine design, production scheduling, capital allocation, and operational guidance. Yet conventional planning processes commonly reduce this uncertainty to a single estimated model, a single mine schedule, and a single production forecast. The result is an appearance of precision that may not adequately represent the range of outcomes that could occur when the deposit is mined.

Resource uncertainty does not arise from one source. It develops through the combined effects of sampling and data quality, geological interpretation, assumptions about mineralisation continuity, statistical treatment, estimation parameters, scale, mining selectivity, dilution, recovery, and plan execution. These uncertainties interact and propagate through the mining value chain. By the time a resource model is converted into a mine plan, uncertainty may have been transformed, compounded, or obscured by subsequent assumptions. In large mining organisations, this challenge is amplified by multiple operations, different orebodies, competing planning horizons, time-bound annual processes, and the use of models developed at different spatial and temporal scales.

A fundamental issue is that average grades do not necessarily produce average mining outcomes. Mine planning involves non-linear decisions, including cut-off selection, destination assignment, stope definition, pit design, sequencing, and production constraints. Therefore, optimising a plan using an averaged resource model is not equivalent to evaluating optimised outcomes across the full range of plausible geological scenarios. Averaging can smooth the local variability that drives mining decisions, conceal downside exposure, and remove potential value opportunities. The problem is consequently not only whether the resource estimate is technically defensible, but whether the planning process uses that estimate in a way that supports robust decisions.

Conditional simulation and related probabilistic methods provide an alternative representation of the deposit. Rather than producing one ‘best’ estimate, simulation generates multiple plausible realisations that collectively describe geological and grade uncertainty. These realisations allow practitioners to evaluate the range, probability, and timing of potential mining outcomes before the orebody is exposed. However, the current application of simulation often stops short of fully integrating uncertainty into planning. A common approach is to run multiple realisations through a mine plan derived from the estimated model. This provides valuable risk profiles for tonnes, grade, metal and cash flow, but it mainly tests the vulnerability of an already selected plan. It does not necessarily identify the plan that performs most reliably across uncertain geological outcomes.

Stochastic mine planning addresses this missing link by evaluating planning decisions across many possible resource models. The objective shifts from finding a plan that is optimal for one assumed version of the orebody to identifying decisions that remain robust and economically resilient across a range of plausible realities. This distinction is especially important in underground mining, where alternative geological realisations can produce different stope shapes, sequences, grades, and economic outcomes, while operational flexibility may be constrained. The practical question is no longer simply, ‘Which plan is optimal?’ It becomes, ‘Which plan can tolerate being wrong about the orebody?’

Moving from deterministic to stochastic planning requires more than improved algorithms. It requires integration across long-, medium-, and short-term planning; commercial tools capable of supporting both open pit and underground applications; efficient updating of probabilistic models; and clear rules defining where probabilistic information improves decisions and where deterministic outputs remain

necessary for execution. It also requires case studies from operating mines that demonstrate whether stochastic approaches produce more reliable plans, better management of downside risk, and decisions that remain actionable for operational teams.

The most significant barriers may therefore be organisational rather than mathematical. Resource modellers, mine planners, operational teams, researchers, software developers, and decision-makers must develop a shared understanding of uncertainty and how it should influence a decision. Capability building must extend beyond generating simulations to interpreting probability, comparing trade-offs, communicating risk, and selecting practical responses. The way forward is a coordinated transformation of people, processes, technology, governance, and mindset. Success will not be defined by producing more models. It will be defined by whether uncertainty is translated into better, more resilient resource decisions and mine plans that are less likely to fail when reality differs from the average.

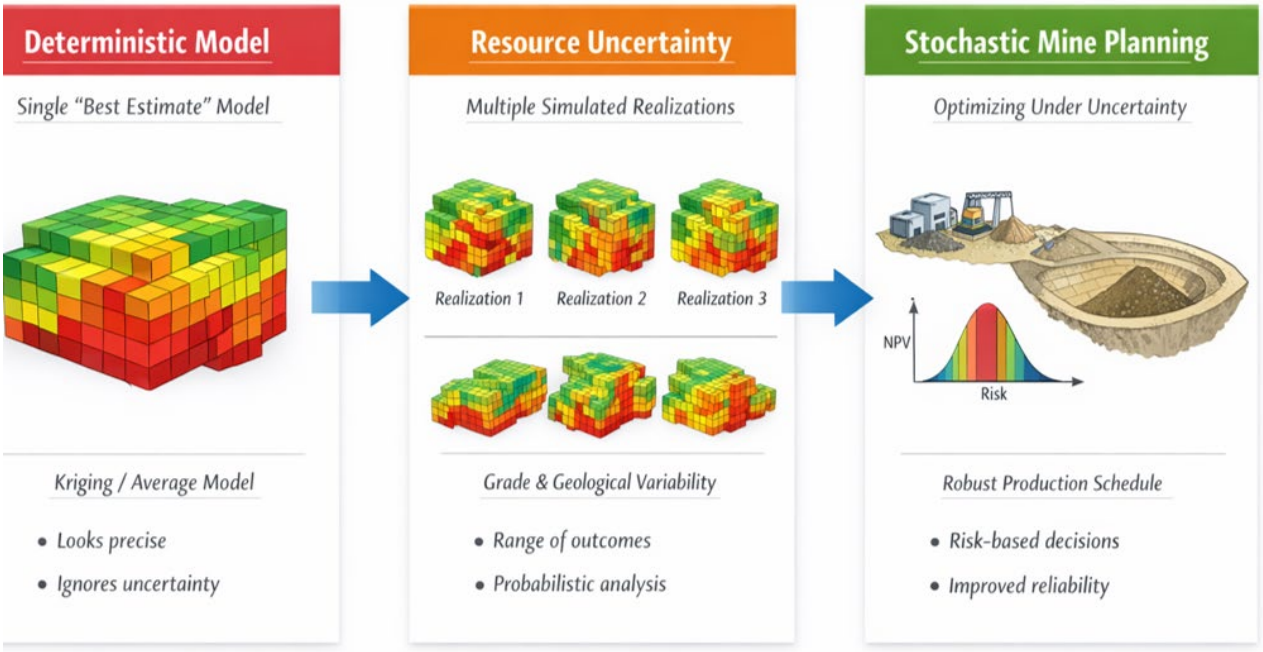


FIG 1 – Conceptual progression from deterministic mine planning, based on a single estimated resource model, to stochastic mine planning, in which decisions are evaluated.

AI and machine learning for mine operations

Reinforcement learning-based stochastic optimisation of short-term production scheduling and shovel allocation incorporating truck dispatching policies and environmental constraints

V Balboa-Espinoza[1,3] and R Dimitrakopoulos[2,3]

1. Research Assistant, COSMO – Stochastic Mine Planning Laboratory, McGill University, Montreal QC H3A 0E8, Canada. Email: victor.balboa-espinoza@mail.mcgill.ca
2. Professor, COSMO – Stochastic Mine Planning Laboratory, McGill University, Montreal QC H3A 0E8, Canada. Email: roussos.dimitrakopoulos@mcgill.ca
3. GERAD – Group for Research in Decision Analysis, Montreal QC H3T 2A7, Canada.

ABSTRACT

Short-term production scheduling and fleet management are complex optimisation problems that must actively manage ore material and equipment performance variability by coordinating material extraction, shovel allocation, truck fleet usage, and downstream processing decisions. Although recent advances have integrated these decisions with greenhouse gas emission considerations, a gap remains in assessing whether strategic long-term decisions can be operationally executed under the stochastic dynamics of daily mine operations. Existing approaches often fail to account for localised ore material variability and simplify complex fleet interactions, making it difficult to verify if strict production and environmental targets remain feasible during short-term execution.

To address this gap, the present work proposes a stochastic optimisation framework for short-term production and shovel allocation that integrates truck dispatching policies and environmental impact indicator estimations through a discrete-event simulator. The framework uses the fleet configuration and extraction targets defined by the long-term schedule and operationalises them through a dynamic simulation engine that explicitly accounts for geological and equipment uncertainty and truck dispatching policies. Within the optimisation environment, a reinforcement learning agent continuously improves an initial solution, which is then evaluated by a discrete-event simulator that captures endogenous interactions between trucks, shovels, and the processing plant. Tested at a copper mining complex, the framework is evaluated across different carbon footprint penalties and multiple truck dispatching policies to determine their effect on material extraction schedules, shovel allocation and fleet performance.

INTRODUCTION

A mining complex is an integrated business system where spatially distributed tasks, decisions, and locations interact to generate and deliver sellable products to specific customers or markets (Hoerger, Seymour and Hoffman, 1999; Stone *et al,* 2007; Whittle, 2007; Whittle, 2010; Dimitrakopoulos, 2011). From a global perspective, the mining industry is a fundamental supplier of materials for infrastructure, electronics, and the energy technologies required for decarbonisation (Calderon *et al,* 2020; De Donno, 2024; Kamran, Raugei and Hutchinson, 2023). As global demand for these minerals grows, the sector faces an increasing need to optimise operations that balance economic outcomes with environmental impacts (Laurence, 2011; Moran *et al,* 2014). Currently, mining companies are responsible for an estimated 8–10 per cent of global carbon emissions (Huo *et al,* 2023; Yokoi, Watari and Motoshita, 2022), making the integration of greenhouse gas (GHG) emission reductions into mine production planning a critical requirement. In open pit operations, nearly 30 per cent of total energy consumption and up to 50 per cent of GHG emissions are attributed to diesel-reliant truck-and-shovel operations, while the comminution stage accounts for another 36–50 per cent of energy use (Ballantyne and Powell, 2014; Jeswiet and Szekeres, 2016; Nikbin *et al,* 2025). Despite its operational and environmental significance, equipment operation is conventionally optimised independently from mine production scheduling, resulting in inefficiencies and higher GHG emissions.

To drive a mining complex toward optimal conditions, mathematical optimisation aims to maximise the net present value (NPV) while ensuring compliance with production targets. However, this process is complicated by the existence of strategic (long-term) and operational (short-term)

planning horizons that are traditionally managed independently. Conventionally, these horizons have relied on deterministic models based on a single estimated orebody model (Cressie, 1990; David, 1977; Journel and Huijbregts, 1978; Krige, 1951). Because mineral exploration data is inherently sparse, ignoring geological uncertainty consistently leads to unrealistic forecasts and sub-optimal designs (Ravenscroft, 1992; Dowd, 1994; Dimitrakopoulos, Farrelly and Godoy, 2002). While stochastic optimisation frameworks have advanced the field by simulating equiprobable orebody scenarios to manage risk across multiple planning components (Both and Dimitrakopoulos, 2020; Dimitrakopoulos, Farrelly and Godoy, 2002; Dimitrakopoulos and Ramazan, 2004; Goodfellow and Dimitrakopoulos, 2016), aligning macroscopic life-of-mine plans with short-term operational execution remains a significant challenge.

Realising strategic economic and environmental forecasts relies entirely on the execution of the short-term operational horizon. Yet most existing approaches fail to acknowledge the endogenous feedback between planning decisions and high-frequency system dynamics. Short-term scheduling must explicitly account for reactive operational policies, such as continuous equipment dispatching, which inherently alter equipment performance variability (Ashtiani *et al,* 2025; de Carvalho and Dimitrakopoulos, 2021; Upadhyay and Askari-Nasab, 2019). When dispatching policies are excluded from the overarching optimisation due to their computational complexity, the resulting models fail to capture the hidden carbon costs of daily operations, such as equipment idling and queuing. Furthermore, while recent studies have incorporated environmental impacts and carbon costing into mathematical optimisation models (Badiozamani *et al,* 2024; Frieire Guimaraes and Dimitrakopoulos, 2026a, 2026b; Levinson and Dimitrakopoulos, 2020, 2024; Pell *et al,* 2019; Rahnema, Amirmoeini and Moradi Afrapoli, 2023; Xu *et al,* 2021), these constraints have primarily been applied to long-term scheduling. This leaves a critical gap in validating whether strict operational and environmental plans can be realised when the stochastic dynamics of daily fleet interactions are simplified or ignored.

To overcome the limitations of current approaches, this paper introduces an integrated stochastic optimisation framework designed for short-term production scheduling and shovel allocation. Unlike previous short-term planning formulations that assign extraction decisions directly to fixed time periods (Balboa-Espinoza and Dimitrakopoulos, 2026; Both and Dimitrakopoulos, 2020; de Carvalho and Dimitrakopoulos, 2024; Kumar and Dimitrakopoulos, 2021; Matamoros and Dimitrakopoulos, 2016; Quigley and Dimitrakopoulos, 2020), the proposed model optimises an invariant and scenario-independent extraction sequence where the actual execution timing emerges dynamically. In this framework, endogenous operational variability acts as a response variable driven entirely by the underlying geological and equipment uncertainty. Rather than treating planning and dispatching separately, the proposed methodology embeds operational truck dispatching policies and environmental impact estimations directly within the optimisation loop. Utilising the fleet configuration and extraction targets dictated by the long-term schedule, the system explicitly accounts for geological and equipment uncertainty through a dynamic simulation engine. At the core of this architecture, a reinforcement learning agent acts as the search mechanism, iteratively refining an initial solution. The operational viability of these refined solutions is then actively evaluated by a discrete-event simulator (DES), which captures the complex, endogenous interactions between the haulage fleet, loading equipment, and downstream processing.

The remainder of this document is organised as follows. Section 2 details the proposed methodology, including the mathematical formulation of the event-based two-stage stochastic optimisation problem and the new proposed reinforcement learning-based hyper-heuristic framework. Section 3 introduces the copper mining complex case study used to validate the approach. Section 4 presents the computational experiments and comparative analyses, evaluating the performance of different reinforcement learning (RL) agents, the impact of carbon footprint penalties, and the effects of varying truck-dispatching policies. Finally, Section 5 summarises the main conclusions drawn from this work.

METHODOLOGY

Mathematical formulation

Decision variables and dynamic sets definition

The proposed framework models short-term scheduling and fleet management as an event-based, two-stage stochastic optimisation problem. First-stage decisions determine scenario-independent execution sequences for the shovels and fleet usage levels for each truck type across the planning periods. The first-stage decision variables are defined as follows:

- $x_{l,j,b} \in \{0,1\}$: equals 1 if digline $b \in \mathbb{B}$ is assigned to the j-th execution position of shovel $l \in \mathbb{L}$, or 0 otherwise.
- $\tau_m^t \in \mathbb{Z}^+$: represent the number of trucks of type $m \in \mathbb{M}$ scheduled in period $t \in \mathbb{T}$.
- $z_{b,i} \in \{0,1\}$: equals 1 if digline $b \in \mathbb{B}$ is assigned to destination $i \in \mathbb{I}$, and 0 otherwise.

Unlike traditional models that pre-assign diglines to fixed time periods, the proposed formulation optimises only the extraction sequence. The actual execution period for each digline emerges dynamically during the discrete-event simulation, driven by equipment availability, truck dispatching logic, processing-plant events, and the realised stochastic scenarios.

The second-stage decisions are associated with the physical execution mapped by the discrete-event simulator based on geological uncertainty ($\mathbb{S}_G$) and equipment uncertainty ($\mathbb{S}_E$). For a joint scenario realisation $s \in \mathbb{S}_G \cup \mathbb{S}_E$ in period $t \in \mathbb{T}$, these variables include:

- $y_{i,k,t,s} \in [0,1]$: the processing-stream decision, representing the proportion of material sent from location $i \in \mathbb{I}$ to destination $k \in \mathbb{K}$.
- $v_{p,i,t,s} \geq 0$ and $v_{h,i,t,s} \geq 0$: the realised values of primary attribute $p \in \mathbb{P}$ and hereditary attribute $h \in \mathbb{H}$, respectively.
- $d^+_{\varepsilon,n,t,s} \geq 0$ and $d^-_{\varepsilon,i,t,s} \geq 0$: continuous variables representing the surplus and shortage, respectively, of primary or hereditary attributes $\varepsilon \in \mathbb{P} \cup \mathbb{H}$ at location $n \in (\mathbb{I} \cup \mathbb{K}) \backslash (\mathbb{I} \cap \mathbb{K})$.

The objective of the event-based two-stage stochastic optimisation model is to maximise the expected economic value while minimising deviations from required production limits across destinations, periods, and scenarios, as well as the environmental footprint associated with fleet operations. Therefore, the model integrates strategic extraction and destination decisions with operational uncertainty captured through the simulation of equipment performance, dispatching policies, and processing-plant events.

To represent the operational dynamics of the system, a global set of discrete operational events, Ψ_s, is defined for each joint scenario realisation s. This set is decomposed into the following subsets during period $t \in \mathbb{T}$:

- $\mathcal{M}^S_{l,t,s} \subset \Psi_s$: movement events executed by shovel $l \in \mathbb{L}$.
- $\mathcal{I}^S_{l,t,s} \subset \Psi_s$: idle events executed by shovel $\in \mathbb{L}$.
- $\mathcal{M}^T_{m,t,s} \subset \Psi_s$: movement events executed by trucks of type $m \in \mathbb{M}$
- $\mathcal{I}^T_{m,t,s} \subset \Psi_s$: idle events executed by trucks of type $m \in \mathbb{M}$.

Each of these event sets is dynamic and depends on the actual operation of the simulated mining complex. In particular, they depend on the truck-dispatching policy, $\Pi(\cdot)$, which allocates variable trucks per period to destinations during the simulation. Moreover, the discrete-event simulator incorporates processing-plant events, such as full-capacity usage, flow redirections, and upstream equipment pauses. These events indirectly affect fleet operations by increasing truck queuing times at specific destinations and modifying the timing of material extraction and transportation.

Objective function

The proposed event-based two-stage stochastic optimisation model evaluates the economic and operational performance of a sequence-based extraction plan through a discrete-event simulator. Therefore, production timing, equipment interactions, queuing, idle times, and processing-plant disruptions are captured endogenously for each joint geological and equipment scenario realisation.

The objective function is formulated as:

$$\max \frac{1}{|\mathbb{S}|}\sum_{s\in\mathbb{S}}\sum_{t\in\mathbb{T}}\sum_{h\in\mathbb{H}}\sum_{i\in\mathbb{I}\cup\mathbb{K}} p_{h,i,t}v_{h,i,t,s}$$

$$-\frac{1}{|\mathbb{S}|}\sum_{s\in\mathbb{S}}\sum_{t\in\mathbb{T}}\sum_{\varepsilon\in\mathbb{P}\cup\mathbb{H}}\sum_{n\in(\mathbb{I}\cup\mathbb{K})\setminus(\mathbb{I}\cap\mathbb{K})}\left(c^{+}_{\varepsilon,t}d^{+}_{\varepsilon,n,t,s}+c^{-}_{\varepsilon,t}d^{-}_{\varepsilon,n,t,s}\right)$$

$$-\frac{1}{|\mathbb{S}|}\sum_{s\in\mathbb{S}}\sum_{t\in\mathbb{T}}\left[\sum_{l\in\mathbb{L}}\sum_{\psi^{s}\in\mathcal{M}^{S}_{l,t,s}} C^{S,mov}_{l,\psi^{s}}(\Pi(\tau^{\mathrm{t}}_{\mathrm{m}}))+\sum_{m\in\mathbb{M}}\sum_{\psi^{T}\in\mathcal{M}^{T}_{m,t,s}} C^{T,mov}_{m,\psi^{t}}(\Pi(\tau^{\mathrm{t}}_{\mathrm{m}}))\right]$$

$$-\frac{1}{|\mathbb{S}|}\sum_{s\in\mathbb{S}}\sum_{t\in\mathbb{T}}\left[\sum_{l\in\mathbb{L}}\sum_{\varphi^{s}\in\mathcal{I}^{S}_{l,t,s}} C^{S,idle}_{l,\varphi^{s}}(\Pi(\tau^{\mathrm{t}}_{\mathrm{m}}))+\sum_{m\in\mathbb{M}}\sum_{\varphi^{T}\in\mathcal{I}^{T}_{m,t,s}} C^{T,idle}_{m,\varphi^{t}}(\Pi(\tau^{\mathrm{t}}_{\mathrm{m}}))\right]$$

$$-\frac{1}{|\mathbb{S}|}\sum_{s\in\mathbb{S}}\sum_{t\in\mathbb{T}}\left[\sum_{m\in\mathbb{M}} C^{CO_2}_{m,t}E_{m,t,s}(\Pi(\tau^{\mathrm{t}}_{\mathrm{m}}))+\sum_{l\in\mathbb{L}} C^{CO_2}_{l,t}E_{l,t,s}(\Pi(\tau^{\mathrm{t}}_{\mathrm{m}}))\right] \quad (1)$$

The first term represents the expected profit generated by hereditary attributes $h \in \mathbb{H}$, such as recovered metal, or other profit-related attributes, at location $i \in \mathbb{I} \cup \mathbb{K}$ in period $t \in \mathbb{T}$, under joint geological and equipment scenario realisations $s \in \mathbb{S} = \mathbb{S}_G \times \mathbb{S}_E$. The parameter $p_{h,i,t}$ represents the unit economic value of hereditary attribute h, which may be positive for revenues or negative for processing, rehandling, or disposal costs. Notably, this term only accounts for the profits incurred at the processing plant or other downstream facilities. It explicitly excludes mining extraction and haulage-related costs, which are instead captured in the third component of the objective function through the dynamic event sets generated by the discrete-event simulator. While traditional short-term scheduling models typically focus on minimising operational costs or meeting production targets, the proposed formulation explicitly maximises expected profit. This design choice is necessary because the framework actively evaluates the blending of extracted material. Since different material blends directly affect processing plant performance and final product quality, a profit-driven objective naturally guides the optimisation to generate the most economically advantageous blending strategies.

The second term penalises deviations ($d^{+}_{\varepsilon,n,t,s}$ and $d^{-}_{\varepsilon,n,t,s}$) from production requirements across locations, periods, attributes, and scenarios. The parameters $c^{+}_{\varepsilon,t}$ and $c^{-}_{\varepsilon,t}$ are the corresponding penalty costs for exceeding or falling short of the required production limits, respectively.

The third term penalises movement-related operating costs generated by the DES. These costs are evaluated over the simulated shovel movement events $\mathcal{M}^{S}_{l,t,s}$ and truck movement events $\mathcal{M}^{T}_{m,t,s}$. The parameters $C^{S,mov}_{l,\psi^{s}}(\Pi(\tau^{\mathrm{t}}_{\mathrm{m}}))$ and $C^{T,mov}_{m,\psi^{t}}(\Pi(\tau^{\mathrm{t}}_{\mathrm{m}}))$ represent the costs associated with shovel and truck movement events ψ^{s} and ψ^{T}, respectively. This component allows the objective function to account for additional operational costs caused by long relocations, inefficient routing, or excessive movement requirements. Notably, these costs depend on the dispatching policy, $\Pi(\tau^{\mathrm{t}}_{\mathrm{m}})$, and on the number of trucks of type m scheduled in period t, $\tau^{\mathrm{t}}_{\mathrm{m}}$. Therefore, they are not fixed input parameters, but simulation-dependent costs generated during the execution of the plan.

The fourth term penalises idle-time costs generated during simulation. The sets $\mathcal{I}^{S}_{l,t,s}$ and $\mathcal{I}^{T}_{m,t,s}$ contain idle events executed by shovels and trucks, respectively. The parameters $C^{S,idle}_{l,\varphi^{s}}(\Pi(\tau^{\mathrm{t}}_{\mathrm{m}}))$ and $C^{T,idle}_{m,\varphi^{t}}(\Pi(\tau^{\mathrm{t}}_{\mathrm{m}}))$ represent the cost associated with idle events φ^{s} and φ^{T}, respectively. This term captures operational inefficiencies such as shovel waiting time, truck queuing, truck starvation, and delays induced by downstream processing constraints.

The final term accounts for the environmental footprint of the fleet. The variables $E_{m,t,s}(\Pi(\tau_{\mathrm{m}}^{\mathrm{t}}))$ and $E_{l,t,s}(\Pi(\tau_{\mathrm{m}}^{\mathrm{t}}))$ represent the total emission deviations generated by shovels and trucks in period t under scenario s, respectively. The parameters $C_{m,t}^{CO_2}$ and $C_{l,t}^{CO_2}$ correspond to the emission costs or environmental penalty factors for trucks and shovels, respectively. This term allows the model to internalise the environmental impact of haulage and loading operations, discouraging solutions that achieve production targets at the expense of excessive fuel consumption or emissions. Notably, emissions also depend on the dispatching policy $\Pi(\tau_{\mathrm{m}}^{\mathrm{t}})$ and the fleet size $\tau_{\mathrm{m}}^{\mathrm{t}}$. It is worth mentioning that only upper limits are considered in this formulation for the environmental footprint, quantified in units of mass of carbon dioxide equivalent (CO_2-eq). Generally, only maximum allowable limits are applied over environmental indicators for industrial operations.

Constraints

To dynamically model the operational behaviour of the mining system, the DES maintains state variables that are required to enforce feasibility during execution. Most of these variables and constraints are related to equipment allocation, dispatching, production sequencing, and downstream processing. Similar operational constraints can be found in Ashtiani *et al* (2025); Both and Dimitrakopoulos (2020); de Carvalho and Dimitrakopoulos (2021, 2024); Kumar, Dimitrakopoulos and Maulen (2020); Kumar and Dimitrakopoulos (2021); Moradi Afrapoli, Tabesh and Askari-Nasab (2019). Related downstream processing and flow-equilibrium constraints can also be found in Goodfellow and Dimitrakopoulos (2016) and Montiel and Dimitrakopoulos (2017).

Most notably from this work, precedence constraints are divided into vertical and horizontal relationships. Vertical precedences are considered mandatory, since the extraction of a digline is not allowed until all its vertical predecessors have been extracted. Horizontal precedences, on the other hand, are defined to preserve a practical extraction direction while still allowing dynamic interaction between the extraction sequence and the equipment during the simulation.

In this regard, Figure 1 shows an example in which digline 9 is depicted in yellow, while its horizontal predecessors are depicted in red. The mathematical constraints for vertical and horizontal precedence relationships are defined as follows:

$$\sum_{l' \in \mathbb{L}} \sum_{j'=1}^{j-1} x_{l',j',b'} \geq x_{l,j,b} \; \forall b \in \mathbb{B}, b' \in \mathbb{B}^{Strong}(b), l \in \mathbb{L}, j \in \mathbb{J} \tag{2}$$

$$\sum_{l' \in \mathbb{L}} \sum_{j'=1}^{j-1} \sum_{b' \in \mathbb{B}^{Weak}(b)} x_{l',j',b'} \geq x_{l,j,b} \; \forall b \in \mathbb{B}, l \in \mathbb{L}, j \in \mathbb{J}, \mathbb{B}^{Weak}(b) \neq \emptyset \tag{3}$$

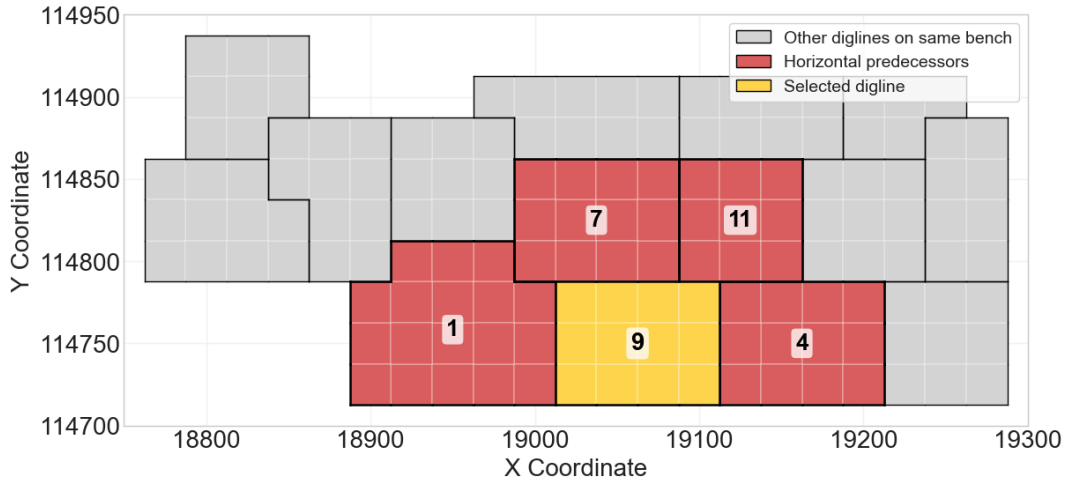


FIG 1 – Horizontal precedences definition.

According to Equation 2, if digline b is assigned to execution position j, then every vertical predecessor b' must have been assigned to an earlier position $j' < j$, possibly by any shovel l'. Similarly, Equation 3 states that if digline b is assigned to execution position j of shovel l, then at least one horizontal predecessor of b must have been scheduled before position j, again possibly by any shovel. This formulation captures the dynamic interaction between precedence feasibility and

equipment allocation, since the availability of a digline for one shovel may depend on the diglines previously extracted by other shovels. Therefore, the extraction sequence is not evaluated independently for each shovel, but as a coordinated system in which the progress of one piece of equipment can enable or restrict the future movements and production of others.

Notably, the DES is designed to respond dynamically to infeasible decisions during execution. This is achieved by replicating realistic operational behaviour. For example, if a digline does not satisfy either its vertical or horizontal precedence constraints, the corresponding shovel enters an idle state and waits until the next digline in its sequence becomes available, or 'open.' Moreover, mass production limits are strictly enforced as hard constraints, while arsenic, copper-to-iron rates, and iron grade are handled as penalised soft constraints. Instead of violating the hard constrained production limits, the DES resolves them dynamically: if more material is scheduled to be dumped into a processing location with no remaining capacity, the truck is forced to wait until spare capacity becomes available.

Solution approach

Due to the highly combinatorial nature of short-term scheduling and the computational intractability of embedding a DES within exact mathematical solvers, a metaheuristic approach is proposed to efficiently explore the solution space. The proposed solution framework addresses this complexity by hybridising a reactive multistart Greedy Randomised Adaptive Search Procedure (GRASP) (Resende and Ribeiro, 2016) with the Neural Diving (ND) hyper-heuristic scheme proposed by Yaakoubi and Dimitrakopoulos (2025). The resulting method, denoted as ND-GRASP, combines the diversification capabilities of a multistart strategy with the learning-guided intensification of reinforcement learning to generate, evaluate, and iteratively refine solutions (Figure 2).

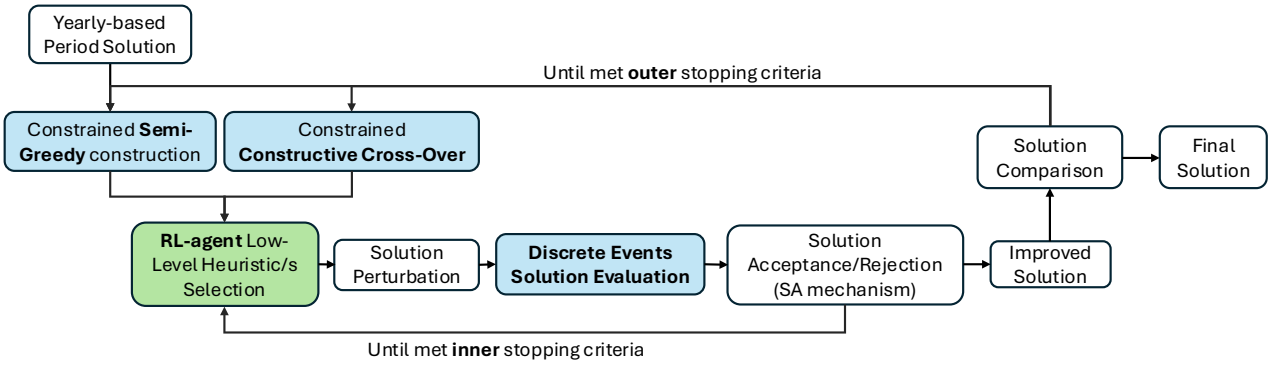


FIG 2 – Main optimisation workflow of the ND-GRASP framework.

The optimisation process begins with a construction phase that generates an initial feasible extraction sequence by assigning diglines to shovels, estimating the required truck fleet, and defining material routing decisions while respecting spatial and operational constraints. To strengthen diversification and escape local optima, a Constructive Crossover mechanism is incorporated (Ahmed, 2020). The algorithm maintains a bounded Elite Pool containing the highest-quality short-term solutions evaluated during the search. Before constructing a new solution, a probabilistic rule determines whether to follow a standard semi-greedy logic or apply crossover. When selected, the algorithm randomly chooses an elite schedule, defines a split period, and inherits the spatial and equipment allocation decisions up to that point. The remaining periods are then reconstructed, allowing the method to exploit high-performing early-period structures while diversifying later-period decisions (Ribeiro and Rosseti, 2007).

Because the constructed solutions are based on static approximations, they must be refined through an improvement phase. However, the search space is highly complex due to the interacting decisions regarding extraction sequencing, dynamic shovel allocation, truck fleet sizing, material destination policy, and downstream routing. Since every local modification must be evaluated under operational dynamics, exhaustive or purely randomised local search becomes computationally inefficient. To address this, ND-GRASP incorporates an RL agent that acts as a hyper-heuristic. Instead of blindly testing perturbations, the agent learns which operational changes, such as digline

sequence swaps, fleet adjustments, or routing modifications, are most likely to improve the current solution. Crucially, this RL agent functions as an online learning mechanism during the search process; rather than relying on offline pre-training across external data sets, the agent continuously adapts its heuristic selection policy dynamically to the specific problem instance being solved.

Candidate solutions are evaluated using a streamlined DES, which models the mining system through active shovel fronts, reactive truck hauling logic, and a dynamic flow graph tracking material routing and spatial depletion across multiple stochastic orebody realisations (Figure 3). To overcome the computational bottleneck of full-horizon evaluations, the framework utilises a state-caching mechanism. The DES stores operational snapshots at the boundary of each time period. When the RL agent modifies a decision in period t, the simulator reloads the cached state from period t−1 and resumes the simulation only from that point onward, drastically reducing computation time.

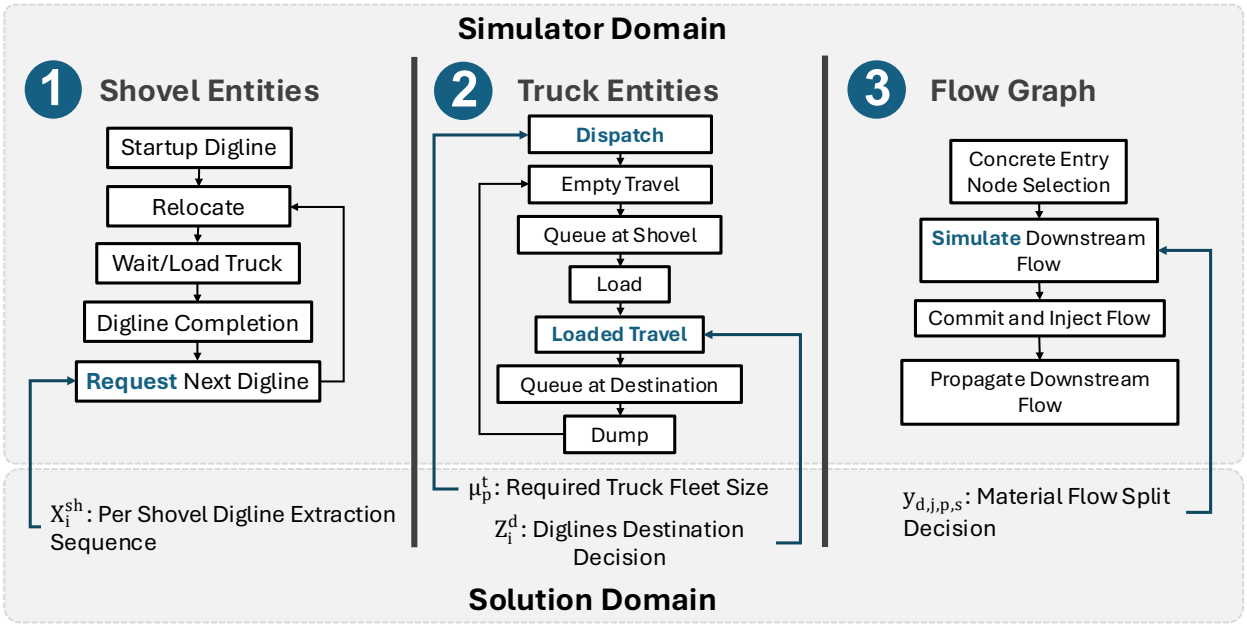


FIG 3 – Discrete-event simulator (DES) components, their main events, and their interaction with the solution.

The simulator outputs are transformed into a state representation comprising global operational indicators, such as expected tonnage deviations and fleet utilisation, alongside graph-based features detailing spatial relationships and digline completion indicators. Using this state, the RL agent dynamically selects low-level heuristics to guide the search. Perturbations are either accepted or not, based on a simulated annealing criterion (Godoy and Dimitrakopoulos, 2004; Goodfellow and Dimitrakopoulos, 2016; Montiel and Dimitrakopoulos, 2017), automatically accepting improvements while allowing non-improving moves according to a cooling temperature to avoid premature convergence. After each step, the RL agent receives a numerical reward proportional to the schedule's performance improvement (Lamghari and Dimitrakopoulos, 2020; Yaakoubi and Dimitrakopoulos, 2023, 2025), progressively refining its policy to better satisfy the complex's operational and sustainability objectives.

A CASE STUDY AT A COPPER MINING COMPLEX

The proposed method is applied to an operating copper mining complex. In this case study, the complex consists of two mines, denoted as Mine 1 and Mine 2. Material is extracted using a heterogeneous fleet of 17 shovels, comprising six different models, and is transported by a fleet of 87 trucks, comprising four different models, to five crushers, three waste dumps, and a sulfide leach pad, as shown in Figure 4. Two auxiliary destinations, Fuel Bay 1 and Fuel Bay 2, are located near Mine 1 and Mine 2, respectively, allowing trucks to refuel during operation. However, when required, trucks are allowed to travel to either fuel bay for refuelling. Material delivered to the crushers is subsequently transported by conveyor belts to three mills or to an oxide leach pad. Notably, the

downstream connections are not all-to-all; instead, spatial restrictions determine the downstream locations to which each crusher can send material.

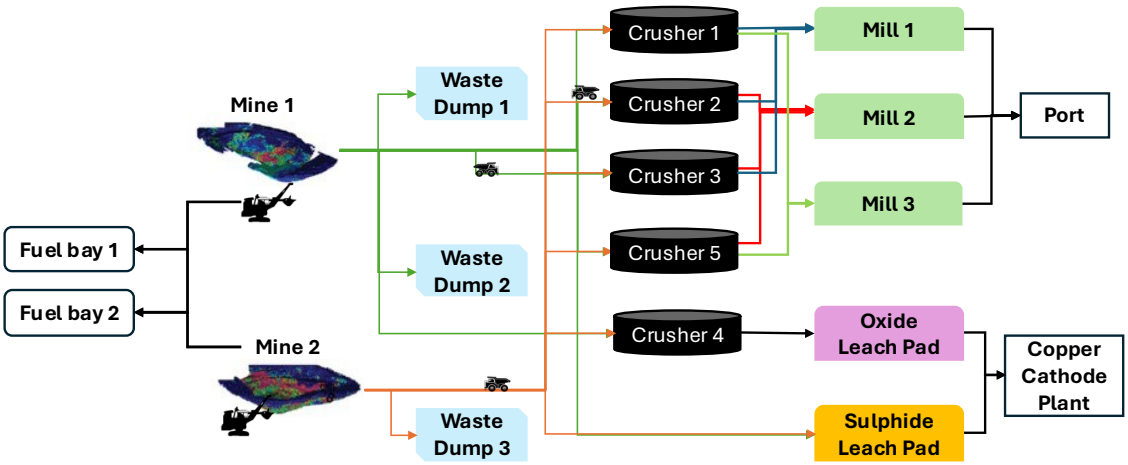


FIG 4 – Diagram of the copper mining complex.

The primary product of the complex is copper. Therefore, the main operational objective is to satisfy the production requirements of the different components of the mining complex while maximising profit by improving copper production. These production requirements are presented and scaled in Table 1 for confidentiality purposes. Geological uncertainty is modelled using a sequential Direct Block-support Simulation method (DBS) that incorporates vector orthogonalisation via minimum/maximum Autocorrelation Factors (MAF) (Boucher and Dimitrakopoulos, 2009; Godoy, 2003). Ten stochastic simulations characterise the spatial distribution of the copper (total and soluble), iron, and arsenic grades in each mine, as illustrated in Figure 5. Equipment-related uncertainty is derived from historical data following Quigley and Dimitrakopoulos (2020). The economic and operational parameters used at the copper mining complex are listed in Table 2 and are also scaled for confidentiality purposes.

TABLE 1

Production limits of the different components of the mining complex.

Location(s)	Attribute	Value (daily)
Crusher 1, 2, 3, 4, and 5	Production capacity (tonnes)	69 000, 125 000, 163 000, 60 000, 157 000
Waste dumps 1, 2, and 3	Handling capacity (tonnes)	160 000, 160 000, 80 000
Sulfide leach pad	Production capacity (tonnes), CuS/CuT ratio	370 000, 0.30
Oxide leach pad	Production capacity (tonnes)	62,200
Mill 1, 2, and 3	Production capacity (tonnes)	155 000, 142 000, 119 000
	CuT/Fe ratio, Fe grade, As grade (ppm)	0.42, 0.026, 75

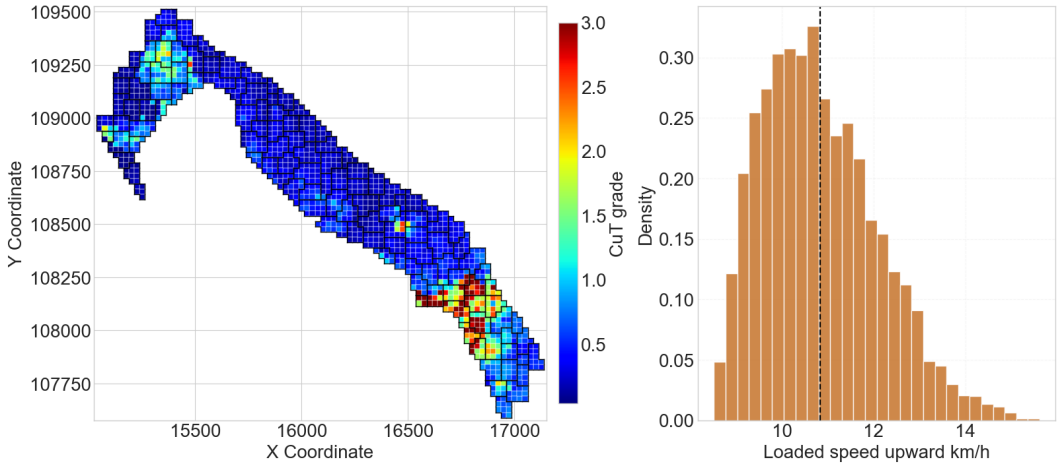


FIG 5 – Example of sources of uncertainty – (left): A simulation of total copper grade; (right): non-normal distribution of the speed for a loaded truck going upwards.

TABLE 2

Economic and operational parameters used at the copper mining complex.

Attribute	Value
Number of mining blocks	Mine 1: 3600 Mine 2: 1200
Production scheduling horizon	30 days
Slope angles	Mine 1: 37 Mine 2: 45
Radius (Mine 1 and 2)	2 blocks (~40 m)
Average road slope	10%
Max. shovels per area	Mine 1: [4, 5, 5] Mine 2: [3]
Copper recovery	Oxide leach pad: 65.0% Sulfide leach pad: 27.0% Mill 1, 2, and 3: 80.4%, 80.0%, 82.6%
Copper price	$5511/ton
Selling cost	Oxide leach pad: $551/ton Sulfide leach pad: $551/ton Mill 1, 2, and 3: $571/ton
Processing cost	Oxide leach pad: $5.81/ton Sulfide leach pad: $1.84/ton Mill 1, 2, and 3: $5.79/ton Crushers: $0.58/ton
Diesel cost and higher heating value (HHV)	$1.14/L, 10.7 kWh/L

The digline-definition procedure used in this work follows the algorithm proposed by de Carvalho and Dimitrakopoulos (2024). Notably, the diglines are generated using initial blasthole data and remain fixed throughout the simulation. While realistic industrial applications continuously update digline boundaries as new grade control data becomes available, the present framework maintains the initial spatial definition; the dynamic updating of diglines during execution is reserved for future

work. Once generated, these static diglines become the basic scheduling units of the short-term optimisation model. An example of the obtained diglines and their destination is shown in Figure 6. As observed, the shapes honour the minimal required space, while the destinations respond to the spatial distribution of the grades of the different attributes considered and the requirements of each specific location.

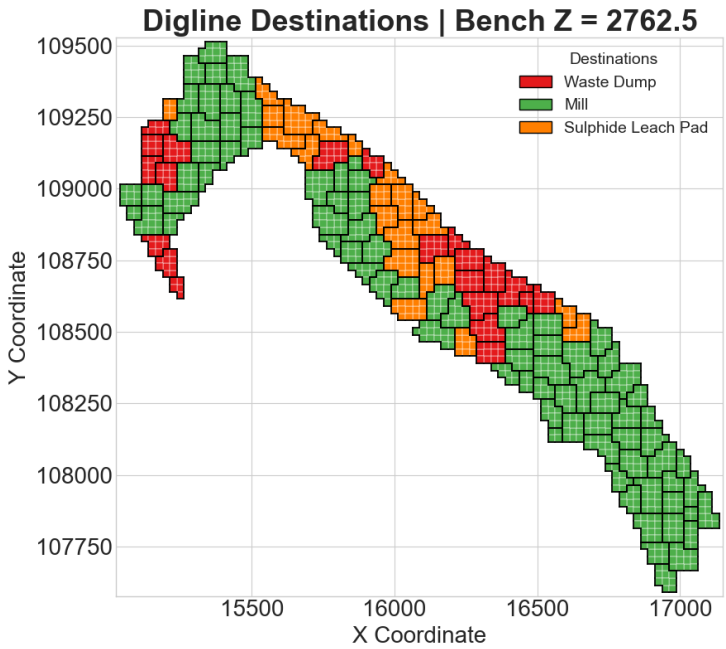


FIG 6 – Example of the defined diglines and their corresponding destination.

EXPERIMENTS AND COMPARISON ANALYSES

The proposed framework is evaluated through a comprehensive set of computational experiments designed to assess the effects of dynamic dispatching logic and environmental penalties on short-term scheduling performance. The optimisation process is driven by a reinforcement learning agent based on the synchronous Advantage Actor-Critic (A2C) architecture (Mnih *et al,* 2016). To enhance adaptability, the agent employs the Permanent/Transient learning principle proposed by Vemgal (2026), where a permanent component captures generalised search heuristics while a transient component continuously adapts to the immediate optimisation context. A concise evaluation and performance comparison of this A2C agent against alternative reinforcement learning architectures is provided in the Appendix, where performance was evaluated over ten independent runs to ensure statistical reliability and account for algorithmic variability across optimisation seeds.

All experiments were conducted on a machine equipped with an Intel Core Ultra 7 155H CPU (1.4 GHz) and 32 GB of RAM, running Windows 10. To balance computational efficiency and modern machine learning capabilities, the core optimisation algorithms and discrete-event simulator were implemented in C++, whereas the reinforcement learning agent was developed in Python 3.10.

The operational results presented in the following three figures represent a baseline scenario optimised without environmental penalties. Fleet movements are governed by the standard *m-trucks-one-shovel* dispatching policy (Alarie and Gamache, 2002), which is dynamically called whenever a truck requires an assignment to prioritise underserved shovels and reduce starvation. Furthermore, the simulation actively resolves operational states by diverting equipment to fuel bays once they reach a critical energy threshold, defined as the minimum energy required to complete a cycle and reach the nearest fuel bay. Additionally, queuing logic captures wait times at occupied shovels or dumping locations, triggering downstream flow redirection to alternative destinations when congestion occurs. Penalty costs are empirically calibrated by systematically varying their values and evaluating the resulting solution quality. The configuration that yields the most favourable overall performance is then retained.

Figures 7 and 8 illustrate the ability of the proposed optimisation framework to satisfy production targets and operational limits under stochastic short-term conditions. In Figure 7, the throughput profiles for Crushers 1 and 5 show that the 50th percentile (P_{50}) trajectory generally remains close to the production target during the 30-day horizon, with most deviations occurring only in isolated periods. These temporary drops are not necessarily caused by poor optimisation performance but rather by short-term operational constraints such as limited material availability, precedence constraints, shovel accessibility, truck dispatch interactions, and downstream congestion. The 10th to 90th percentiles (P_{10}–P_{90}) envelope also shows that the uncertainty around realised throughput is relatively contained for most periods, indicating that the schedule remains robust across equiprobable stochastic scenarios.

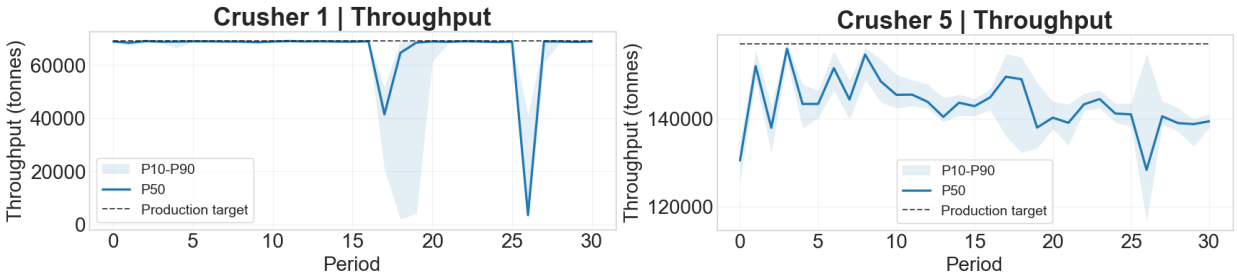


FIG 7 – Example of throughput material at Crushers 1 and 5 during the 30 days of operation.

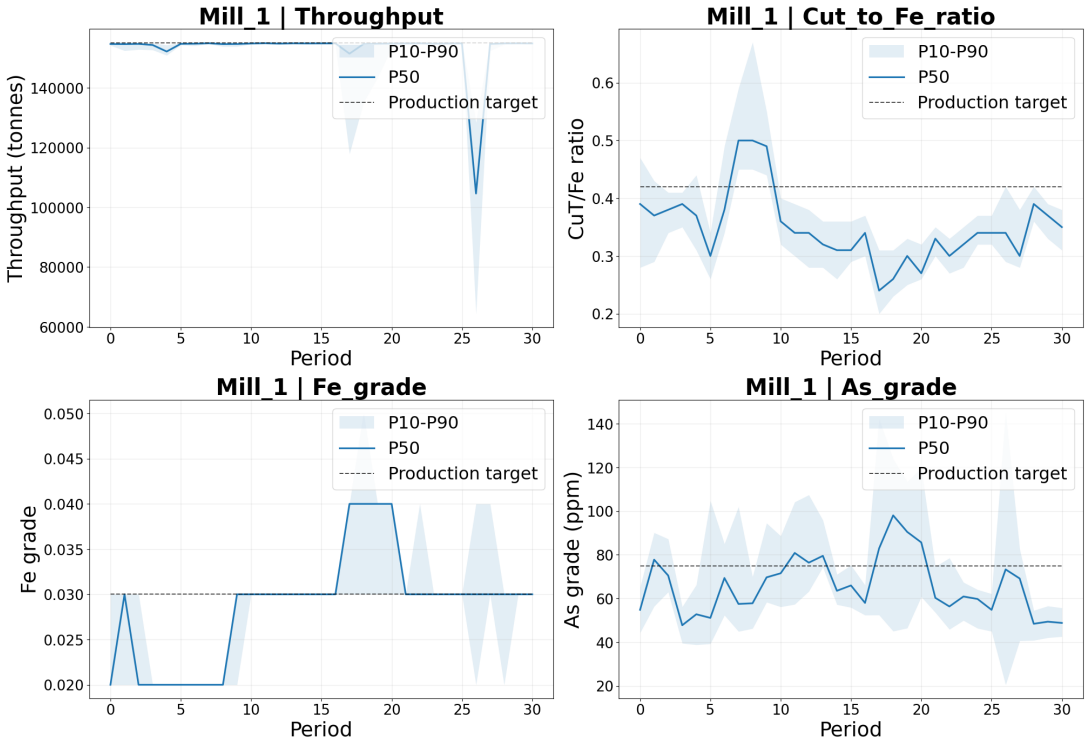


FIG 8 – Example of production limits and risk analysis across the 30 days for Mill 1.

A similar behaviour is observed in Figure 8 for Mill 1, where the throughput is maintained close to the target in most periods, while quality-related limits such as copper-to-iron (Cu/Fe) ratio, iron (Fe) grade, and arsenic (As) grade are explicitly monitored. Although some temporary violations or near-limit conditions appear, especially for the Cu/Fe ratio and Fe grade, these occur in specific periods where the available material blend constrains the feasible operating range. This highlights the advantage of the simulation-optimisation framework, which not only seeks high-value schedules but also evaluates whether the selected extraction sequence and material routing decisions can satisfy downstream production and quality constraints under uncertainty. The spatial extraction schedule map shown in Figure 9 further supports this, showing that the framework selects a feasible set of started diglines distributed across the active mining areas, while respecting spatial availability and extraction sequencing requirements. While these figures detail specific high-priority locations,

proportional target-compliance was successfully maintained across all other crushers, waste dumps, and leach pads, with average throughput deviations remaining within acceptable operational limits.

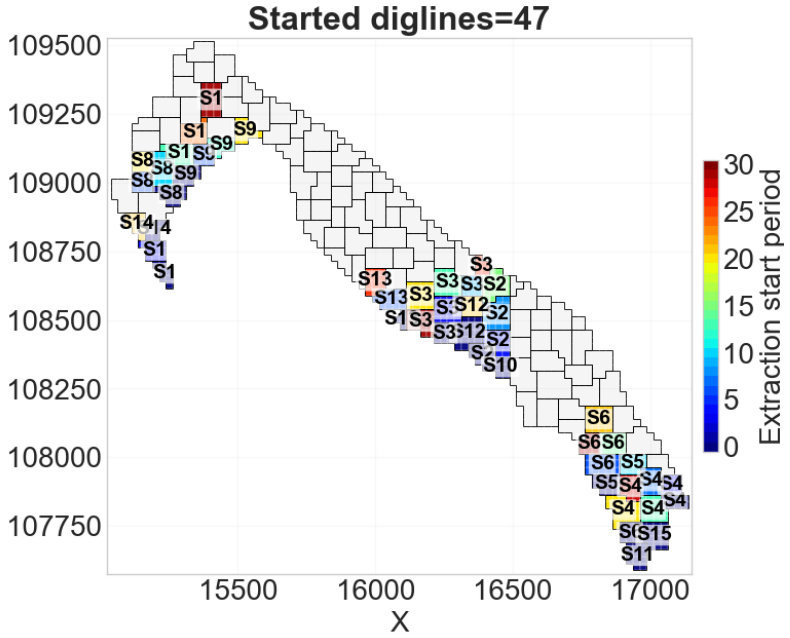


FIG 9 – Production schedule for bench Z=2762. White diglines are not extracted within the 30-day window.

A first experiment is performed to evaluate the impact of including carbon footprint penalties. This experiment evaluates whether accounting for fuel-related emissions changes the selected production schedule, fleet behaviour, and economic performance under dynamic systems. Three scenarios are run: (1) no carbon footprint penalties; (2) moderated carbon footprint penalties ($50/kg of CO_2 eq); (3) high carbon footprint penalties ($100/kg of CO_2 eq).

Figures 10 and 11 compare the effect of incorporating carbon-footprint penalties into the objective function. Emissions are calculated by first converting operational energy requirements into equivalent fuel consumption, which is then multiplied by standard emission factors to determine the total emissions output. As expected, increasing the emission penalty reduces both the period-by-period and cumulative carbon footprint, with the high-penalty scenario achieving the lowest cumulative emissions by the end of the 30-day horizon. This reduction is generally obtained at the expense of lower cumulative profit compared with the no-penalty case, indicating a trade-off between economic performance and environmental impact. However, the moderate-penalty scenario provides a more balanced outcome, reducing cumulative emissions by approximately 5.6 per cent while increasing cumulative profit by 4.6 per cent relative to the no-penalty case. This suggests that a moderate emission penalty can discourage inefficient haulage and equipment-use patterns without overly restricting production performance.

As the carbon-footprint penalty increases, the total tonnage moved tends to decrease, suggesting that the model is not simply reducing emissions but also limiting production. Instead, the penalty encourages more efficient fleet usage and routing decisions, allowing less material to be moved with a lower emission intensity. The utilisation profiles in Figure 11 further show that carbon penalties modify equipment usage patterns. Truck utilisation tends to decrease under stronger penalties, reflecting shorter or fewer haulage cycles, while shovel utilisation remains within a comparable range across scenarios. Nevertheless, under moderate penalties, average utilisation tends to increase while the number of trucks used per period decreases. This indicates a more efficient use of the available fleet, thereby reducing excessive operating costs and emissions. Overall, these results show that the proposed framework can internalise emission-related costs and adapt fleet behaviour accordingly, allowing the evaluation of profit-emission trade-offs rather than relying on a single economic objective.

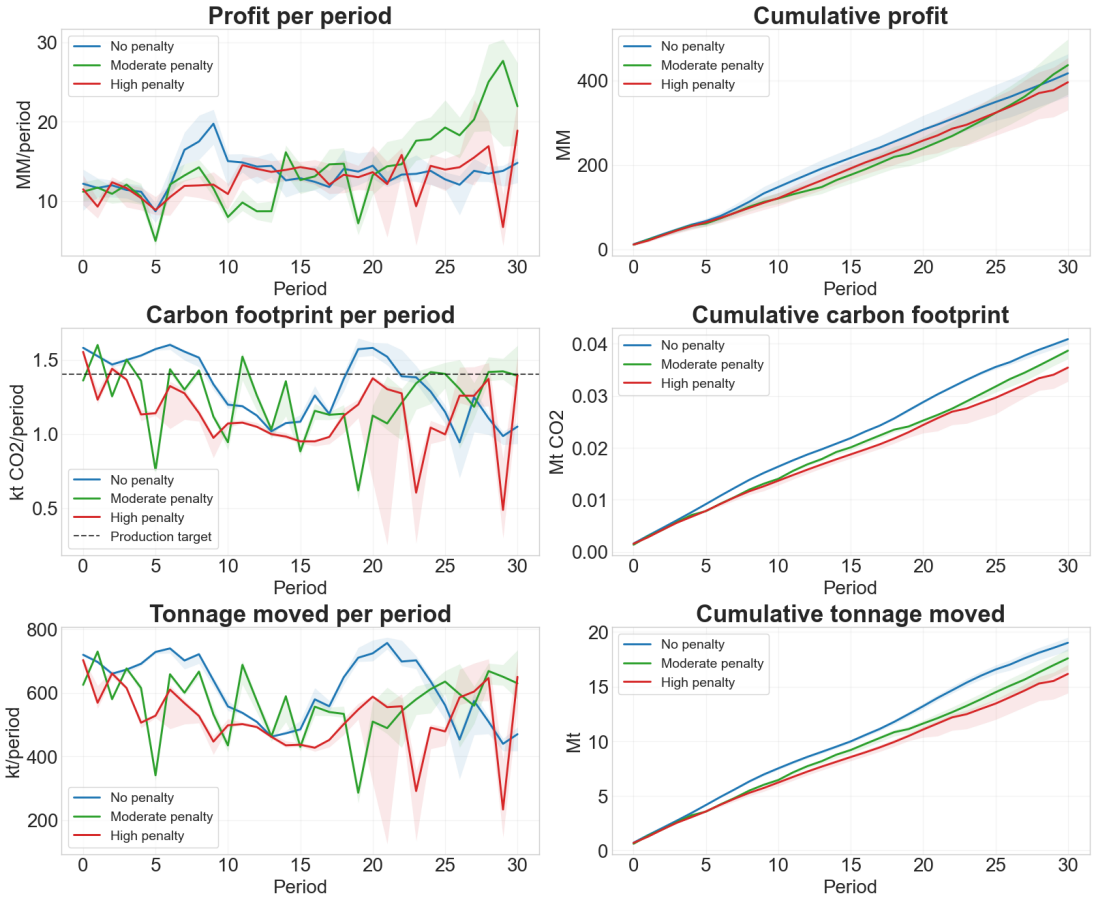


FIG 10 – Carbon footprint, profit, and moved tonnages risk analysis for the three carbon footprint scenarios evaluated.

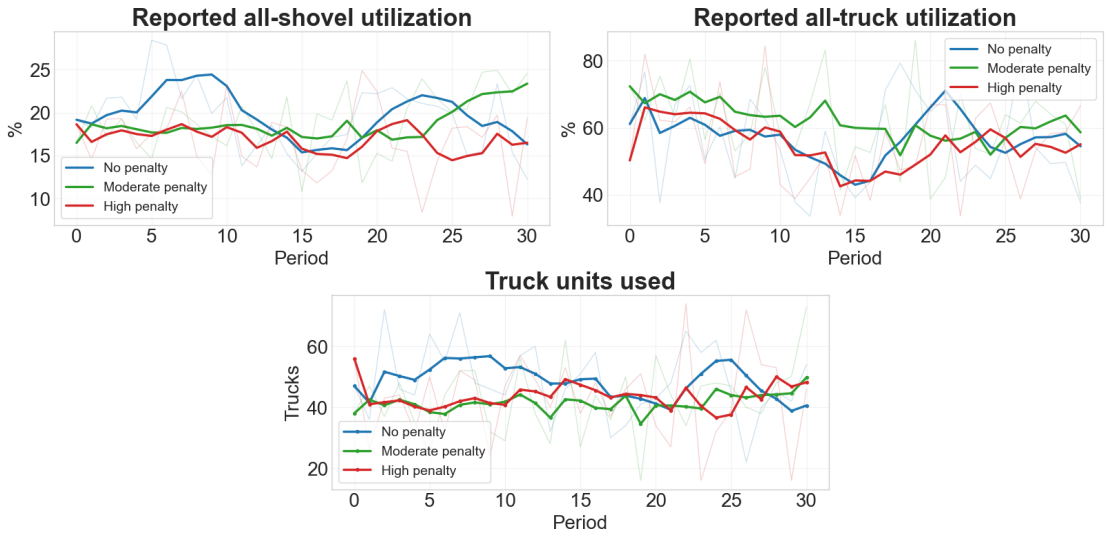


FIG 11 – Equipment utilisations for each of the three carbon footprint scenarios evaluated.

A second experiment is also developed so two truck-dispatching policies are compared: a *greedy-shovel* policy and the baseline *m-trucks-one-shovel* policy. The *greedy-shovel* policy assigns each available truck to the closest feasible shovel, prioritising short travel distances and reduced truck cycle times. This comparison evaluates whether local distance-based dispatching or need-based shovel balancing provides better operational performance under the proposed simulation-optimisation framework.

Figures 12 and 13 compare the performance of the *greedy-shovel* and *m-trucks-one-shovel* dispatching policies. The *greedy-shovel* policy achieves a higher cumulative profit, increasing the

final cumulative profit by approximately 9.7 per cent relative to the *m-trucks-one-shovel* policy. This improvement is obtained despite moving approximately 19.2 per cent less cumulative tonnage, suggesting that the greedy policy does not improve performance by increasing total material movement, but by selecting more spatially efficient and economically valuable truck–shovel interactions. By assigning trucks to the closest feasible shovel, the greedy policy reduces unnecessary haulage, shortens cycle times, and lowers cumulative emissions by approximately 15.4 per cent.

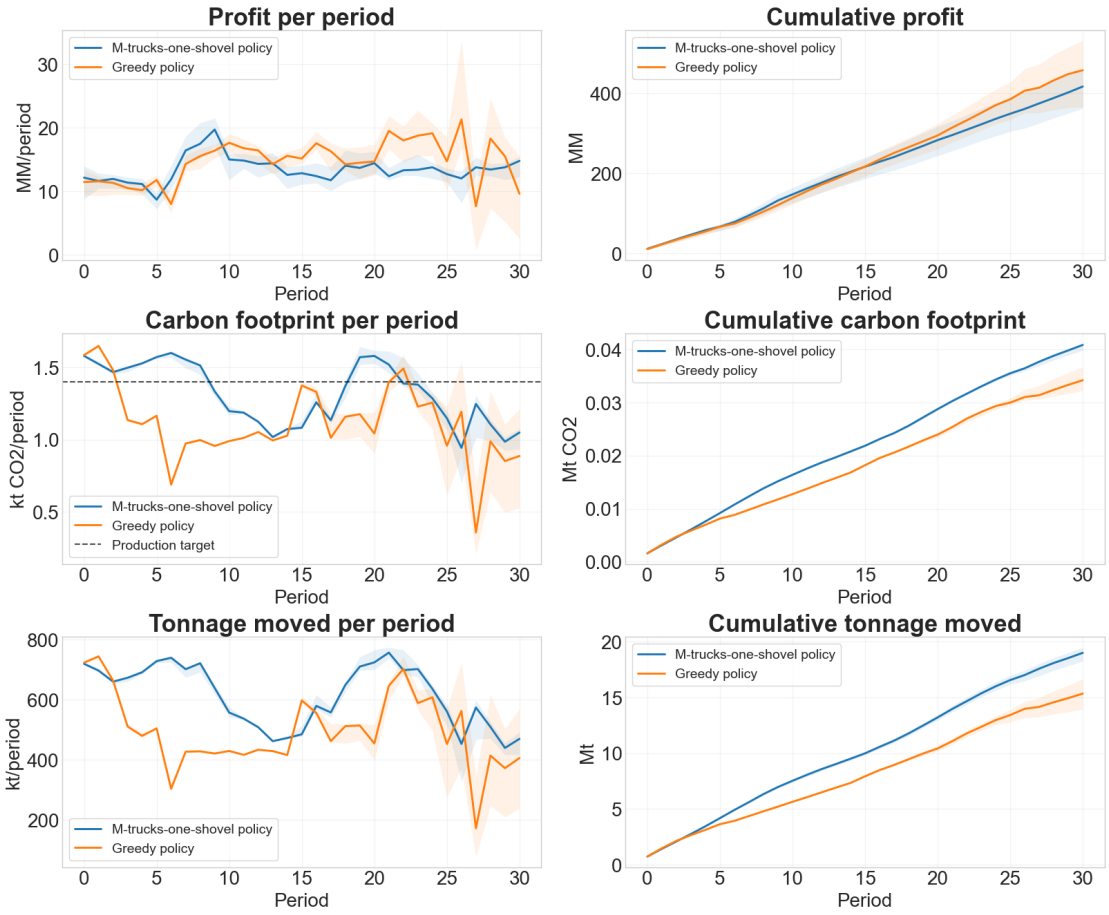


FIG 12 – Carbon footprint, profit, and moved tonnages risk analysis for the two dispatching scenarios evaluated.

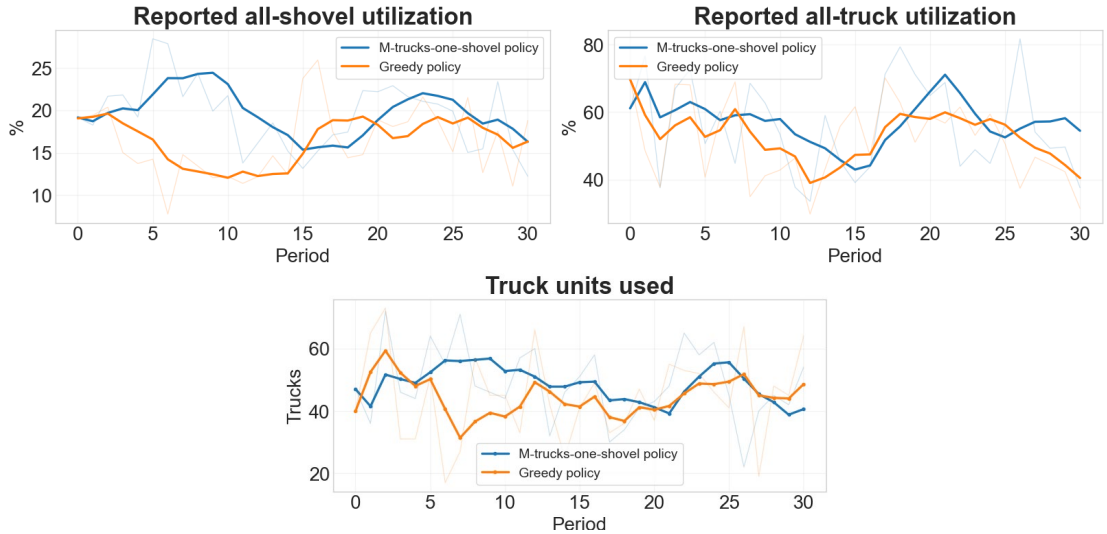


FIG 13 – Equipment utilisations for each of the two dispatching scenarios evaluated.

The higher cumulative tonnage moved by the *m-trucks-one-shovel* policy indicates that need-based shovel balancing can reduce shovel starvation and increase material movement but may also assign trucks to spatially less efficient loading units, increasing haulage effort and reducing profitability. The utilisation profiles in Figure 13 support this since the *m-trucks-one-shovel* policy generally maintains higher shovel and truck utilisation, whereas the greedy policy uses the truck fleet more selectively, reducing the average number of truck units used by approximately 7.0 per cent.

CONCLUSIONS

This work proposed a novel event-based optimisation framework for short-term production scheduling and fleet management under geological and equipment uncertainty, optimised using a reinforcement learning-based metaheuristic approach based on GRASP and Neural Diving (ND-GRASP). The framework integrates digline-based extraction decisions, shovel allocation, truck fleet usage, truck dispatching logic, downstream processing constraints, and environmental penalties within a stochastic DES-based evaluation.

The production-limit analysis demonstrates that the proposed framework is able to generate schedules that generally satisfy throughput and quality requirements across crushers, leach pads, and mills during the 30-day horizon. Temporary deviations occur in specific periods, mainly due to material availability, precedence relationships, shovel accessibility, equipment interactions, and downstream restrictions. However, the results show that the optimisation framework can explicitly monitor these operational constraints and adjust extraction and routing decisions under uncertainty.

The carbon-footprint experiments show that incorporating emission penalties changes fleet behaviour and material movement decisions. Higher penalties reduce cumulative emissions, while moderate penalties provide the best balance between economic and environmental performance. In particular, the moderate-penalty scenario reduces emissions and deviations from production targets while improving profit, suggesting that carbon pricing can discourage inefficient haulage patterns without necessarily reducing production. This indicates that environmental objectives can be internalised in the short-term scheduling process while maintaining operational viability.

Finally, the comparison between dispatching policies shows that truck-dispatching logic has a significant effect on both economic and environmental performance. The greedy-shovel policy achieves higher cumulative profit and lower cumulative emissions, despite moving less total material than the m-trucks–one-shovel policy. This result suggests that, for the analysed case study, haulage efficiency and material selectivity are more important than maximising total truck utilisation or total tonnage moved. In contrast, the m-trucks–one-shovel policy increases material movement but also increases emissions and reduces economic efficiency. Overall, the results confirm that the proposed framework can support short-term decision-making by improving operational feasibility, economic performance, and environmental control.

ACKNOWLEDGEMENTS

The work in this paper was funded by the NSERC Discovery Grant 239019, the industry consortium members of COSMO Stochastic Mine Planning Laboratory (AngloGold Ashanti, Anglo American/DeBeers, Agnico Eagle, BHP, IAMGOLD, Kinross Gold, Newmont, Vale SA and Vale Base Metals), and the Canada Research Chairs Program.

REFERENCES

Ahmed, Z H, 2020. Genetic Algorithm with Comprehensive Sequential Constructive Crossover for the Travelling Salesman Problem, *International Journal of Advanced Computer Science and Applications*, 11(5). https://doi.org/10.14569/ijacsa.2020.0110533

Alarie, S and Gamache, M, 2002. Overview of Solution Strategies Used in Truck Dispatching Systems for Open Pit Mines, *International Journal of Surface Mining, Reclamation and Environment*, 16(1):59–76. https://doi.org/10.1076/ijsm.16.1.59.3408

Ashtiani, M K, Afrapoli, A M, Doucette, J and Askari-Nasab, H, 2025. A stochastic energy-efficient robust simulation-based truck dispatching optimization for simultaneous ghg mitigation and operational excellence in open-pit mines, *Simulation Modelling Practice and Theory*, 138:103026. https://doi.org/10.1016/j.simpat.2024.103026

Badiozamani, M M, Pourrahimian, Y, Ben-Awuah, E, Afrapoli, A M and Askari-Nasab, H, 2024. MOL Report Twelve—2023/2024 (with University of Alberta Library and Mining Optimization Laboratory (MOL) Researchers/Graduate Students), University of Alberta Library. https://doi.org/10.7939/84202

Balboa-Espinoza, V and Dimitrakopoulos, R, 2026. Joint fleet management and stochastic short-term production planning in mining complexes with carbon footprint reduction, *International Journal of Mining, Reclamation and Environment*, 1–22. https://doi.org/10.1080/17480930.2026.2628757

Ballantyne, G R and Powell, M S, 2014. Benchmarking comminution energy consumption for the processing of copper and gold ores, *Minerals Engineering*, 65:109–114.

Both, C and Dimitrakopoulos, R, 2020. Joint stochastic short-term production scheduling and fleet management optimization for mining complexes, *Optimization and Engineering*, 21(4):1717–1743. https://doi.org/10.1007/s11081-020-09495-x

Boucher, A and Dimitrakopoulos, R, 2009. Block simulation of multiple correlated variables, *Mathematical Geosciences*, 41(2):215–237. https://doi.org/10.1007/s11004-008-9178-0

Calderon, J L, Bazilian, M, Sovacool, B, Hund, K, Jowitt, S M, Nguyen, T P, Månberger, A, Kah, M, Greene, S and Galeazzi, C, 2020. Reviewing the material and metal security of low-carbon energy transitions, *Renewable and Sustainable Energy Reviews*, 124:109789.

Cressie, N, 1990. The origins of kriging, *Mathematical Geology*, 22(3):239–252. https://doi.org/10.1007/BF00889887

David, M, 1977. *Geostatistical ore reserve estimation* (Elsevier Scientific Pub, Co).

de Carvalho, J P and Dimitrakopoulos, R, 2021. Integrating production planning with truck-dispatching decisions through reinforcement learning while managing uncertainty, *Minerals*, 11(6):587.

de Carvalho, J P and Dimitrakopoulos, R, 2024. Simultaneous shovel allocation and grade control decisions for short-term production planning of industrial mining complexes – an actor-critic approach, *International Journal of Mining, Reclamation and Environment*, 38(1):53–78. https://doi.org/10.1080/17480930.2023.2247196

De Donno, M G, 2024. Metals for the energy transition: Exploring opportunities amidst supply-demand imbalance, *SPE Europe Energy Conference and Exhibition*. https://doi.org/10.2118/219988-MS

Dimitrakopoulos, R and Ramazan, S, 2004. Uncertainty based production scheduling in open pit mining, *SME Transactions*, 316.

Dimitrakopoulos, R, 2011. Stochastic optimization for strategic mine planning: A decade of developments, *Journal of Mining Science*, 47(2):138–150. https://doi.org/10.1134/S1062739147020018

Dimitrakopoulos, R, Farrelly, C T and Godoy, M, 2002. Moving forward from traditional optimization: Grade uncertainty and risk effects in open-pit design, *Mining Technology*, 111(1):82–88. https://doi.org/10.1179/mnt.2002.111.1.82

Dowd, P, 1994. Risk assessment in reserve estimation and open-pit planning, *Transactions of the Institution of Mining and Metallurgy, Mining Technology*, 103:A148–A154.

Frieire Guimaraes, V and Dimitrakopoulos, R, 2026a. Integrating haulage management with waste encapsulation and progressive reclamation in long-term production planning of mining complexes under uncertainty, *Trans of the Inst of Mining and Metallurgy: Mining Technology*, 135(3)305–317. https://doi.org/10.1177/25726668261459286

Frieire Guimaraes, V and Dimitrakopoulos, R, 2026b. Simultaneous stochastic optimisation of mining complexes: Integrating waste management and progressive reclamation with encapsulation, *International Journal of Mining, Reclamation and Environment*, 14 p. https://doi.org/10.1080/17480930.2026.2654147

Godoy, M and Dimitrakopoulos, R, 2004. Managing risk and waste mining in long-term production scheduling of open-pit mines, *SME Transactions*, 316(3):43–50.

Godoy, M, 2003. The effective management of geological risk in long-term production scheduling of open pit mines, PhD thesis, The University of Queensland. https://doi.org/10.14264/106571

Goodfellow, R C and Dimitrakopoulos, R, 2016. Global optimization of open pit mining complexes with uncertainty, *Applied Soft Computing*, 40:292–304. https://doi.org/10.1016/j.asoc.2015.11.038

Hessel, M, Modayil, J, van Hasselt, H, Schaul, T, Ostrovski, G, Dabney, W, Horgan, D, Piot, B, Azar, M and Silver, D, 2018. Rainbow: Combining Improvements in Deep Reinforcement Learning, *Proceedings of the 32nd AAAI Conference on Artificial Intelligence, AAAI 2018*, 32(1):3215–3222. https://doi.org/10.1609/aaai.v32i1.11796

Hoerger, S, Seymour, F and Hoffman, L, 1999. Mine planning at Newmont's Nevada operations, *Mining Engineering*, 51(10):26–30.

Huo, D, Sari, Y A, Kealey, R and Zhang, Q, 2023. Reinforcement learning-based fleet dispatching for greenhouse gas emission reduction in open-pit mining operations, *Resources, Conservation and Recycling*, 188:106664.

Jeswiet, J and Szekeres, A, 2016. Energy consumption in mining comminution, *Procedia CIRP*, 48:140–145.

Journel, A G and Huijbregts, C J, 1978. *Mining geostatistics* (Academic Press).

Kamran, M, Raugei, M and Hutchinson, A, 2023. Critical elements for a successful energy transition: A systematic review, *Renewable and Sustainable Energy Transition*, 4:100068.

Krige, D G, 1951. A statistical approach to some basic mine valuation problems on the Witwatersrand, *Journal of the Southern African Institute of Mining and Metallurgy*, 52(6):119–139.

Kumar, A and Dimitrakopoulos, R, 2021. Production scheduling in industrial mining complexes with incoming new information using tree search and deep reinforcement learning, *Applied Soft Computing*, 110:107644. https://doi.org/10.1016/j.asoc.2021.107644

Kumar, A, Dimitrakopoulos, R and Maulen, M, 2020. Adaptive self-learning mechanisms for updating short-term production decisions in an industrial mining complex, *Journal of Intelligent Manufacturing*, 31(7):1795–1811. https://doi.org/10.1007/s10845-020-01562-5

Lamghari, A and Dimitrakopoulos, R, 2020. Hyper-heuristic approaches for strategic mine planning under uncertainty, *Computers and Operations Research*, 115:104590. https://doi.org/10.1016/j.cor.2018.11.010

Laurence, D, 2011. Establishing a sustainable mining operation: An overview, *Journal of Cleaner Production*, 19(2–3):278–284.

Levinson, Z and Dimitrakopoulos, R, 2020. Simultaneous stochastic optimisation of an open-pit gold mining complex with waste management, *International Journal of Mining, Reclamation and Environment*, 34(6):415–429. https://doi.org/10.1080/17480930.2019.1621441

Levinson, Z and Dimitrakopoulos, R, 2024. Simultaneous stochastic optimisation of mining complexes: Integrating progressive reclamation and waste management with contextual bandits, *International Journal of Mining, Reclamation and Environment*, 38(9):701–721. https://doi.org/10.1080/17480930.2024.2342687

Matamoros, M E V and Dimitrakopoulos, R, 2016. Stochastic short-term mine production schedule accounting for fleet allocation, operational considerations and blending restrictions, *European Journal of Operational Research*, 255(3):911–921. https://doi.org/10.1016/j.ejor.2016.05.050

Mnih, V, Badia, A P, Mirza, M, Graves, A, Lillicrap, T, Harley, T, Silver, D and Kavukcuoglu, K, 2016. Asynchronous methods for deep reinforcement learning, in *Proceedings of The 33rd International Conference on Machine Learning*, 48:1928–1937 (Proceedings of Machine Learning Research (PMLR)). Available from: <https://proceedings.mlr.press/v48/mniha16.html>

Mnih, V, Kavukcuoglu, K, Silver, D, Rusu, A A, Veness, J, Bellemare, M G, Graves, A, Riedmiller, M, Fidjeland, A K and Ostrovski, G, 2015. Human-level control through deep reinforcement learning, *Nature*, 5187540. 529–533.

Montiel, L and Dimitrakopoulos, R, 2017. A heuristic approach for the stochastic optimization of mine production schedules, *Journal of Heuristics*, 23(5):397–415. https://doi.org/10.1007/s10732-017-9349-6

Moradi Afrapoli, A, Tabesh, M and Askari-Nasab, H, 2019. A multiple objective transportation problem approach to dynamic truck dispatching in surface mines, *European Journal of Operational Research*, 276(1):331–342. https://doi.org/10.1016/j.ejor.2019.01.008

Moran, C J, Lodhia, S, Kunz, N C and Huisingh, D, 2014. Sustainability in mining, minerals and energy: New processes, pathways and human interactions for a cautiously optimistic future, *Journal of Cleaner Production*, 84:1–15.

Nikbin, R, Bagherpour, R, Purhamadani, E and Taherinia, A, 2025. Enhancing sustainability in mining by reducing hauling energy consumption through optimization of distance and slope with semi-mobile in-pit crushers and conveyors, *Scientific Reports*, 15(1):22119.

Pell, R, Tijsseling, L, Palmer, L W, Glass, H J, Yan, X, Wall, F, Zeng, X and Li, J, 2019. Environmental optimisation of mine scheduling through life cycle assessment integration, *Resources, Conservation and Recycling*, 142:267–276. https://doi.org/10.1016/j.resconrec.2018.11.022

Quigley, M and Dimitrakopoulos, R, 2020. Incorporating geological and equipment performance uncertainty while optimising short-term mine production schedules, *International Journal of Mining, Reclamation and Environment*, 34(5):362–383. https://doi.org/10.1080/17480930.2019.1658923

Rahnema, M, Amirmoeini, B and Moradi Afrapoli, A, 2023. Incorporating Environmental Impacts into Short-Term Mine Planning: A Literature Survey, *Mining*, 3(1):163–175. https://doi.org/10.3390/mining3010010

Ravenscroft, P J, 1992. Risk analysis for mine scheduling by conditional simulation, *Transactions of the Institution of Mining and Metallurgy, Mining Industry*, 101:A104–A108.

Resende, M G C and Ribeiro, C C, 2016. *Optimization by GRASP*: *Greedy Randomized Adaptive Search Procedures,* (Springer: New York). https://doi.org/10.1007/978-1-4939-6530-4

Ribeiro, C C and Rosseti, I, 2007. Efficient parallel cooperative implementations of GRASP heuristics, *Parallel Computing*, 33(1):21–35.

Schulman, J, Wolski, F, Dhariwal, P, Radford, A and Klimov, O, 2017. Proximal policy optimization algorithms. *arXiv Preprint,* 1707.06347. Available from: <https://arxiv.org/abs/1707.06347>

Stone, P, Froyland, G, Menabde, M, Law, B, Pasyar, R and Monkhouse, P H L, 2007. Blasor—Blended iron ore mine planning optimisation at Yandi, Western Australia, in *Advances in Applied Strategic Mine Planning,* pp 39–46 (Springer).

Upadhyay, S P and Askari-Nasab, H, 2019. Dynamic shovel allocation approach to short-term production planning in open-pit mines, *International Journal of Mining, Reclamation and Environment*, 33(1):1–20. https://doi.org/10.1080/17480930.2017.1315524

Vemgal, N, 2026. The Permanent and Transient Learning Framework for Continual Reinforcement Learning, PhD thesis, McGill University. https://doi.org/10.82308/55720

Whittle, G, 2007. Global asset optimisation, *Orebody Modelling and Strategic Mine Planning*, 14:361–366.

Whittle, J, 2010. The global optimizer works–what next, in *Advances in Orebody Modelling and Strategic Mine Planning*, Spectrum Series 17, 3–5 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Xu, X, Gu, X, Wang, Q, Zhao, Y and Wang, Z, 2021. Open pit limit optimization considering economic profit, ecological costs and social benefits, *Transactions of Nonferrous Metals Society of China*, 31(12):3847–3861. https://doi.org/10.1016/S1003-6326(21)65769-2

Yaakoubi, Y and Dimitrakopoulos, R, 2023. Learning to schedule heuristics for the simultaneous stochastic optimization of mining complexes, *Computers and Operations Research*, 159:106349. https://doi.org/10.1016/j.cor.2023.106349

Yaakoubi, Y and Dimitrakopoulos, R, 2025. Decision-focused neural adaptive search and diving for optimizing mining complexes, *European Journal of Operational Research*, 320(3):699–719. https://doi.org/10.1016/j.ejor.2024.07.024

Yokoi, R, Watari, T and Motoshita, M, 2022. Future greenhouse gas emissions from metal production: Gaps and opportunities towards climate goals, *Energy and Environmental Science*, 15(1):146–157.

APPENDIX

To determine the most effective learning strategy for the proposed framework, different reinforcement learning agents were compared within the ND-GRASP framework against a baseline Adaptive Large Neighbourhood Search (ALNS) hyper-heuristic proposed by Lamghari and Dimitrakopoulos (2020).

The evaluated RL agents include an ND selector based on the synchronous Advantage Actor-Critic (A2C) method (Mnih *et al,* 2016); an ND_PPO selector, based on Proximal Policy Optimisation (Schulman *et al,* 2017); and a value-based ND_RAINBOW selector inspired by the Rainbow deep Q-learning framework of Hessel *et al* (2018), which combines several improvements to the original Deep Q-Network proposed by Mnih *et al* (2015). In this work, the value-based agent does not include the spectral-normalisation component because of its additional computational overhead (Yaakoubi and Dimitrakopoulos, 2025).

To ensure statistical reliability, the results are reported based on ten independent runs for each algorithmic variant. Furthermore, each optimisation run was executed for a strict 120 min limit to approximate a realistic industrial deployment setting.

As shown in Figure 14, all ND-GRASP variants outperform the baseline ALNS, with an average improvement of approximately 10 per cent. The best performance is obtained by the ND-GRASP variant using the A2C agent, which achieves an average improvement of 13 per cent over the P_{50} objective function value obtained by the ALNS hyper-heuristic. This result is consistent with the ability of ND-GRASP to explore the solution space more efficiently and escape local optima through its intrinsic diversification strategy. However, as shown on the right-hand side of Figure 14, the ND-PPO variant exhibits the lowest variability across independent runs (-2.0 per cent), indicating greater consistency and robustness. Despite this, the ND-GRASP variant using the A2C agent was selected for all operational analyses in the main study because it consistently achieves the highest absolute objective function values.

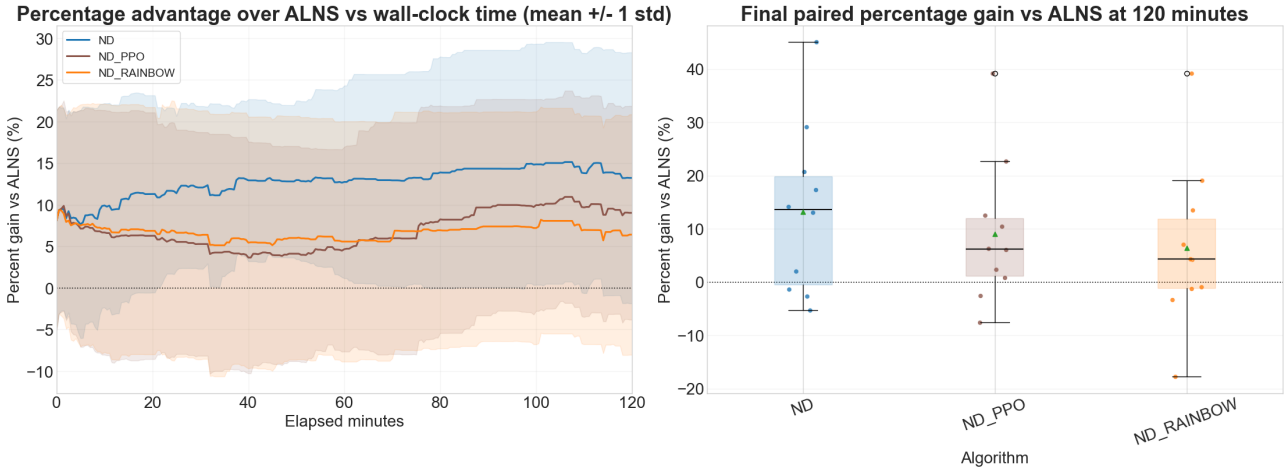


FIG 14 – Performance analysis of each algorithm and percentage gain versus ALNS baseline.

Learning interpretable dispatch heuristics for real-time open pit truck dispatching under uncertainty

J Luo[1] and R Dimitrakopoulos[2]

1. Postdoc, McGill University, Montreal QC H3A 2A7, Canada. Email: jingyu.luo@mcgill.ca
2. Professor, McGill University, Montreal QC H3A 2A7, Canada. Email: roussos.dimitrakopoulos@mcgill.ca

INTRODUCTION

Efficient real-time truck dispatching is important to short-term production control in open pit mines. Dispatching decisions are made repeatedly during a shift while truck locations, shovel queues, remaining muck pile tonnage, destination congestion, and production status evolve dynamically. These decisions must account for geological grade uncertainty, stochastic cycle times, and equipment availability. Because each decision must be made in real time, solving the underlying stochastic optimisation problem at every dispatch event is computationally impractical, and heuristic rules are widely used. Manually configured rules are fast and transparent, but their site-specific logic may not balance production, grade control, and congestion effectively under uncertainty.

This study proposes a large language model (LLM)-driven evolutionary framework for the offline design of explicit, interpretable truck-dispatching rules. The framework generates executable scoring functions, evaluates them in a discrete-event simulator (DES) on training instances, and revises them using performance feedback. Each instance specifies the muck pile layout, grades, equipment fleet, travel-time structure, and production targets. After the search, the rule with the best training fitness is returned and applied unchanged to unseen instances. Deployment requires only arithmetic and conditional logic over the current operational state, with no further LLM calls or search iterations. Learning-based real-time dispatch policies have also recently been studied using deep reinforcement learning (Noriega, Pourrahimian and Askari-Nasab, 2025).

DISPATCHING PROCESS AND OBJECTIVE

For rule design and evaluation, the truck-dispatching problem is modelled as a sequence of real-time assignment decisions over a production shift, with performance evaluated from the resulting production and grade outcomes. The DES tracks unloaded travel, shovel queuing, loading, loaded travel, destination queuing, dumping, cumulative mill feed, and realised delivered grade. A dispatch event occurs at the start of the shift for an initially available truck or when an operational truck completes dumping and becomes idle. The rule then selects one feasible muck pile for the truck to visit next. The truck travels unloaded to the selected pile, queues if its assigned shovel is busy, receives a load, and hauls the material to the destination determined by the pile's pre-shift ore/waste classification. Ore is sent to the mill, while waste is sent to the dump.

A dispatch rule is represented as a priority scoring function. At each dispatch event e, the rule assigns a score to every feasible muck pile for the idle truck, and the selected muck pile x_e is:

$$x_e = \underset{m \in \mathbb{M}^{\text{feas}}(\tau_e, t_e)}{\arg\max} \; h\big(\phi(\tau_e, t_e, m)\big)$$

Here, τ_e is the event time, t_e is the idle truck, $\mathbb{M}^{\text{feas}}(\tau_e, t_e)$ is the feasible muck pile set, and φ is the observable feature vector for truck t_e and candidate pile m. Features cover estimated pile grade, remaining tonnage, effective load, destination, shovel congestion, nominal cycle time, shift progress, and production and grade target gaps. Operational uncertainty is represented by log-normal variability in loading, haul, return, and dumping times ($CV = 0.10$) and independent shift-level truck and shovel availability (0.85), with unavailable equipment excluded for the shift.

The objective rewards contained metal delivered to the mill while penalising production shortfall, delivered-grade deviations on either side of the target, and mill feed above its upper bound. With unit penalty weights, the normalised objective F for one simulated shift is:

$$F = \widehat{M}_H - \widetilde{D}_P^- - \widetilde{D}_G^- - \widetilde{D}_G^+ - \widetilde{D}_C$$

where the first term is normalised contained metal, followed by normalised penalties for production shortfall, below-target grade, above-target grade, and upper-bound violation. A higher value indicates better performance.

Across the modelled uncertainty, rule design seeks:

$$\max_{h \in \mathcal{H}} \mathbb{E}_{\omega,\xi,A}[F(h)]$$

where $\mathcal{H}$ denotes the admissible rule class, and ω, ξ, and A denote geological realisations, stochastic operating times, and equipment-availability outcomes. Candidate fitness is estimated by DES over the training pool. Per-instance mean objective values are aggregated using a winsorised mean with $q = 0.1$, which limits the influence of extreme availability or geological outcomes on the search signal.

LLM-BASED SEARCH FRAMEWORK AND BENCHMARK COMPARISON

The LLM-based search maintains an archive of high-performing rules, initialised with three hand-crafted seed rules. At each generation, the LLM receives the problem description, the scaled feature interface, selected archive rules, and DES-based fitness feedback. Each request uses one of four instructions: explore, refine, simplify, or restructure. Returned Python functions are validated, deduplicated, and evaluated on the training instances. The archive is updated after each generation, and the rule with the best training fitness is returned. The learned rules retain explicit, inspectable arithmetic and conditional logic that can be related to operational states. Figure 1 gives a compact equivalent form of the selected rule.

```python
def score(features):
    prod, dev, is_mill = features["PROD_frac"], features["GRAD_dev_norm"], features["is_mill"]
    congestion = 1.0 + 1.9 * features["committed_frac"]
    throughput = features["load_frac"] / max(features["cycle_frac"] * congestion, 1e-9)
    if is_mill > 0.5:
        grade_correction = -features["G0_rel"] * dev * (25.0 + 15.0 * abs(dev))
        metal_incentive = features["G0_rel"] * max(0.0, 1.25 - prod)
        if prod < 0.85:
            score = (2.8 * (1.11 - prod) + features["production_urgency"] + metal_incentive + grade_correction) * throughput
        elif prod < 1.25:
            score = ((grade_correction * (1.5 + abs(dev)) + metal_incentive * (1.25 - prod)) * throughput
                     - 350.0 * features["projected_overprod_frac"])
        else:
            score = (-1500.0 * features["projected_overprod_frac"]
                     - 800.0 * features["PROD_gap_upper_frac"] + grade_correction)
    else:
        if prod < 0.75:
            waste_priority = -18.0
        elif prod < 1.11:
            waste_priority = 12.0 * (prod - 1.11)
        else:
            waste_priority = 5.0 + 55.0 * (prod - 1.11) + 2.5 * (1.0 - features["RTC_frac"])
        score = waste_priority * throughput
    return float(score)
```

FIG 1 – Compact equivalent form of the selected LLM-generated dispatching rule.

The selected rule has an explicit production-dependent structure. It discounts candidate trips by congestion-adjusted cycle time; while mill feed is short, it prioritises mill-bound material; near the production target, it emphasises grade correction; and near or above the upper production bound, it suppresses additional mill feed and shifts dispatch priority towards waste.

We benchmark the LLM search against Genetic Programming (GP), a well-established evolutionary hyper-heuristic for automated rule design, and against 14 predefined dispatching rules used as baselines. GP uses the same feature interface and DES-based fitness evaluation as the LLM search. The baseline rules include classical and single-attribute heuristics, three hand-crafted composite seed rules, and a tuned DISPATCH-style real-time rule combining shovel waiting, destination congestion, and production-plan shortfall (Moradi Afrapoli, Tabesh and Askari-Nasab, 2019).

EXPERIMENTAL SETTING AND RESULTS

Rule learning and out-of-sample evaluation require many diverse instances, so we use the available case data to calibrate a synthetic instance generator. Standard instances span a 10-hour shift and contain 6–12 muck piles, 2–4 shovels, and 12–24 trucks in 150- and 200-ton payload classes. The standard training and test pools each contain 30 instances and are disjoint. The large training and

test pools likewise contain 30 instances each, with 20–30 muck piles, 5–8 shovels, and 40–60 trucks. The same sampling mechanisms are used at both scales, making instance size the primary systematic difference.

Each instance has 15 spatially correlated geological realisations generated using a Gaussian copula over pile locations. Search fitness uses eight realisations and three operational replications per realisation (24 replications per instance), while testing uses all 15 realisations and five replications per realisation (75 per instance). Both methods are run independently 30 times. The LLM search uses 30 generations of 16 candidates and averages 408 distinct DES-evaluated rules per run; GP uses 300 individuals for 70 generations and averages 17 843. Because the evaluation scenarios and random-seed grids are fixed, reevaluating an unchanged rule would reproduce the same fitness; both methods therefore reuse recorded fitness when a previously evaluated rule appears again.

Test performance is measured by the mean gap to the best observed value on each test instance, with lower values indicating better performance. For a rule h, let $\bar{F}_j(h)$ denote its mean end-of-shift objective over all test replications on instance j, and let F_j^* denote the highest such mean attained by any rule evaluated on the same test pool. With N_{test} denoting the number of instances in the test pool, $N_{\text{test}} = 30$ for both test pools. The gap is:

$$\text{Gap}(h) = \frac{1}{N_{\text{test}}} \sum_{j=1}^{N_{\text{test}}} \left(F_j^* - \bar{F}_j(h) \right).$$

Paired comparisons use two-sided Wilcoxon signed-rank tests on per-instance method means across the 30 test instances. Table 1 summarises held-out performance on the standard and large-instance test pools. For brevity, the table reports only the tuned DISPATCH-style rule among the 14 baseline dispatching rules; comparisons with all 14 are summarised in the text below. The large-trained LLM row reports five exploratory runs; all other learned-method rows report 30 runs. Each rule evaluation covers the full training pool and its uncertainty replications.

TABLE 1

Held-out performance on the standard and large-instance test pools.

Method	Training pool	Rule evaluations per run	Standard gap	Large gap
LLM search (30 runs)	Standard	408	0.257 ± 0.049	0.389 ± 0.212
GP (30 runs)	Standard	17 843	0.290 ± 0.085	0.718 ± 0.470
LLM search (5 runs)	Large	408	0.292 ± 0.059	0.216 ± 0.021
Tuned DISPATCH-style rule	Standard	n/a	0.993	0.895

On the standard test pool, the LLM search achieves a numerically lower mean gap than GP (0.257 versus 0.290), although the difference is not statistically significant ($p = 0.237$). The LLM search uses only 2.3 per cent of GP's rule evaluations and shows lower observed run-to-run variability, with standard deviations of 0.049 and 0.085, respectively. It also achieves lower gaps than all 14 baseline dispatching rules. Relative to the tuned DISPATCH-style rule, its gap is 0.736 lower, and it is better on 29 of 30 instances ($p < 0.001$).

On the large-instance test pool, the standard-trained LLM search achieves a mean gap of 0.389 versus 0.718 for GP and is better on 29 of 30 instances ($p < 0.001$). Five exploratory LLM runs trained directly on large instances further reduce the mean gap to 0.216.

Compared with GP, the LLM search achieves a numerically lower standard-scale gap using only 2.3 per cent of the rule evaluations, shows lower run-to-run variability, and transfers more effectively to larger instances. Compared with the 14 baseline dispatching rules, it achieves substantially lower gaps while retaining explicit and inspectable dispatch logic.

REFERENCES

Moradi Afrapoli, A, Tabesh, M and Askari-Nasab, H, 2019. A multiple objective transportation problem approach to dynamic truck dispatching in surface mines, *European Journal of Operational Research*, 276(1):331–342.

Noriega, R, Pourrahimian, Y and Askari-Nasab, H, 2025. Deep reinforcement learning based real-time open-pit mining truck dispatching system, *Computers and Operations Research*, 173:106815.

LOforSP – Learning-to-optimise for two-stage stochastic programs with continuous recourse

L Shchichko[1], S S Chauhan[2], E Delage[3] and R Dimitrakopoulos[4]

1. PhD Candidate, Concordia University, Montréal QC H3G 2V4, Canada. Email: l_shchic@live.concordia.ca
2. Full Professor, Concordia University, Montréal QC H3G 2V4, Canada. Email: satyaveer.chauhan@concordia.ca
3. Full Professor, HEC Montréal, Montréal QC H3T 2A7, Canada. Email: erick.delage@hec.ca
4. Full Professor, McGill University, Montréal QC H3A 0G4, Canada. Email: roussos.dimitrakopoulos@mcgill.ca

INTRODUCTION

Two-stage stochastic programs (2SP) with continuous recourse are a cornerstone of optimisation under uncertainty, but solving them to optimality through the extensive form (EF) becomes computationally prohibitive as the number of scenarios and first-stage variables grows. This paper introduces L2Ofor2SP, a learning-to-optimise (L2O) framework whose aim is to find a good first-stage solution x* for a 2SP directly from problem data, replacing repeated online optimisation with a single fast forward pass. Two complementary learning strategies are proposed: (i) imitation-based learning, a supervised approach that uses pre-computed 'expert' solutions as labels together with conventional losses such as mean squared error or negative log-likelihood; and (ii) objective-based learning, a self-supervised approach (Tang, Khalil and Drgoňa, 2024) that requires no labels and instead minimises the average objective function value of the stochastic program across the training samples. Both strategies embed a constrained optimisation layer inside the neural network (NN), differentiating the solution of the optimisation model with respect to its input.

The contributions are summarised as follows:

- L2Ofor2SP reduces computational complexity by providing a compact surrogate representation of the 2SP that is leveraged to obtain a high-quality first-stage solution; this surrogate is easier to solve online (at inference/deployment time) than the EF and than earlier L2O proposals (Bengio *et al,* 2020; Patel *et al,* 2022; Larsen *et al,* 2024).
- The approach is generic and scalable — the network is trained at one set of variable dimensions yet, once trained, can be used with any (potentially much larger) dimensions, because every learned weight matrix acts on a single block, period, material type or scenario at a time rather than on the full index set, with material-type and scenario information instead summarised into a fixed-length, permutation-invariant context vector (mean, standard deviation, extremes and higher moments).
- A new 2SP benchmark from the mining industry is introduced, in which one decides the period in which to extract each block and how the extracted material is allocated across the value chain, supplied with ready-to-go parametrisation (Zhang, Nault and Dimitrakopoulos, 2019).
- To the best of the authors' knowledge, this is the first work to apply this objective-based (self-supervised) L2O paradigm to two-stage stochastic programming, as distinct from Tang, Khalil and Drgoňa's (2024) self-supervised L2O for mixed-integer non-linear programming.

In the mining application, the deposit is discretised into blocks $i \in I$, each extractable in a period $t \in T$; a block's material is split across material types $j \in J$ (eg ore, waste), with prices and grades subject to scenario-based uncertainty $s \in S$. The first-stage binary decisions x fix, for every block, the period in which it is extracted. The second-stage (recourse) decisions allocate each period's extracted material across the value chain (direct-to-plant processing, stockpiling, and later rehandling from stockpile to plant), subject to per-period processing-capacity, blending grade-bound and stockpile mass-balance constraints. The objective is the expected discounted NPV of the resulting extraction and processing schedule.

PROBLEM STATEMENT

A standard formulation of the two-stage stochastic program (Ahmed and Shapiro, 2002) is:

$$g(x) := c^T x + \mathbb{E}_\xi[\, Q(x, \xi(\omega)) \,] \quad \text{s.t.} \quad Ax \le b, \quad x \in \{0,1\}^{nx}$$

$$\text{with} \quad Q(x, \xi(.)) := \min \{\, q^T y \;:\; Tx + Wy \le h, \; y \in \mathbb{R}^{ny} \,\}$$

where $\xi := (q, T, W, h)$ collects the parameters of the second-stage problem, some or all of which are random, written $\xi(\omega)$, and the expectation is taken with respect to the (known) distribution $\mathcal{P}$ of $\xi(\omega)$ with support Ω. The value function Q returns the cost of the optimal recourse decisions $y \in \mathbb{R}^{ny}$ under a realisation of ξ, given the first-stage decisions x. The expectation is approximated by the sample average approximation (SAA): using K scenarios $\xi_1, \ldots, \xi_K$ drawn from $\mathcal{P}$ and, for ease of exposition, a uniform distribution, the EF creates an explicit copy of the second-stage variables for each scenario:

$$\min \quad c^T x + K^{-1} \Sigma_k \, q_k^T y_k \quad (k = 1, \ldots, K) \quad \text{s.t.} \quad Ax \le b, \;\; x \in \{0,1\}^{nx}; \quad T_k x + W_k y_k \le h_k, \;\; y_k \in \mathbb{R}^{ny}$$

Each problem instance is a realisation of data $\psi \in \Psi$ comprising deterministic first-stage parameters (c, A, b) and a finite set of random second-stage parameters $\{\xi_k = (q_k, W_k, T_k, h_k)\}$. When optimal labels x^* are available, the data set supports imitation-based learning; in their absence the data is used in a self-supervised fashion. A parametric prediction model maps the data ψ, through hidden layers of arbitrary width, to a first-stage problem coefficient vector $\hat{c}_\theta(\psi)$. A first optimisation layer then produces a relaxed first-stage solution $\hat{x}_\theta(\psi) \in [0,1]^{nx}$ by solving $\operatorname{argmin} -\hat{c}_\theta(\psi)^T x + ||x||_2^2$ over $\mathcal{X} := \{\, x \in [0,1]^{nx} : Ax \le b \,\}$, and a second optimisation layer recovers the recourse solution $\hat{y}$. A feasible binary first-stage decision is then recovered at inference time by re-solving the original mixed-integer program with $\hat{c}_\theta(\psi)$ as the objective coefficients over the original constraint set ($Ax \le b$, $x \in \{0,1\}^{nx}$); this assumes relatively complete recourse, ie that a feasible second-stage response exists for every first-stage decision reachable by this step. Two losses are used to train the network's parameters θ. The imitation (supervised) loss minimises the discrepancy between predicted and expert decisions:

$$\mathcal{L}_{\text{imit}}(\theta) := \mathbb{E}_{\psi \sim \mathcal{P}}[\, l(\, \hat{x}_\theta(\psi), x^* \,) \,]$$

while the objective-based (self-supervised) loss directly minimises the expected cost of the stochastic program without requiring labels:

$$\mathcal{L}_{\text{obj}}(\theta) := \mathbb{E}_{\psi \sim \mathcal{P}}[\, c^T \hat{x}_\theta(\psi) + K^{-1} \Sigma_k \, q_k^T \, \hat{y}_k(\, \hat{x}_\theta(\psi), \xi_k \,) \,]$$

where $\hat{y}_k$ denotes the optimal second-stage recourse associated with the estimated first-stage solution $\hat{x}_\theta(\psi)$ and scenario ξ_k. This formulation explicitly ties the learning task to the underlying stochastic optimisation objective. Gradients of either loss are propagated back through the constrained optimisation layers via the envelope theorem.

SELECTED EXPERIMENTAL RESULTS

Figure 1 presents the training progress when using the imitation and objective losses, respectively (validation curves follow the same trends, with no evidence of overfitting). The loss unit in the objective-based approach is net present value (NPV); thus, higher is better. The four curves correspond to four variants of the score network: a block-only variant using only block-related features; a time-only variant using only value-chain (downstream) related features; and two variants ('block dot time' and 'time dot block') that compute the score as the dot product of two parallel sub-networks, one over block features and one over downstream features. The 'block dot time' variant, hereafter 'Our', gave the best performance on both the training and validation sets and is used throughout the results below.

Once the model is trained and the best parameters θ are found, the approach is tested on an independent set of samples drawn from the same distribution. For benchmarking, the authors: (i) compute a best-known objective value of the 2SP in its EF using a commercial solver to a mixed-integer programming (MIP) gap of 2 per cent, and (ii) compute the objective value using a mining field-specific heuristic, the block economic value (BEV) approach (adjusted here to incorporate aggregated scenario information), which decomposes the 2SP into a first-stage model and a second-stage model and solves them sequentially, but only once.

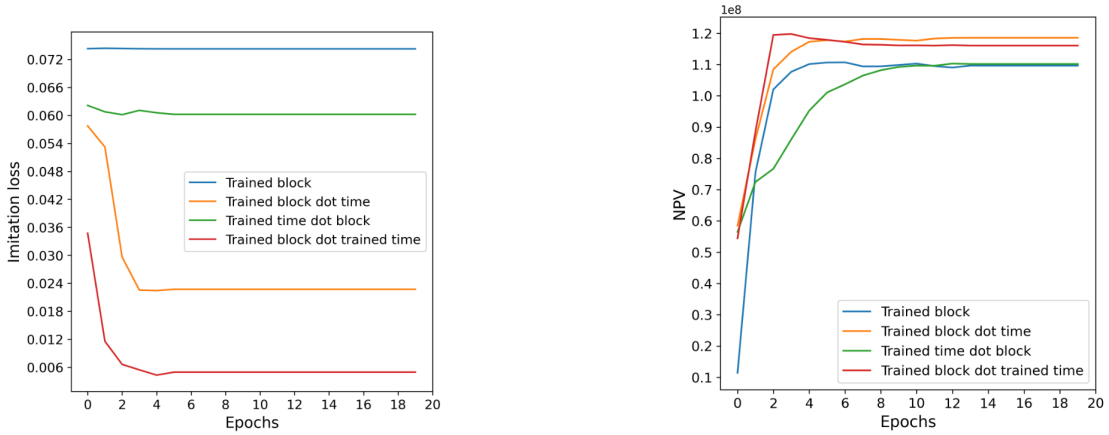


FIG 1 – Training progress with: (a) imitation loss; and (b) objective-based loss; for four score-network variants ('Trained block', 'Trained block dot time', 'Trained time dot block', 'Trained block dot trained time'). 'Our' uses the 'block dot time' variant.

The benchmark comparison is reported in Table 1. |I|, |T|, |J| and |S| denote the number of blocks, periods, material types and scenarios, respectively (|J| = |S| = 10 in all settings shown). 'Value' is the mean NPV, in units of $\$10^4$, across 200 test instances, rounded to the nearest whole unit for display; gaps are computed from the underlying (unrounded) means, so recomputing them from the rounded, displayed values will not reproduce the reported figures exactly. Gap_{BEV} = (Our value − BEV value) / Our value, and $\text{Gap}_{\text{E}}\text{F}$ = (EF value − Our value) / EF value, both expressed as percentages. 'Time' is the cumulative solving time over the 200 test instances; for 'Our' this comprises the forward pass to compute the score and the re-solve of the surrogate optimisation problem described above, excluding offline data-generation and training time (amortised across deployment on many instances). The approach produces near-best-known solutions in considerably less time than the EF, with time savings especially pronounced on larger instances. Relative to the BEV benchmark, the approach exhibits slightly higher compute time but a directionally consistent NPV advantage across all six test configurations (Gap_{BEV} between 0.02 per cent and 0.11 per cent); this is reported as a possible runtime-quality trade-off rather than a claim of a statistically confirmed advantage, as no per-instance dispersion measures have been computed to date. The lower BEV time at |I| = 150 than at |I| = 100 for |T| = 10 (97 sec versus 135 sec) is most likely explained by the fact that each instance size is generated independently, so problem difficulty varies from size to size by chance, alongside the fixed |T| potentially giving larger |I| more scheduling flexibility.

TABLE 1

Comparison of mean objective values ($\$10^4$, over 200 test instances) of the trained solution ('Our'), the BEV baseline, and the best-known EF solution (2 per cent MIP gap), with cumulative solving times, gaps and mean dollar savings on the test set.

|I|	|T|	Value (×10^4 \$)			Time (sec)			Gap_{BEV} (%)	$\text{Gap}_{\text{E}}\text{F}$ (%)	\$ saved
		Our	BEV	EF	Our	BEV	EF			
60	5	11 127	11 115	11 134	76	25	196	0.108	0.063	108 890
100	5	17 785	17 766	17 793	93	31	249	0.107	0.048	189 790
150	5	22 494	22 485	22 558	127	39	345	0.041	0.286	92 520
60	10	11 042	11 032	11 044	217	102	1732	0.094	0.023	103 390
100	10	17 945	17 941	17 960	178	135	1972	0.024	0.082	43 430
150	10	24 317	24 307	24 328	246	97	1900	0.041	0.047	100 040

CONCLUSION

This paper proposes a learning approach that finds a near-best-known first-stage 2SP solution in a short amount of time. Within the proposed approach, two decision-focused loss functions are

designed, and gradients are propagated through the constrained optimisation layers via the envelope theorem. The approach is empirically validated on a real-world, complex problem from the mining industry.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the support of Concordia University, IVADO (Montreal, Canada) and the COSMO Stochastic Mine Planning Laboratory at McGill University (Montreal, Canada).

REFERENCES

Ahmed, S and Shapiro, A, 2002. The sample average approximation method for stochastic programs with integer recourse, *Optimization Online*, pp 1–24.

Bengio, Y, Frejinger, E, Lodi, A, Patel, R and Sankaranarayanan, S, 2020. A learning-based algorithm to quickly compute good primal solutions for stochastic integer programs, in *Proceedings of the International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, pp 99–111 (Springer).

Larsen, E, Frejinger, E, Gendron, B and Lodi, A, 2024. Fast continuous and integer L-shaped heuristics through supervised learning, *INFORMS Journal on Computing*, 36(1):203–223.

Patel, R M, Dumouchelle, J, Khalil, E B and Bodur, M, 2022. Neur2SP: neural two-stage stochastic programming, *Advances in Neural Information Processing Systems*, 35:23992–24005.

Tang, B, Khalil, E B and Drgoňa, J, 2024. Learning to optimize for mixed-integer non-linear programming with feasibility guarantees, *arXiv preprint*, arXiv:2410.11061.

Zhang, J, Nault, B R and Dimitrakopoulos, R G, 2019. Optimizing a mineral value chain with market uncertainty using Benders decomposition, *European Journal of Operational Research*, 274(1):227–239.

Closure and sustainable mine planning

Regenerative mine planning – embedding closure, carbon sequestration and circular land use into surface mine design

E Chauke[1] and M Mochubele[2]

1. Mining Engineering Student, University of the Witwatersrand, School of Mining Engineering, Johannesburg Gauteng 2001, Republic of South Africa. Email: 2354602@students.wits.ac.za
2. Lecturer, University of the Witwatersrand, School of Mining Engineering, Johannesburg Gauteng 2001, Republic of South Africa. Email: mothusi.mochubele@wits.ac.za

ABSTRACT

Mine rehabilitation in South Africa is treated, in practice, as a residual liability discharged at end of life rather than as a design problem. The legislative framework, anchored in National Environmental Management Act (NEMA) (Republic of South Africa, 1998a) and the Mineral and Petroleum Resources Development Act (MPRDA) (Republic of South Africa, 2002), has required progressive rehabilitation for over two decades but does not specify what the rehabilitated land must do. This paper proposes a regenerative mine planning framework that closes that gap by treating productive post-mining land use, carbon performance and progressive rehabilitation as state variables inside the strategic mine plan, on the same footing as grade, recovery and net present value. The framework allocates each parcel of the rehabilitated surface to one of three land capability tiers. Tier 1 runs controlled rotational grazing coupled to on-site anaerobic digestion of livestock manure, yielding biogas, digestate and a measurable soil organic carbon pathway. Tier 2, on water-constrained or shallow-soil areas, runs closed-loop hydroponic production at approximately 90 per cent water recirculation. Tier 3, on severely degraded areas, runs phytoremediation followed by vertical growing on engineered platforms. Energy, water and nutrient flows are coupled so the system runs as a single integrated farm. Four credit-eligible carbon pathways operate simultaneously and map onto methodologies recognised under section 19(c) of the Carbon Tax Act 15 of 2019 (Republic of South Africa, 2019). The framework reads naturally into simultaneous stochastic optimisation by adding five surface variables per cell. An illustrative numerical example on a hypothetical 500 hectare (ha) surface mine indicates that internalising these variables converts rehabilitation from a depreciating liability into a continuing revenue stream, with annual revenue exceeding cost at approximately half of the productive footprint, breaking dependency on the closure provision fund. The framework is developed against South African statute but its planning logic and objective function are jurisdiction-agnostic and internationally applicable.

INTRODUCTION

A surface mine occupies the land for a defined, finite period. When the orebody is exhausted, the local economy the operation supported is exhausted with it, while the physical and biogeochemical consequences of extraction remain: disturbed topography, lost topsoil, suppressed vegetation, atmospheric carbon debt, disturbed aquifers, oxidised sulfide minerals, contaminated groundwater and altered geotechnical conditions. Conventional closure programmes treat these as sunk costs at end of life. Rehabilitated land, where it exists, rarely supports productive economic activity. South Africa's state-funded legacy mine closure liability is estimated at R30 to R60 billion against approximately 6100 derelict and ownerless sites (Council for Geoscience, 2024; Watson and Olalde, 2019). Successive Auditor-General audits show that even active operations frequently under-provide against statutory closure obligations (Bench Marks Foundation, 2022; Cole, 2024).

The legislative position has been clear for more than 20 years. The Minerals Act (50 of 1991) (Republic of South Africa, 1991), the Mineral and Petroleum Resources Development Act (MPRDA; 28 of 2002) (Republic of South Africa, 2002) and National Environmental Management Act (NEMA; 107 of 1998) (Republic of South Africa, 1998a) established environmental obligations and made the closure certificate the legal end point of a mining right. The Financial Provisioning Regulations (Government Notice (GN) R1147 of 2015; Department of Environmental Affairs, 2015) require an annual rehabilitation plan, a final rehabilitation, decommissioning and closure plan, and an environmental risk assessment report each updated annually, each costed, each demanding evidence of concurrent rehabilitation. What the legislation does not specify is what the progressively

rehabilitated land must do. The legislative gap is in the productive destination of the rehabilitated surface.

This paper closes that gap. The framework rests on three positions. First, the rehabilitated mine surface is heterogeneous and must be treated as such: each parcel is allocated to a tier defined by soil depth, drainage, residual contamination and proximity to water and energy infrastructure. Second, the tiers are coupled: manure feeds a digester; the digester feeds biogas to combined heat and power (CHP) and digestate back to soil; the CHP drives the hydroponic loop; phytoremediation biomass is harvested away from the food chain. The system runs as one integrated farm. Third, closure quality, carbon performance and post-mining productivity are specified at the strategic planning stage as inputs to surface design, dump shaping, topsoil management, drainage and the production schedule. In weak commodity cycles, the discounted value of progressively rehabilitated productive land can become a more important determinant of the schedule than the marginal grade in the next strip.

The framework is developed against South African statutes and the Just Transition debate because that is where the execution gap is most severely documented, but the planning logic itself is jurisdiction-agnostic: the three-tier differentiation, the subsystem coupling, and the augmented objective function attach to any surface mine closure regime, with only tier definitions and the regulatory anchoring of carbon and rehabilitation accounting requiring local adaptation, as is already the case across Western Australian (Government of Western Australia, 2020), Queensland and Canadian provincial closure frameworks. Its four contributions: (1) formalising three-tier capability differentiation and coupling design; (2) setting out four credit-eligible carbon pathways aligned with section 19(c) Carbon Tax Act methodologies; (3) specifying surface variables that allow simultaneous stochastic optimisation on the same footing as grade and recovery; (4) identifying regulatory and capability gaps for moving from concept to audited practice.

REGULATORY AND CONCEPTUAL CONTEXT

The South African statutory framework

South African mine closure law has tightened steadily since 1991. The MPRDA formalised closure planning; the 2014 One Environmental System reform brought environmental governance under NEMA, where section 24P establishes the duty to make financial provision for rehabilitation, decommissioning and management of residual impacts. GN R1147 of 2015 specifies the three plans accompanying every mining right. The post-2015 regime is built on a single design principle: rehabilitation runs in parallel with extraction, not at the end. This is the regulatory hook on which the regenerative framework hangs. Other statutes that must be honoured include the National Water Act (Republic of South Africa, 1998b) and GN R704/1999 (Department of Water Affairs and Forestry, 1999), the Waste Act 59 of 2008 (Republic of South Africa, 2008), and the Biodiversity Act 10 of 2004 (Republic of South Africa, 2004); engagement with the Department of Mineral and Petroleum Resources (DMPR), the Department of Forestry, Fisheries and the Environment (DFFE) and provincial agriculture/water authorities must begin at design stage.

The execution gap – closure certificates and the rehabilitation fund

The statutory framework is respectable; its execution is not. Three sets of statistics delineate the gap.

First, derelict and ownerless (D&O) mines. The Council for Geoscience documented approximately 6100 D&O sites as of 2021, up from 5906 in 2009. Of these, the Auditor-General South Africa (AGSA) recorded in its 2021/2022 audit 2568 as high-risk, 1170 carrying dangerous mine openings, and only 555 cumulatively rehabilitated under the state-funded programme. The DMPR reported rehabilitating four D&O sites against a target of three in 2024/2025. At the current state-funded execution rate the backlog requires more than a millennium to close.

Second, closure certificates. Watson and Olalde (2019) documented 787 certificates issued 2011 to 2016 (approximately 157 per annum), overwhelmingly for prospecting sites and small operations. Mugunyani (2026) identified fewer than six closure certificates for Mpumalanga mines despite it being the most heavily mined coal province. The closure certificate, the legal end point of a mining

right, is in practice almost never issued for substantive operations. Companies retain their environmental liabilities indefinitely, and the liability cascades into the D&O category whenever they liquidate (Mpanza, Adam and Moolla, 2021).

Third, the financial provisioning fund. Aggregate provisions held by current right-holders are estimated at approximately R60 billion (around R14 billion against Mpumalanga operations). South Africa relies almost exclusively on bank and insurance guarantees rather than ring-fenced cash or trust funds (Matola, 2016). The 2022 collapse of Constantia Insurance exposed the fragility: approximately R45.5 million in Mpumalanga-linked guarantees became concurrent claims with no clear recovery mechanism. The funds are also strictly ring-fenced to the operations that posted them: the R60 billion cannot be used to rehabilitate the 6100 D&O sites, which require R30 to R60 billion of separate state-funded expenditure.

TABLE 1

Headline statistics on the South African mine closure execution gap.

Indicator	Reported figure	Source
D&O mines on the National Register (2021)	~6100	Human Rights Watch (2022); Council for Geoscience (2024)
Of which classified as high-risk	2568	Auditor-General South Africa (2022)
Of which carry dangerous openings	1170	Auditor-General South Africa (2022)
Cumulative D&O mines rehabilitated under state programme to 2021/2022	555	Auditor-General South Africa (2022)
D&O mines rehabilitated in 2024/2025 financial year	4	Department of Mineral and Petroleum Resources (2025)
Closure certificates issued, DMRE, 2011–2016	787 (~157/a)	Watson and Olalde (2019)
Estimated cleanup cost for the 6100 D&O mines	R30–R60 billion	Council for Geoscience (2024); Pulitzer Center (2017)
Aggregate financial provision held by current right holders	~R60 billion	Oxpeckers (2017); Mugunyani (2026)
Financial provision held against Mpumalanga operations	~R14 billion	Mugunyani (2026); Centre for Environmental Rights (2022)

The pattern is unambiguous (Table 1): the D&O backlog is growing; state-funded rehabilitation runs roughly three orders of magnitude slower than the backlog; closure certificates are almost never issued; the funded provisions are segregated, vulnerable, and cannot reach the legacy sites. The framework addresses this along three lines: shifting design philosophy higher in the hierarchy of control; creating a productive rehabilitation pathway whose revenue can substitute for some residual financial provision; and giving the regulator measurable, audited relinquishment criteria.

The Just Transition and the coal phase-down

Since 2021, the Presidential Climate Commission's (PCC) Just Transition Framework (PCC, 2022), the Just Energy Transition (JET) Investment Plan committed at the Conference of the Parties (COP_{26}), and the updated Paris Nationally Determined Contribution (NDC) have placed a defined decarbonisation trajectory on the South African energy system, with material consequences for the coal value chain in Mpumalanga, KwaZulu-Natal and the Waterberg. The PCC has been explicit that closure-driven employment loss is the principal political and social risk of the transition. The framework speaks to this by converting the post-mining footprint from a depleting cost centre into a continuing productive landscape commissioned during rather than after production. Three South

African pilots (Bokamoso Ba Rona, Green Engine, Impact Catalyst) were established to test this substitution logic.

Integrated and progressive closure planning

The international literature has converged on the principle the South African statutes already encode: closure obligations must be embedded in core business, legal and operational processes from pre-feasibility (Limpitlaw and Briel, 2014; Young *et al,* 2022). The International Council on Mining and Metals (ICMM) integrated closure framework, the Western Australian leading-practice guidelines, and the Queensland Progressive Rehabilitation and Closure Plans regime all enforce the same logic. Coaltech (2023) concluded that the South African regime is comprehensive in scope but weak in enforcement and in specification of measurable relinquishment criteria. The gap is execution, not law. Three South African pilots (Bokamoso Ba Rona, Green Engine, Impact Catalyst) demonstrate the practical infrastructure exists at project scale. What is missing is a planning logic that connects the pilot to the strategic mine plan and the orebody model.

Strategic mine planning under uncertainty

Strategic mine planning has moved past deterministic single-net present value (NPV) optimisation. Simultaneous stochastic optimisation integrates extraction sequence, processing destinations, blending and waste handling under geological, metallurgical and market uncertainty, with NPV improvements of 10–25 per cent over conventional sequential optimisation (Dimitrakopoulos, 2011; Dimitrakopoulos and Lamghari, 2022; Jiang and Dimitrakopoulos, 2024). Environmental, social and governance (ESG) variables are now being introduced into the same envelope (Goycoolea, 2025; Kazadi, Cawood and Tholana, 2024). The regenerative framework presented here is the logical next step: post-mining land productivity, carbon balance and rehabilitation cost are state variables that can be modelled and optimised inside the same simultaneous stochastic framework.

SITUATIONAL ANALYSIS – ROOT CAUSES OF THE SOUTH AFRICAN REHABILITATION DEFICIT

Six interacting root causes reproduce the headline symptoms:

1. **Design-stage gap:** rehabilitation is not in the strategic mine plan; it enters cash flow only as a provisioning charge. The framework intervenes by internalising rehabilitation cost, post-mining revenue and carbon balance as block-model attributes.
2. **Legislative gap:** R1147 requires concurrent rehabilitation but does not specify what the land must do; relinquishment criteria are unmeasured. The framework specifies tier-defined end-states with measurable, auditable thresholds.
3. **Financial mechanism gap:** approximately R60 billion held in fragile insurance and bank guarantees; the Constantia collapse exposed the risk; provisions cannot be cross-applied to D&O sites. The framework generates a self-sustaining revenue stream crossing the rehabilitation cost curve at approximately 50 per cent of productive footprint.
4. **Institutional capacity gap:** regulator and operator. DMPR rehabilitated four D&O sites in 2024/2025 against a target of three; operators lack agronomic and biogas engineering capability. The framework defines data and competence requirements explicitly; favours regional shared-capability platforms.
5. **Hierarchy of control gap:** conventional closure operates at the lowest mitigation levels (restoration, offset) after damage is done. The framework shifts closure decisions to design stage and higher hierarchy levels.
6. **Social licence and livelihood transition gap:** closure produces a step change in community livelihood; residual ownership of rehabilitated land is rarely specified. The framework brings productive land into operation during the production phase, smoothing transition.

The six causes are not independent. The design-stage and legislative gaps explain why rehabilitation appears only as a provisioning charge; the institutional and financial gaps explain why provisioning fails to convert into rehabilitated land; the hierarchy gap explains the philosophical posture behind

all of the above; the social licence gap is the consequence that finally makes the others politically untenable. The framework addresses design-stage, legislative and hierarchy gaps directly; the financial mechanism gap via the self-sustaining model; and surfaces the institutional and social licence gaps as explicit prerequisites in Implementation Considerations.

THE REGENERATIVE MINE PLANNING FRAMEWORK

Hierarchy of control as the design philosophy

Two convergent control hierarchies establish the framework's design posture: the occupational hierarchy of control (ISO 45001:2018: elimination, substitution, engineering controls, administrative controls and personal protective equipment (PPE)) and the ICMM mitigation hierarchy (avoid, minimise, restore, offset) (Cross-Sector Biodiversity Initiative, 2013; ICMM, 2020). Both insist that higher-level decisions prevent the need for lower-level interventions, and that lower-level interventions are progressively more expensive, more uncertain and less effective.

Conventional closure operates at the lower levels: restoration after irreversible disturbance, offset of residual impact, administrative compliance. The South African statistics are the consequence compounded over decades. The framework shifts decision-making to the upper levels (Table 2).

TABLE 2

Mapping of the regenerative framework onto the occupational hierarchy of control (ISO 45001:2018) and the ICMM mitigation hierarchy.

Hierarchy level	OHS hierarchy	ICMM hierarchy	Action under the framework
Highest	Elimination	Avoidance	Strategic pit-shell, dump and haul-road design avoiding the most ecologically valuable parcels from pre-feasibility
High	Substitution	Minimisation	Selective topsoil and subsoil handling; hydroponic Tier 2 substitution; biogas CHP substitution for fossil-fired site energy
Medium-high	Engineering controls	Minimisation/ Restoration	Designed-in three-tier differentiation; closed-loop coupling; passive AMD treatment integrated with Tier 3
Medium	Administrative	Restoration	Progressive rehabilitation scheduling; auditable relinquishment criteria; MRV architecture; community agreements
Lowest	PPE/end-of-pipe	Offset	Residual after-the-fact remediation; minimised by upstream design choices

* OHS: Occupational Health and Safety. AMD: acid mine drainage. MRV: monitoring, reporting and verification.

The conventional closure budget expresses the lower-hierarchy posture financially; the regenerative framework's capital and operating expenditure expresses higher-hierarchy investment, which in steady state is cheaper as well as more effective.

Scope and assumptions

The framework applies to surface mines, directly to South African strip and opencast coal (Witbank, Highveld, Waterberg), platinum, chrome, iron-ore, mineral sands and quarry operations; transferable to surface bauxite, copper and gold internationally. It does not currently address tailings storage facilities, sub-surface workings or *in situ* leach operations (though the logic adapts). Climates below approximately 400 mm mean annual rainfall require a redesigned Tier 1 weighted toward dryland systems.

Three-tier land capability differentiation

A rehabilitated mine surface is not uniform. Treating the whole footprint as a single rehabilitation target produces an average outcome that fits no parcel well. The framework allocates each parcel to one of three capability tiers, summarised in Table 3.

TABLE 3

Three-tier land capability differentiation in the regenerative mine planning framework.

Tier	Land condition	Post-mining use	Primary outputs
Tier 1	Backfilled, topsoil reinstated, drainage stable	Rotational livestock grazing coupled to anaerobic digestion	Meat, milk, biogas, digestate, soil organic carbon
Tier 2	Water-constrained or shallow-soil rehabilitated areas	Closed-loop hydroponic production under protected structures	Leafy greens, herbs, recirculated water, continuous biomass
Tier 3	Severely degraded, soil-limited or chemically compromised	Phytoremediation followed by vertical growing on engineered platforms	Stabilised soil, harvested biomass, contained metals, woody carbon

Tier 1 – Rotational grazing and anaerobic digestion

Tier 1 is the productive heart of the post-mining surface. Rotational grazing moves livestock between paddocks on a schedule set by plant recovery; evidence indicates that, under moderate stocking density, this maintains or increases soil organic carbon stocks (Conant *et al,* 2017; Ren *et al,* 2025). Manure is processed in an on-site anaerobic digester, producing biogas (typically 50–75 per cent methane) and nutrient-rich digestate (Atelge *et al,* 2020; United States Environmental Protection Agency, 2025). Biogas runs a CHP unit producing electricity and heat for site use, including Tier 2 pumps and lighting. Digestate is applied to Tier 1 soils and, after balancing, used as partial Tier 2 nutrient input, closing the on-site nutrient and energy loop.

Tier 2 – Closed-loop hydroponics

Tier 2 covers areas where soil reinstatement is shallow, drainage is poor, or surface water is constrained, a description fitting many South African coalfield post-mining surfaces. Recirculating nutrient film technique (NFT) and deep water culture (DWC) systems consume approximately 90 per cent less water per unit biomass than open-field cultivation (1–3 L/kg lettuce versus 30–250 L/kg conventional; Barbosa *et al,* 2015; Pomoni *et al,* 2023; Regmi *et al,* 2024). Vertically stacked systems support 10–20 times the production per unit floor area of open-field horticulture. The energy penalty is offset by biogas-fired CHP and supplemented by site-mounted solar where economic. The hydroponic and grazing loops form a single integrated farm with a single energy and nutrient budget.

Tier 3 – Phytoremediation and vertical growing

Tier 3 is the difficult land: severely degraded surfaces with elevated arsenic, cadmium, lead or other potentially toxic elements, and surfaces where compaction and chemical disturbance preclude conventional revegetation. A two-stage pathway: first, phytoremediation using hyperaccumulator species and plant growth-promoting rhizobacteria to extract, stabilise or volatilise contaminants (Bhat *et al,* 2022; DalCorso *et al,* 2019; Yan *et al,* 2023). Second, once chemistry is stable, vertical growing systems on engineered platforms independent of the underlying soil. Long-lived woody species established as part of the phytostabilisation cover contribute the credit-eligible biomass carbon component.

Coupling between subsystems

The three tiers are not independent silos. Manure from Tier 1 grazing feeds the digester; biogas runs the CHP, supplying electricity to Tier 2 pumps and lighting and Tier 3 vertical systems; digestate

returns to Tier 1 soils as a low-carbon fertiliser substitute and supplies part of the Tier 2 nutrient solution after filtration and balancing; hydroponic water recirculates within Tier 2 with periodic blow-down absorbed by Tier 1 irrigation, closing the on-site water loop in all but the wettest and driest extremes; Tier 3 phytoremediation biomass is processed for safe metal recovery or contained disposal and never returns to the food chain. The system runs as one farm, not three.

SOIL REHABILITATION METHODS

Biological rehabilitation underpins the entire framework. Without restored soil structure, organic matter, microbial activity and chemistry, none of the three tiers can be sustained. Two complementary classes of technique do the work, bioremediation and phytoremediation, both deployed in designed succession rather than as one-shot treatments. Acid mine drainage (AMD) is treated separately because it integrates across all three tiers.

Bioremediation

Bioremediation uses bacteria, fungi, actinomycetes and cyanobacteria to degrade, transform or immobilise pollutants. Mechanisms include biosorption and bioaccumulation, bioleaching, oxidation–reduction reactions, biosurfactant production, biovolatilisation and biomineralisation (Bhat *et al,* 2022; DalCorso *et al,* 2019). Effective taxa in mine soils include Pseudomonas (notably P. fluorescens and P. putida), Rhizobium, Bacillus, Achromobacter species, and several mycorrhizal fungi. Single-strain inoculations seldom hold under field conditions; stable consortia are required (Khanna *et al,* 2022; Hnini, Rabeh and Oubohssaine, 2024).

Phytoremediation

Phytoremediation uses higher plants and rhizosphere microbes via six mechanisms: phytoextraction, phytostabilisation, phytodegradation, phytovolatilisation, phytostimulation (rhizodegradation) and rhizofiltration. For mine soils with heavy metals, phytoextraction and phytostabilisation are most relevant. Approximately 450–500 hyperaccumulator species are characterised worldwide: Pteris vittata (arsenic), Thlaspi caerulescens and Noccaea caerulescens (cadmium, zinc), Berkheya coddii (a South African nickel accumulator), Brassica species (broader range). Native pseudometallophytes are particularly attractive for South African application (Pratas, Favas and Paula, 2014). Limits: low biomass and decade-scale timescales; harvested contaminant-laden biomass requires careful disposal or phytomining. The framework absorbs these limits by Tier 3's sequential pathway and by keeping Tier 3 biomass out of the food chain.

Acid mine drainage as a specific challenge

Acid mine drainage is the dominant long-term water quality risk on South African coal and gold operations. The framework treats AMD as a site-scale water management problem upstream of all three tiers. Passive treatment (constructed wetlands and sulfate-reducing bioreactors) is well documented for low-to-moderate loads and integrates with Tier 3 phytoremediation: anaerobic bacteria reduce sulfate to sulfide and precipitate metals as insoluble sulfides; wetland vegetation polishes effluent. Higher AMD loads require active chemical neutralisation (lime dosing then clarification) upstream. The water management train must be in the closure plan and its energy demand in the framework's energy budget, with biogas-fired electricity contributing.

Combined approach within the framework

Initial bioremediation consortia condition the soil microbiome; phytoremediation species are then established to extract or stabilise residual metals while building organic matter; productive vegetation follows once chemical and physical thresholds are met. Coaltech (2023): relinquishment without measurable criteria is relinquishment without enforcement. The closure plan must publish explicit numerical thresholds, by tier, for residual metal concentrations, soil organic carbon, infiltration rate and vegetation cover, with third-party verification before a tier is declared productive.

CARBON BALANCE OF THE INTEGRATED SYSTEM

The framework does measurable work against atmospheric CO_2: real, multi-pathway, and credit-eligible under the Carbon Tax Act 15 of 2019 (Republic of South Africa, 2019).

Four pathways operate simultaneously

Four pathways are active in the integrated system: (1) direct fixation in standing vegetation across all three tiers; (2) transfer to stable soil organic carbon via root turnover, root exudates, and partial decomposition of above-ground litter and excreta; (3) avoided methane from open manure decomposition, since anaerobic digestion captures the methane and combusts it; and (4) fossil-energy displacement through biogas-fired CHP offsetting diesel pumping and grid electricity.

The pathways reinforce one another: continuous vegetation feeds the manure pipeline that feeds the digester that produces the energy that runs the irrigation that maintains the vegetation. The carbon flux through the integrated system is materially larger than the sum of the fluxes the same parcels would carry under disconnected, single-purpose rehabilitation.

Illustrative hectare arithmetic

For a hypothetical 500 ha lease with 350 ha actively mined and 150 ha progressively rehabilitated (100 ha Tier 1, 5 ha Tier 2, 45 ha Tier 3), the order-of-magnitude carbon arithmetic is shown in Table 4.

TABLE 4

Illustrative order-of-magnitude carbon arithmetic for a 500 ha hypothetical surface mine with 150 ha rehabilitated footprint allocated across the three tiers. Ranges are deliberately wide and must be replaced with site-specific values.

Pathway	Tier	Annual flux (t CO_2e/a)	Credit eligibility under SA Carbon Tax Act
Avoided methane (anaerobic digestion of manure)	Tier 1 + digester	300–900	Eligible (CDM/VCS livestock waste methodology); strong, easily verified
Soil organic carbon under managed grazing	Tier 1	100–400	Eligible up to mineral-associated organic matter capacity; requires baseline sampling
Standing and woody biomass	Tier 3	150–500	Eligible for long-lived woody fraction (afforestation/revegetation methodology)
Fossil-energy displacement (biogas CHP)	Digester + CHP	200–600	Eligible (energy efficiency/ renewable energy); no double-counting with section 12L
Short-cycle biogenic carbon (food crops, pasture forage)	Tier 1 + Tier 2	Not creditable	Net-zero under standard GHG accounting; drives the system but is not the credit basis
Indicative credit-eligible total	All	750–2400	Subject to project-specific MRV and registry validation

CDM: Clean Development Mechanism. GHG: Greenhouse Gas. MRV: monitoring, reporting and verification. VCS: Verified Carbon Standard.

High-density production amplifies the flux

Vertically stacked hydroponic structures on Tier 2 support 10–20 times the plant density of conventional open-field horticulture. Cyclic crop rotation maintains continuous standing biomass and a continuous flow of plant residue into the digester or compost loop. On Tier 1, rotational grazing

follows the same logic: rest paddocks fix carbon while grazed paddocks return manure to the digester.

What counts as a credit, and what does not

Under IPCC 2019 Refinement and ISO 14067:2018, short-cycle biogenic carbon (hydroponic produce, pasture forage harvested and consumed within months) returns to the atmosphere: large, real, but net-zero in standard greenhouse gas (GHG) accounting. The credit-eligible fractions are: avoided methane from digestion of manure; soil organic carbon accumulated under rotational grazing and digestate application (up to mineral-associated organic matter capacity); the one-off increase in standing biomass at first establishment relative to the degraded baseline; long-lived woody biomass on Tier 3 phytostabilisation areas.

A caveat from the recent literature must be recorded. Chang *et al* (2023) showed that, globally, soil carbon sequestration in grasslands cannot offset the warming from methane and nitrous oxide of current ruminant systems considered in isolation (approximately 135 Gt C, roughly twice the current global stock in managed grasslands). This strengthens the design choice to pair grazing with anaerobic digestion: the digester intercepts the methane that would otherwise emerge from open manure decomposition. The grazing-and-digester pairing, not the grazing alone, is what keeps the framework's carbon balance defensible.

Monitoring, reporting and verification

A four-layer monitoring, reporting and verification architecture is required. The first layer is *in situ* instrumentation: continuous biogas flow and composition at the digester, electricity metering at CHP, soil moisture and temperature probes on Tier 1, and weight, growth and yield logging on livestock and hydroponic loops. The second layer is periodic soil sampling for organic carbon stock on a stratified spatial grid, with deep-core sampling per ISO 14064–2:2019 and the relevant offset methodology. The third layer is remote sensing: satellite-derived above-ground biomass and surface temperature monitoring. The fourth layer is third-party verification of baselines, additionality arguments and ongoing performance. The data architecture must be planned alongside the physical design and specified explicitly in the closure plan submitted under R1147.

INTEGRATION INTO THE STRATEGIC MINE PLANNING WORKFLOW

The conceptual case for regenerative design only earns its place in strategic mine planning if its parameters can be expressed in the same mathematical language as the rest of the plan.

Regenerative variables in the surface model

Each surface cell carries five additional attributes: the assigned post-mining capability tier; the time at which the cell is released from active mining to rehabilitation; the rehabilitation cost as a function of starting condition and target tier; the post-mining revenue stream as a function of tier and operating year; and the cumulative carbon balance over the planning horizon, decomposed by pathway. These variables are not exotic; most are already implicit in NEMA closure plans. The change is to bring them inside the optimisation envelope rather than calculate them separately at the end. Where simultaneous stochastic optimisation is already in use, marginal computational cost is small.

Augmented objective function

The augmented objective adds three terms to conventional extraction NPV: a discounted rehabilitation cost at the time a cell is released; a discounted post-mining revenue stream from release to end of planning horizon; and a discounted carbon-credit revenue stream from credit-eligible pathways on the tier. The augmented objective is given by Equation 1.

$$\text{maximise} \quad E\left[\Sigma_t \left(V_t^{EXT} + V_t^{REH} + V_t^{CSR} \right) \cdot (1+i)^{-t} \right] \tag{1}$$

Where:

E[·] = expectation across the joint distribution of geological, market, post-mining productivity and carbon price scenarios

t = planning period (year)

i = discount rate

V_t^{EXT} = extraction net cash flow at time t, summed over blocks B

V_t^{REH} = rehabilitation net cash flow at time t, summed over surface cells S

V_t^{CSR} = carbon credit revenue at time t

The first term is the conventional extraction objective. The second term captures progressive rehabilitation cost and post-mining revenue. The third term captures credit-eligible carbon flux monetised at the prevailing carbon price. The constraint set is enlarged beyond conventional capacity, blending, slope and mining width constraints by tier-assignment, release-timing, water-budget and energy-balance constraints, and the optimisation is solved using the same metaheuristics that already handle large simultaneous stochastic models.

Sequencing and progressive rehabilitation

This is the most important planning lever the framework introduces. Conventional operations defer rehabilitation because reshaping, topsoiling and revegetation are charged to closure provisions only when a cell can no longer be reused. The framework treats every cell as a candidate for early release: the earlier the transition, the longer the cell generates tier revenue and accumulates organic carbon. In weak commodity cycles, the productive value of rehabilitated land shifts the optimal mining sequence, a result conventional planning cannot produce.

Stochastic treatment

Two principal new uncertainty sources are introduced: post-mining productivity (yield per hectare depending on weather, market access and management) and carbon price (depending on tax rate trajectory and section 13 credit limits). Both are handled through scenario distributions in the same simultaneous stochastic framework that already handles grade and metal price uncertainty. The output is a family of robust schedules whose performance is assessed across the joint distribution of geological, market and post-mining outcomes.

An illustrative numerical example

Consider a hypothetical opencast coal operation with a 500 ha lease, a 20-year life-of-mine, a peak active footprint of 350 ha, and a closure liability calculated under R1147 at approximately R150 million undiscounted. Under conventional planning, rehabilitation is back-loaded to years 18–20, the closure provision accumulates as a depreciating financial liability, and the rehabilitated land carries no revenue stream. Under the regenerative framework, the operator commits to a release schedule that brings 7–10 ha per annum into productive Tier 1 to Tier 3 use from year 4 onward. Table 5 sets out an illustrative cost and revenue comparison.

TABLE 5

Illustrative comparison of conventional and regenerative scheduling on a hypothetical 500 ha South African coal operation over a 20-year planning horizon. All figures are illustrative.

Item	Conventional schedule	Regenerative schedule
Closure liability at end of life-of-mine (R million)	~150	~70–90
Capital investment in productive infrastructure (R million)	Nil	~25–40
Annual rehabilitation operating cost in steady state (R million/a)	Charged to closure provision	~3–6
Annual post-mining revenue from years 5–20 (R million/a)	Nil	~8–18

Credit-eligible CO_2e (t/a) once mature	Nil	~750–2400
Community livelihood continuity at closure	Step change	Continuous
Treatment in section 24P financial provision	Full provision	Reduced provision against demonstrated pathway

Two qualitative observations survive any reasonable refinement: the regenerative schedule converts a single end-of-life cash outflow into a multi-decade revenue stream compounding with progressive cell release; the residual closure liability against which financial provision must be made under section 24P is materially smaller. The shift is strongest in the weak-commodity-price scenarios where conventional NPV is itself most fragile.

ECONOMIC AND STATUTORY IMPLICATIONS

Cost and revenue structure

Three effects underpin the economic case. First, progressive rehabilitation with a defined productive end-state reduces the residual liability for section 24P financial provision. Second, rehabilitated land becomes a revenue source through livestock and crop sales, biogas energy displacement, and offset credits. Third, the social licence dividend is real.

Against these stand non-trivial capital and operating costs: topsoil management, livestock fencing and water reticulation, an anaerobic digester, protected hydroponic structures, and monitoring systems. Under R1147 obligations, much of this capital belongs to the closure budget rather than operating expenditure, being demonstrably part of the rehabilitation pathway; this is a conversation operators must open with DMPR and DFFE, where legislative guidance would help. Steady-state annual rehabilitation operating cost (approximately R3 to R6 million for a 500 ha operation) is materially less than back-loaded closure when discounting and contingency are properly accounted.

Carbon offsets under the Carbon Tax Act

Section 13 of the Carbon Tax Act (National Treasury, 2019) permits a taxpayer to reduce carbon tax liability by up to 10 per cent of total greenhouse gas emissions through qualifying offset credits. Eligible methodologies under the section 19(c) regulations currently include the Clean Development Mechanism (CDM), Verified Carbon Standard and Gold Standard projects. Three framework pathways map directly: avoided methane from digestion of manure; soil organic carbon sequestration on rehabilitated grazing land; afforestation and revegetation on Tier 3. Each requires a registered project, third-party verification, and ongoing monitoring. The framework does not by itself generate offset credits; it builds the project boundary and physical system around which one can be designed. For a large emitter the 10 per cent allowance is substantial, with voluntary market credits adding revenue. Double-counting with the section 12L energy efficiency allowance under the Income Tax Act is prohibited.

Statutory alignment in the South African context

The framework is consistent with the direction of South African mining and environmental law: the 2015 Financial Provisioning Regulations require concurrent rehabilitation; the Draft National Mine Closure Strategy of 2021 endorses regional, collaborative closure with substitutive economies; the One Environmental System brings closure under NEMA's broader sustainability mandate. It is also consistent with the Joint Ore Reserves Committee (JORC) Code's underlying principle that resource and reserve reporting must reflect realistic modifying factors, including environmental and social factors (JORC, 2012). A regenerative plan that reduces residual closure risk and supports social licence strengthens the modifying factors underpinning reserve reporting.

Global applicability

The framework is anchored in South African statutes but is jurisdiction-agnostic in principle. Western Australian, Queensland and Canadian provincial frameworks all converge on integrated, progressive, productive closure. Adaptation changes tier definitions (depending on local climate, soil, dominant agriculture, water availability) and the regulatory anchoring of carbon and rehabilitation accounting.

THE SELF-SUSTAINING REHABILITATION MODEL

The illustrative arithmetic of Table 5 indicates that the regenerative framework converts an end-of-life cash outflow into a multi-decade stream of post-mining revenue and offset credits. The next question is whether the revenue stream can become large enough to fund the rehabilitation programme itself, breaking the dependency on a separately accumulated closure provision fund. The answer is yes, beyond a definable threshold, and that result reframes the entire economic argument for the framework.

The funding crossover

The funding profile passes through three time-dependent regimes. In the capital-dominant regime (early years) the programme is in net rehabilitation deficit, with capital committed before any cell is productive and the conventional closure provision still the primary funding instrument. In the transition regime (early to mid years) the first cells come into production, annual revenue is positive but smaller than annual cost, and the deficit is shrinking on a clear trajectory. In the self-funding regime (mid years onward) annual revenue exceeds annual cost, the programme funds its own continuing activity from the cash flow of cells already in production, and no further call on the closure provision fund is needed except for catastrophic or non-routine events.

Crossover arithmetic

Using midpoint figures from Tables 4 and 5, steady-state annual revenue is approximately R13 million across 150 hectares (around R87 000/ha/a net after operating cost); steady-state annual rehabilitation cost is approximately R5 million; capital cost is spread over the first five years at approximately R6 million per annum. The crossover occurs when productive hectares exceed approximately 75, or 50 per cent of the eventual productive footprint. Below the crossover the programme is in net deficit and the conventional closure provision must remain. Above the crossover the programme generates more cash than it consumes, and the surplus is available to fund additional cells in the same operation, retire residual closure provision early, or cross-fund rehabilitation in adjacent operations under a regional shared-platform arrangement.

From mine-level to portfolio-level self-funding

The mine-level crossover understates available leverage. At the portfolio level, capital can be sequenced so that operations past crossover subsidise newer operations in capital-dominant regime. Regional collaboration extends this further: the Impact Catalyst in Limpopo is structured to share closure capability and capital across multiple operators; layering a shared rehabilitation revenue stream over the same coordinating structure produces a regional self-funding platform. In principle, surplus revenue from regenerated active sites could begin to address the legacy derelict and ownerless backlog that no individual operator is responsible for, complementing rather than replacing the state-funded Council for Geoscience programme.

Academic backing for the self-sustaining logic

The funding logic instantiates the circular economy as defined by Geissdoerfer *et al* (2017): residual outputs of one process become inputs to another, eliminating the requirement for external waste sinks. Lèbre, Corder and Golev (2017) extend this to mine site resource management. Costanza *et al* (2014) and the natural capital literature provide the valuation framework in which the rehabilitated land's ecosystem services can be priced and counted as an asset on the operator's balance sheet rather than a liability. The Ellen MacArthur Foundation (2021) and WBCSD (2020) document analogous crossover thresholds in closed-loop industrial systems. Two differences from

generic circular economy applications matter for mining: the disturbance is large and irreversible, and the timescale is multi-decadal. Both argue for internalising the self-sustaining logic at the planning stage.

Two cautions attach to the self-sustaining model. First, crossover figures are derived from order-of-magnitude assumptions and are sensitive to project-specific variables; a real operation must produce its own crossover arithmetic from site-specific yield, cost and price data. Second, the self-sustaining model addresses the recurring rehabilitation cost, not the catastrophic-event or unforeseen-liability cost. The section 24P closure provision must remain in place to cover residual liability that cannot be reasonably foreseen, though its quantum may be revised downward in proportion to the demonstrated self-funding pathway.

IMPLEMENTATION CONSIDERATIONS

Five practical considerations bear on whether the framework is deployable at a real operation.

Capability and competence

The framework requires specialist agronomic, ecological and biogas engineering capability not part of the conventional mining skill set. Companies must hire or contract it; recent South African closure assessments document this capacity gap explicitly (van Wyk and Haagner, 2025). The most efficient medium-term route is shared service arrangements at regional scale, of the type the Impact Catalyst is piloting in Limpopo.

Stakeholder engagement

Productive post-mining land use is durable only with meaningful community stake in design, operation and residual ownership. Co-design of tier allocation with community structures, transparent benefit-sharing arrangements, and credible succession plans for residual ownership are minimum conditions, not optional extras.

Monitoring and data architecture

The framework runs on measurable inputs and outputs at every tier: soil organic carbon stocks, metal concentrations, water recirculation rates, biogas yields, livestock weights, crop yields, carbon balance. Modern sensor networks, satellite-derived biomass measurements and digital twins of the mine surface make this monitoring tractable. The data architecture must be planned alongside the physical design and specified in the closure plan submitted under R1147 with sufficient detail to support the third-party verification required for offset registration.

Risk allocation

Productive rehabilitation introduces market risk (commodity prices for meat, milk, vegetables and carbon credits) into a domain where the historical risk has been technical. Risk allocation between the mining company, post-mining operators and the host community must be specified contractually, and the optimisation must reflect realistic distributions rather than point estimates. The simultaneous stochastic optimisation framework already used for grade and metal price uncertainty handles these additional uncertainty sources naturally.

Pilot before full deployment

The framework is integrated by design but should not be deployed at full scale on the first attempt. The most prudent pathway is a pilot of a single tier coupling on a recently backfilled cell, paired with a household-scale digester and continuous monitoring. Scale-up follows once the pilot data confirm the assumed yields, costs and carbon balance. The Bokamoso Ba Rona, Green Engine and Impact Catalyst projects collectively offer the empirical platform from which such pilots can be designed.

CONCLUSIONS

Mine closure in South Africa is a design problem, not a residual cost problem. The legislation has said so since 2002, but the publicly reported execution statistics make clear that the duty to

rehabilitate is not, in practice, being discharged: 6100 derelict and ownerless sites on the national register; fewer than 200 closure certificates issued per annum at the peak of recent practice; only four state-funded derelict-site rehabilitations completed in 2024/2025; and an estimated R30–R60 billion legacy cleanup cost held outside the financial provisioning system that exists for active operations. These statistics are symptoms of six interacting root causes, of which the most fundamental is the design-stage gap.

The regenerative framework presented in this paper addresses the design-stage gap directly by re-positioning rehabilitation at the upper levels of the engineering hierarchy of control and the ICMM mitigation hierarchy. Three-tier capability differentiation allocates every parcel of the rehabilitated surface to a productive use. Energy, water and nutrient couplings turn the post-mining footprint into a single integrated farm. The carbon performance is real, multi-pathway and credit-eligible under the Carbon Tax Act 15 of 2019. The framework reads naturally into simultaneous stochastic optimisation with five additional surface variables per cell and an augmented objective function capturing progressive rehabilitation cost, post-mining revenue and credit-eligible carbon flux.

The framework's most consequential economic claim is the self-sustaining funding mechanism. At approximately 50 per cent of productive footprint in operation, annual revenue exceeds annual rehabilitation cost, and the programme funds its continuing activity from its own cash flow. Above the crossover, the closure provision can be progressively retired and surplus revenue can begin to address the legacy backlog through portfolio-level or regional shared-platform arrangements. The R60 billion currently held in fragile insurance and bank guarantees is not a stranded asset under this framing; it is a transitional instrument whose quantum decreases as the regenerative pathway matures.

Three streams of work follow. The first is quantitative: a worked numerical case study, anchored to a real South African coal or platinum operation, in which the regenerative variables enter a simultaneous stochastic optimisation model and the resulting NPV, closure liability, carbon balance and self-funding crossover point are compared against a conventional baseline. The second is experimental: pilot-scale field trials of the Tier 1 to digester to Tier 2 coupling on a recently backfilled site, with continuous monitoring of soil organic carbon, water balance and biogas yield over at least three growing seasons. The third is regulatory: formal engagement with the DMPR, the DFFE and the JORC Committee on the treatment of regenerative variables in financial provision calculations and reserve-reporting modifying factors; parallel engagement with the National Treasury and DFFE on registration under section 19(c) offset regulations; and engagement with the Council for Geoscience and the Auditor-General on whether surplus revenue from regenerative active sites can, at portfolio scale, begin to retire a portion of the derelict and ownerless backlog. Together, these three streams move regenerative mine planning from a coherent concept toward demonstrable, audited practice.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the School of Mining Engineering at the University of the Witwatersrand for the academic environment in which this work was developed.

REFERENCES

Atelge, M R, Krisa, D, Kumar, G, Eskicioglu, C, Nguyen, D D, Chang, S W, Atabani, A E, Al-Muhtaseb, A H and Unalan, S, 2020. Biogas production from organic waste: recent progress and perspectives, *Waste and Biomass Valorization*, 11(3):1019–1040.

Auditor-General South Africa (AGSA), 2022. General report on national and provincial audit outcomes 2021/22, Auditor-General of South Africa, Pretoria.

Barbosa, G L, Gadelha, F D A, Kublik, N, Proctor, A, Reichelm, L, Weissinger, E, Wohlleb, G M and Halden, R U, 2015. Comparison of land, water and energy requirements of lettuce grown using hydroponic vs conventional agricultural methods, *International Journal of Environmental Research and Public Health*, 12(6):6879–6891.

Bench Marks Foundation, 2022. Mine closure: the elephant in the room, Bench Marks Foundation, Johannesburg.

Bhat, S A, Bashir, O, Haq, S A U, Amin, T, Rafiq, A, Ali, M, Americo-Pinheiro, J H P and Sher, F, 2022. Phytoremediation of heavy metals in soil and water: an eco-friendly, sustainable and multidisciplinary approach, *Chemosphere*, 303:134788.

Centre for Environmental Rights, 2022. The rehabilitation of derelict and ownerless mines, Centre for Environmental Rights, Cape Town.

Chang, J, Ciais, P, Gasser, T, Smith, P, Herrero, M, Havlik, P, Obersteiner, M, Guenet, B, Goll, D S, Li, W, Naipal, V, Peng, S, Qiu, C, Tian, H, Viovy, N, Yue, C and Zhu, D, 2023. Risk to rely on soil carbon sequestration to offset global ruminant emissions, *Nature Communications*, 14:7430.

Coaltech, 2023. Guideline for agreed relinquishment criteria for mine closure, Coaltech Research Association Report E2020-2, Johannesburg.

Cole, M J, 2024. A mine closure risk rating system for South Africa, *Mining*, 4(1):58–78.

Conant, R T, Cerri, C E P, Osborne, B B and Paustian, K, 2017. Grassland management impacts on soil carbon stocks: a new synthesis, *Ecological Applications*, 27(2):662–668.

Costanza, R, de Groot, R, Sutton, P, van der Ploeg, S, Anderson, S J, Kubiszewski, I, Farber, S and Turner, R K, 2014. Changes in the global value of ecosystem services, *Global Environmental Change*, 26:152–158.

Council for Geoscience, 2024. Derelict and Ownerless Mines Project [online], Council for Geoscience, Pretoria. Available from: <https://www.geoscience.org.za> [Accessed: 9 June 2026].

Cross-Sector Biodiversity Initiative (CSBI), 2013. A cross-sector guide for implementing the Mitigation Hierarchy, CSBI (IPIECA, ICMM and Equator Principles Association), London.

DalCorso, G, Fasani, E, Manara, A, Visioli, G and Furini, A, 2019. Heavy metal pollution: state of the art and innovation in phytoremediation, *International Journal of Molecular Sciences*, 20(14):3412.

Department of Environmental Affairs, 2015. National Environmental Management Act 1998 (Act No 107 of 1998), Regulations pertaining to the financial provision for prospecting, exploration, mining or production operations, Government Notice (GN) R1147. Government Gazette No 39425, 20 November, Pretoria.

Department of Mineral and Petroleum Resources (DMPR), 2025. Annual Performance Report 2024/25, Department of Mineral and Petroleum Resources, Pretoria.

Department of Water Affairs and Forestry, 1999. Regulations on use of water for mining and related activities aimed at the protection of water resources, Government Notice (GN) R704. Government Gazette No 20118, 4 June, Pretoria.

Dimitrakopoulos, R and Lamghari, A, 2022. Simultaneous stochastic optimisation of mining complexes, mineral value chains: an overview of concepts, examples and comparisons, *International Journal of Mining, Reclamation and Environment*, 36(7):443–460.

Dimitrakopoulos, R, 2011. Stochastic optimization for strategic mine planning: a decade of developments, *Journal of Mining Science*, 47(2):138–150.

Ellen MacArthur Foundation, 2021. Universal Circular Economy Policy Goals: Enabling the Transition to Scale, Ellen MacArthur Foundation, Cowes.

Geissdoerfer, M, Savaget, P, Bocken, N M P and Hultink, E J, 2017. The Circular Economy, a new sustainability paradigm?, *Journal of Cleaner Production*, 143:757–768.

Government of Western Australia, 2020. Statutory guidelines for mine closure plans, Department of Mines, Industry Regulation and Safety and Environmental Protection Authority, Perth.

Goycoolea, M, 2025. Co-optimisation of profitability and ESG variables in open-pit and quarry mine planning, working paper, Universidad Adolfo Ibañez, Santiago.

Hnini, M, Rabeh, K and Oubohssaine, M, 2024. Interactions between beneficial soil microorganisms (PGPR and AMF) and host plants for environmental restoration: a systematic review, *Plant Stress*, 11:100391.

Human Rights Watch, 2022. “The Forever Mines”: Perpetual Rights Risks from Unrehabilitated Coal Mines in South Africa, Human Rights Watch, New York.

Intergovernmental Panel on Climate Change (IPCC), 2019. 2019 Refinement to the 2006 IPCC Guidelines for National Greenhouse Gas Inventories (eds: E Calvo Buendia, K Tanabe, A Kranjc, J Baasansuren, M Fukuda, S Ngarize, A Osako, Y Pyrozhenko, P Shermanau and S Federici), Intergovernmental Panel on Climate Change, Geneva.

International Council on Mining and Metals (ICMM), 2020. Mining Principles: Performance Expectations, International Council on Mining and Metals, London.

International Organization for Standardization (ISO), 2018. ISO 14067:2018, Greenhouse gases, Carbon footprint of products, Requirements and guidelines for quantification, International Organization for Standardization, Geneva.

International Organization for Standardization (ISO), 2018. ISO 45001:2018 – Occupational health and safety management systems, Requirements with guidance for use, International Organization for Standardization, Geneva.

International Organization for Standardization (ISO), 2019. ISO 14064-2:2019 – Greenhouse gases, Part 2, Specification with guidance at the project level for quantification, monitoring and reporting of greenhouse gas emission reductions or removal enhancements, International Organization for Standardization, Geneva.

Jiang, Y and Dimitrakopoulos, R, 2024. Simultaneous stochastic optimisation of mining complexes with equipment uncertainty: application at an open-pit copper mining complex, *Mining, Metallurgy and Exploration*, 41(5):2287–2305.

Joint Ore Reserves Committee (JORC), The, 2012. Australasian Code for Reporting of Exploration Results, Mineral Resources and Ore Reserves (The JORC Code) [online], The Joint Ore Reserves Committee of The Australasian Institute of Mining and Metallurgy, Australian Institute of Geoscientists and Minerals Council of Australia. Available from: <http://www.jorc.org> [Accessed: 25 April 2026].

Kazadi, K, Cawood, F T and Tholana, T, 2024. A systems engineering approach to incorporate ESG risks and opportunities in early-stage mine design and planning, *Mining*, 4(3):553–575.

Khanna, K, Kohli, S K, Kaur, R, Handa, N, Bakshi, P, Sharma, P, Ohri, P, Bhardwaj, R, Yousaf, B and Rinklebe, J, 2022. Reconnoitering the efficacy of plant growth promoting rhizobacteria in expediting phytoremediation potential of heavy metals, *Journal of Plant Growth Regulation*, 42(10):6474–6502.

Lèbre, É, Corder, G D and Golev, A, 2017. The role of the mining industry in a circular economy: a framework for resource management at the mine site level, *Journal of Industrial Ecology*, 21(3):662–672.

Limpitlaw, D and Briel, A, 2014. Post-mining land use opportunities in developing countries, a review, *Journal of the Southern African Institute of Mining and Metallurgy*, 114(11):899–903.

Matola, A, 2016. Investigation of Western Australia's rehabilitation fund as a fiscal policy solution for South African abandoned mines, MSc thesis, University of KwaZulu-Natal, Durban.

Mpanza, M, Adam, E and Moolla, R, 2021. A critical review of the impact of South Africa's mine closure policy and the winding-up process of mining companies, *The Journal for Transdisciplinary Research in Southern Africa*, 17(1):a985.

Mugunyani, S, 2026. Ghost mines, deceased children and a R14 billion problem, *Cape Times investigation* [online], 31 May. Available from: <https://capetimes.co.za> [Accessed: 9 June 2026].

National Treasury, 2019. Carbon Tax Act 2019 (Act No 15 of 2019), Government Gazette No 42483. 23 May, Pretoria.

Oxpeckers, 2017. R60-billion held for mines that are never closed, Oxpeckers Centre for Investigative Environmental Journalism [online], 5 May. Available from: <https://oxpeckers.org> [Accessed: 9 June 2026].

Pomoni, D I, Koukou, M K, Vrachopoulos, M G and Vasiliadis, L, 2023. A review of hydroponics and conventional agriculture based on energy and water consumption, environmental impact and land use, *Energies*, 16(4):1690.

Pratas, J, Favas, P J C and Paula, D, 2014. Phytoremediation of soils contaminated with metals and metalloids at mining areas: potential of native flora, in *Environmental Risk Assessment of Soil Contamination* (ed: M C Hernandez-Soriano), pp 485–517 (IntechOpen: Rijeka).

Presidential Climate Commission (PCC), 2022. A Framework for a Just Transition in South Africa, Presidential Climate Commission, Pretoria.

Pulitzer Center, 2017. South Africa: R60-billion held for mines that are never closed [online], Pulitzer Center, Washington DC. Available from: <https://pulitzercenter.org> [Accessed: 9 June 2026].

Regmi, A, Rueda-Kunz, D, Liu, H, Trevino, J, Kathi, S and Simpson, C, 2024. Comparing resource use efficiencies in hydroponic and aeroponic production systems, *Technology in Horticulture*, 4:e005.

Ren, S, Wang, L, Zhou, Y P, Wei, Z J, Zhang, S P and Xin, X P, 2025. More than a decade of moderate grazing: no impact on soil organic carbon stocks and enhancement of mineral-associated organic carbon via livestock diversification, *Global Change Biology*, 31(9):e70466.

Republic of South Africa, 1991. Minerals Act 1991 (Act No 50 of 1991), Government Gazette, Pretoria.

Republic of South Africa, 1998a. National Environmental Management Act 1998 (NEMA; Act No 107 of 1998), Government Gazette, Pretoria.

Republic of South Africa, 1998b. National Water Act 1998 (Act No 36 of 1998), Government Gazette No 19182, 26 August, Pretoria.

Republic of South Africa, 2002. Mineral and Petroleum Resources Development Act 2002 (MPRDA; Act No 28 of 2002), Government Gazette, Pretoria.

Republic of South Africa, 2004. National Environmental Management: Biodiversity Act 2004 (Act No 10 of 2004), Government Gazette No 26436, 7 June, Pretoria.

Republic of South Africa, 2008. National Environmental Management: Waste Act 2008 (Act No 59 of 2008), Government Gazette No 32000, 10 March 2009, Pretoria.

Republic of South Africa, 2019. Regulations on carbon offsets in terms of section 19(c) of the Carbon Tax Act 2019, Government Gazette No 42873, Notice 1556, 29 November, Pretoria.

United States Environmental Protection Agency (US EPA), 2025. AgSTAR: anaerobic digestion on livestock farms [online], United States Environmental Protection Agency, Washington DC. Available from: <https://www.epa.gov/agstar> [Accessed: 25 April 2026].

van Wyk, S J and Haagner, A S H, 2025. A review of mine land rehabilitation outcomes: culture, procurement and practice, *Journal of the Southern African Institute of Mining and Metallurgy*, 125(4):209–216.

Watson, I and Olalde, M, 2019. The state of mine closure in South Africa, what the numbers say, *Journal of the Southern African Institute of Mining and Metallurgy*, 119(7):639–645.

World Business Council for Sustainable Development (WBCSD), 2020. Circular Transition Indicators V2.0, World Business Council for Sustainable Development, Geneva.

Yan, A, Wang, Y, Tan, S N, Mohd Yusof, M L, Ghosh, S and Chen, Z, 2023. Phytoremediation technologies and their mechanism for removal of heavy metal from contaminated soil: an approach for a sustainable environment, *Frontiers in Plant Science*, 14:1076876.

Young, R E, Gann, G D, Walder, B, Liu, J, Cui, W, Newton, V, Nelson, C R, Tashe, N, Jasper, D, Silveira, F A O, Carrick, P J, Hagger, V, Hallett, L, Eisenberg, C, Guariguata, M R and Dixon, K, 2022. International principles and standards for the ecological restoration and recovery of mine sites, *Restoration Ecology*, 30(S2):e13771.

Digital solutions for sustainable mining – an optimisation methodology for pit-limit and extraction-sequence definition to reduce mining impacts and closure liabilities

F Egger[1] and L Fouché[2]

1. Principal Advisor – Strategic Production Planning, Rio Tinto, Perth WA 6000. Email: florian.egger@riotinto.com
2. Manager Strategic Production Planning, Integrated Planning and Scheduling, Rio Tinto, Perth WA 6000. Email: leon.fouche@riotinto.com

ABSTRACT

Over the past two decades, mining legislation and societal expectations of the industry have changed substantially. The industry has responded with a stronger focus on environmental stewardship, community expectations and transparency. Industry-standard mine planning workflows have evolved to reflect these requirements but continue to rely heavily on manual iterations because of limitations in available digital solutions. These iterations typically occur during mine design and production scheduling, where multiple pit and pushback designs, extraction sequences, ex-pit dumps and pit-void backfill options must be assessed to understand the range and sensitivity of mining impacts and closure requirements. This paper presents a reduced impact mining (RIM) digital solution that combines direct block scheduling (DBS) and phase/bench scheduling (PBS) to consider land disturbance and closure commitments while defining pit limits and internal extraction sequences.

The methodology combines block models with environmental, heritage, hydrogeological and closure attributes and applies modified material-flow networks to represent waste-rock placement relative to the water table within released pit-void capacity. Minimum, maximum and ratio constraints prevent waste-rock allocation to backfill until sufficient pit-floor area is exposed. Ground-disturbance, rehabilitation and closure-activity triggers, constraints and costs are linked to the optimisation objective. The resulting DBS schedules guide initial pit-stage and backfill geometry for PBS, where operational constraints are introduced and the practical schedules are reconciled with the strategic DBS outcomes.

Pilot and proof-of-concept work for three anonymised Pilbara applications demonstrated that alternative extraction sequences can release pit-void capacity earlier, increase in-pit backfill and reduce reliance on ex-pit dumping. For Mine A, the DBS RIM Case reduced ex-pit dumping ground disturbance by 31 per cent and truck requirements by 14 per cent, while increasing above water table (AWT) and below water table (BWT) backfill by 105 per cent and 202 per cent, respectively. For Mine B, the DBS RIM Case reduced ex-pit dumping ground disturbance by 79 per cent and truck requirements by 11 per cent, while increasing AWT backfill by 63 per cent. For Mine C, the tested scenarios achieved in-pit backfill rates of approximately 83–92 per cent and combined disturbance reductions of approximately 42–47 per cent relative to fully external waste placement.

The results indicate that DBS-RIM provides a transparent basis for comparing economic, operational, disturbance and closure outcomes before detailed geometries are fixed. The benefits are sequence-dependent and must be confirmed through practical PBS translation, multidisciplinary review and the applicable approvals and consent processes.

INTRODUCTION

Open pit strategic planning has traditionally focused on maximising net present value (NPV) while satisfying geotechnical, processing and production constraints (Newman *et al,* 2010). Lerchs–Grossmann and pseudoflow methods efficiently identify nested economic shells (Lerchs and Grossmann, 1965). However, selecting the ultimate pit, designing pushbacks, assigning material destinations and developing the time-based schedule remain largely separate steps. Although each step is technically rigorous, the overall workflow relies on repeated manual iterations between planning disciplines. Environmental impacts and closure outcomes are therefore often assessed after major spatial decisions have already been made.

The new reduced impact mining (RIM) methodology brings these considerations into the strategic planning process. The objective is to reduce avoidable clearing, external waste landforms and residual void liabilities while maintaining ore delivery, operability and value. This approach responds to stronger environmental, social and governance expectations from regulators, Traditional Owners, communities, capital providers and boards**,** which increasingly align with global frameworks for sustainable development (Jenkins and Yakovleva, 2006; Franks *et al,* 2014; United Nations, 2015). It also supports closure in design by considering final landforms, water, rehabilitation and relinquishment requirements when pit limits, sequences and destinations are selected, rather than adding these requirements after the mine plan has matured (Australian Government, 2016a; ICMM, 2019).

In Australia, expectations for the early, transparent and auditable consideration of cultural heritage in project planning have increased. This requires planning processes that can incorporate heritage values, access conditions, cumulative disturbance and future optionality from the strategic stage, rather than relying only on later design reviews. RIM supports this approach by providing a reproducible optimisation framework through which agreed heritage and environmental conditions can inform scenarios, constraints and decision records. It is intended to complement, rather than replace, engagement, consent processes and statutory assessment.

Direct block scheduling (DBS) was developed within strategic mine planning to address limitations in multi-step optimisation workflows, where relevant constraints and costs may not be considered during stage generation. Rather than taking a fixed pit and pushback sequence as an input, DBS assigns extraction timing and material destinations directly to blocks while enforcing precedence and system constraints. Rio Tinto began investigating DBS in 2018 and added dedicated reduced-impact modelling logic in late 2024. Pilot and proof-of-concept work across three Pilbara hubs tested whether mine/leave, timing, destination, ground-disturbance and closure decisions could be assessed within one optimisation environment and then translated into executable phase/bench scheduling (PBS) plans (Rio Tinto, 2026).

The DBS-RIM methodology described in this paper is the subject of a Rio Tinto patent application. The paper presents the methodology at a level suitable for technical discussion while recognising the associated intellectual property.

INPUTS

The optimisation uses a common block-scale data structure so that value, material movement and mining impacts are calculated for the same decision units. Geological inputs include coordinates, dimensions, density, lithology, weathering, slope sectors, ore/waste classification, grade and confidence classifications. Economic inputs include commodity price assumptions, royalties, mining, processing and haulage costs, discount rates, fixed costs, capital triggers and penalties. These inputs retain the conventional economic basis while allowing reduced-impact outcomes to be represented as costs, benefits or constraints.

Geometallurgical attributes include product yield, recovery, deleterious elements, hardness and blending rules. Hydrogeological inputs include pre- and post-mining water tables, dewatering domains and the distinction between backfill placed above and below the water table. Environmental inputs include approved development envelopes, watercourses, exclusion areas and associated buffers, as well as rehabilitation objectives. Heritage inputs include agreed avoidance zones and conditional-access areas.

Closure inputs include final-landform requirements, minimum backfill commitments, pit-lake assumptions, rehabilitation costs, haulage requirements, topsoil and subsoil inventories, progressive rehabilitation windows and residual-liability factors. Ground-disturbance variables represent newly cleared areas for mining, dumping, infrastructure and soil storage. Each surface cell is activated only once, preventing disturbance from being counted again when deeper blocks are mined. Table 1 summarises the main input groups.

TABLE 1

Principal DBS-RIM input groups and their planning roles.

Input group	Representative attributes	Optimisation role
Geological and geometallurgical	Tonnage, grade, density, lithology, yield, recovery and blending class	Value, precedence, quality and process routing
Economic and operational	Prices, costs, discounting, capacities, access and rate limits	NPV and executable production profile
Water, environment and heritage	Water tables, buffers, habitats, approvals and exclusion domains	Hard constraints, penalties and scenario boundaries
Closure and disturbance	Void capacity, landform work, soil, rehabilitation and clearing areas	Backfill release, liability, progressive rehabilitation and footprint

METHODOLOGY AND PROCESS

The workflow in Figure 1 consists of data preparation, block-model enhancement, DBS optimisation, waste-flow modelling, trade-off assessment and translation to PBS. One or more block models can be flagged with multiple sets of pit extraction limits. The top and bottom blocks within each vertical stack are then identified. The Top Block represents the first physical disturbance associated with mining the stack, while the Bottom Block represents pit-floor exposure and the potential release of void capacity. Assigning these properties only to the top and bottom blocks prevents surface area and closure benefits from being counted repeatedly through the stack.

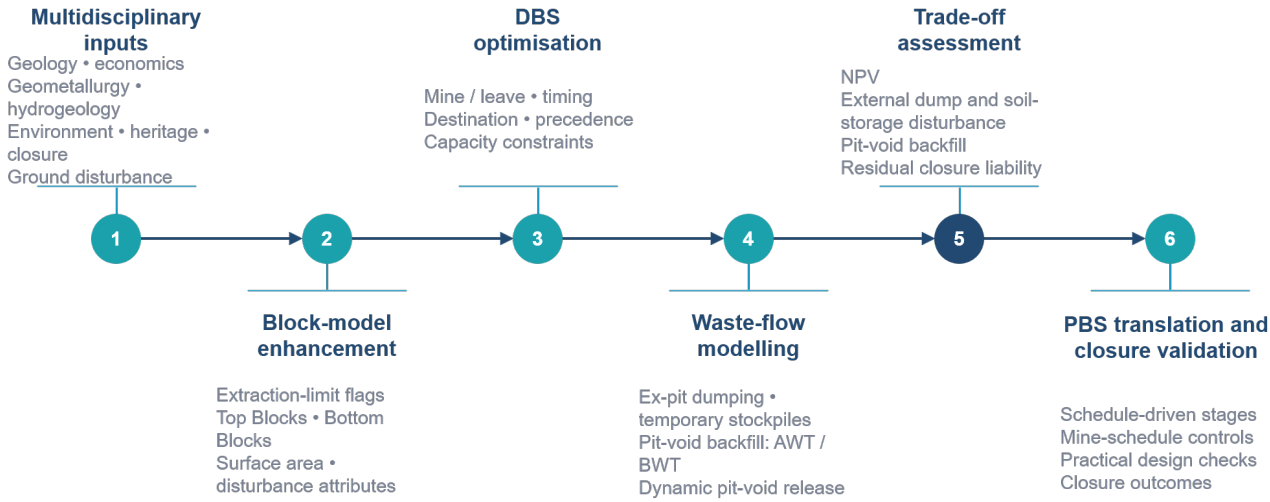


FIG 1 – DBS-RIM workflow from multidisciplinary inputs and block-model enhancement through optimisation, trade-off assessment, PBS translation and closure validation.

Waste movement is represented through a modified material-flow network. Mining blocks direct ore and waste rock to processing facilities, ex-pit dumps, temporary ore and waste stockpiles, or pit-void backfill destinations. Above water table (AWT) refers to waste rock placed within released pit-void capacity above the applicable water table elevation, while below water table (BWT) refers to waste rock placed below that elevation. Material-balance constraints ensure that material is neither created nor lost, while capacity and timing constraints control movement along each flow path.

Pit-void availability changes through time. When extraction reaches a bottom block, a defined quantity of backfill capacity is released. Separate material-flow rules apply where closure commitments distinguish between placement above and below the water table. Several Rio Tinto iron ore mines in the Pilbara have specific closure obligations for the post-mining pit void below the applicable water table elevation. Ratio constraints link released capacity to realised backfill placement and ensure that waste rock is allocated to a pit-void destination only after sufficient pit-floor area has been exposed, as illustrated in Figure 2.

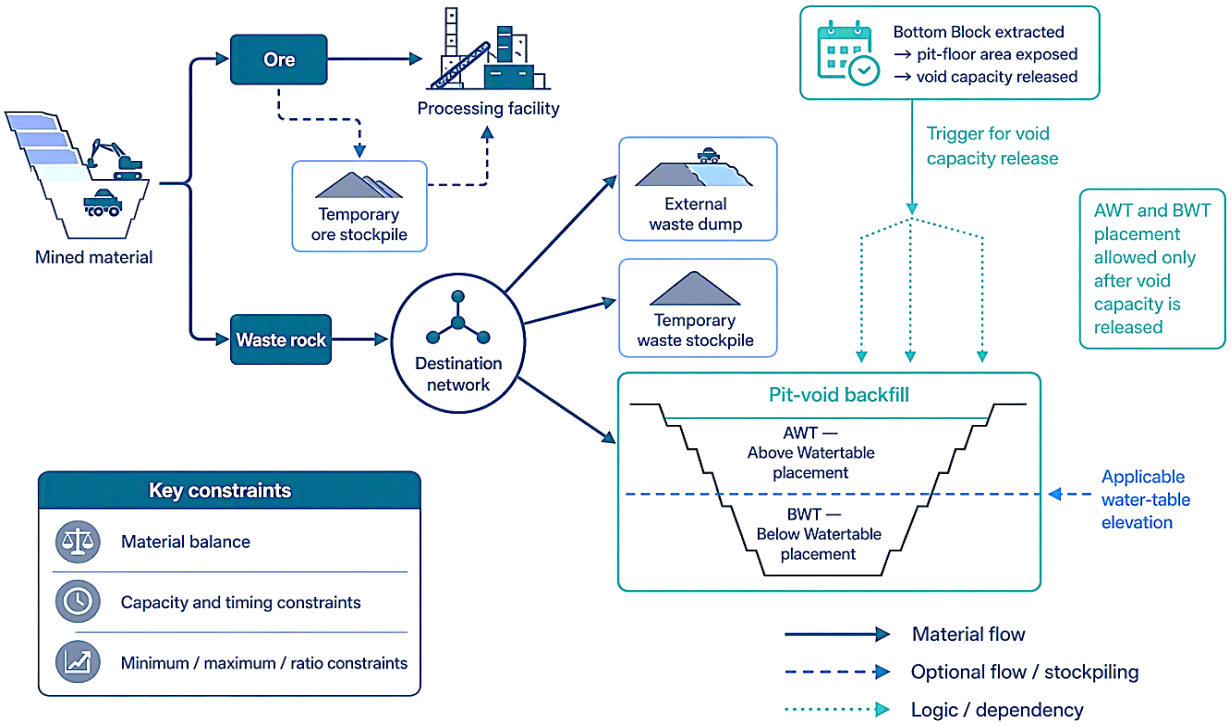


FIG 2 – Material flow network and pit-void-release logic. Generated waste enters the destination network; AWT and BWT placement can occur only after the corresponding pit-floor area and capacity have been released.

Ground disturbance is modelled by activating the relevant surface block for mining areas and the corresponding dump or soil-storage cell for dumping activities. The model can represent disturbance through a hard limit, a period-specific access constraint, a monetary penalty or a rehabilitation credit. In DBS, ex-pit dumps are treated as terminal destinations, while temporary waste stockpiles may be included as modelling objects where subsequent movement is required. Pit-void backfill destinations can be governed by rules for elevation, water table position, material suitability and closure commitments.

The closure-aware objective function maximises discounted economic value and quantified closure and rehabilitation benefits, less costs associated with ground disturbance from external waste dumps and soil-storage areas, additional haulage, rehandling, landform construction and residual liabilities. Hard exclusion areas remain constraints and are not traded against value. Economic-priority, balanced and closure-priority scenarios are generated by varying selected penalties and closure parameters. This produces a transparent trade-off relationship between NPV, ground disturbance associated with external waste dumps and soil-storage areas, pit-void backfill and residual closure liability rather than a single combined sustainability score (Figure 3).

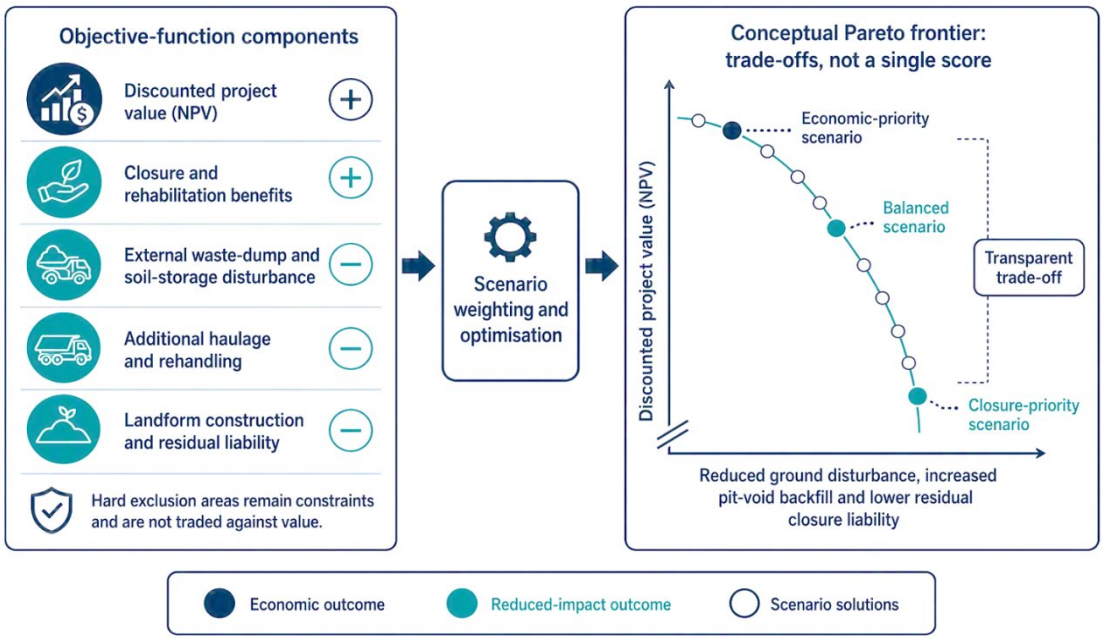


FIG 3 – Closure-aware objective and trade-off assessment. Scenario weights generate a transparent frontier between discounted project value, ground disturbance associated with external waste dumps and soil-storage areas, in-pit placement and residual closure liability.

Direct block scheduling

DBS optimises block extraction decisions through time and can therefore incorporate time-dependent inputs directly. In contrast, pseudoflow workflows generally introduce timing and time-dependent inputs only in a subsequent scheduling step. Each candidate block can be left *in situ* or mined. If mined, the block is assigned a period and destination subject to slope precedence, mining and processing capacities, blending, access, fixed-cost and capital-trigger constraints. Mine/leave, timing and destination decisions are therefore resolved together rather than passed sequentially between pit optimisation, design and scheduling applications.

Conventional Lerchs–Grossmann or pseudoflow pit optimisation determines a maximum-value closure under precedence but does not normally include the full time value of money, multi-period capacities or changing pit-void destinations in the initial shell solution. Nested revenue-factor shells remain useful design guides, but numerous additional considerations affecting the internal extraction sequence can be incorporated only through subsequent manual iterations. DBS uses decomposition and network-based optimisation to evaluate the schedule directly. In the RIM implementation, the same decision space also includes first disturbance, void release, AWT and BWT placement, and closure activities. The output is not a final operational design; it is a high-resolution strategic schedule that shows where value and reduced-impact opportunities can be achieved together.

Phase/bench scheduling integration

PBS converts the strategic DBS result into mineable geometry and a controlled long-term production plan. Blocks with similar extraction periods are grouped and reviewed against bench continuity, minimum mining width, ramp access, geotechnical sectors, equipment deployment and practical waste routes. Changes in mining-front direction, exposure of the final pit-floor and release of backfill capacity are used to identify candidate stage boundaries. The stages are therefore derived from the schedule rather than the schedule being constrained by stages selected before timing and destination trade-offs are understood.

During PBS, additional operational inputs, rules, triggers and constraints are introduced. Preliminary pit-void designs support sequencing rules that coordinate mining and dumping. Ex-pit waste-rock dumps can be modelled as stockpiles from which material may be reclaimed to satisfy closure commitments, whereas pit-void backfill locations remain terminal destinations and may be subdivided by elevation to represent water table considerations. Because land clearing is typically

undertaken in campaigns, ground-disturbance and terminal rehabilitation costs are aggregated through campaign trigger variables.

The PBS schedule scenarios are compared with the DBS reference to determine whether practical design translation preserves the strategic outcomes for ore delivery, value, disturbance, void release and waste placement. Where a design change materially erodes these outcomes, stage boundaries or scheduling parameters are revised. Figure 4 shows this iterative hand-off and validation loop.

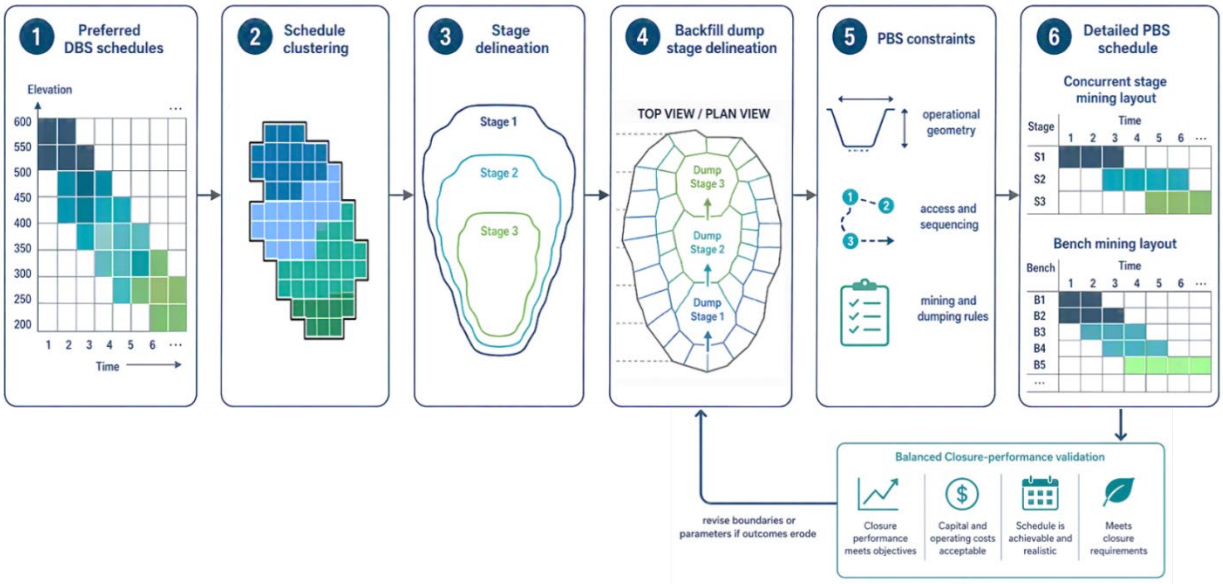


FIG 4 – DBS-PBS integration. Preferred block schedules guide stage delineation; operational geometries and constraints are introduced in PBS; the resulting plan is checked against the value, disturbance and backfill outcomes of the strategic reference.

CASE STUDIES

Three anonymised Pilbara applications evaluate DBS-RIM against site-specific planning baselines across mature, multi-deposit and greenfield settings. The comparisons focus on extraction sequence, waste routing, pit-void backfill, ground disturbance, truck requirements and design optionality. Reported quantitative outcomes are derived from PBS outputs using practical, mineable pit-stage and in-pit backfill geometries.

Mine A

Mine A is a mature Pilbara operation with established pits and substantial ex-pit waste-storage requirements. For a planned satellite mining area, the DBS-RIM assessment identified cutback sequences that exposed Bottom Blocks and released pit-void capacity earlier than the Reference Case. This enabled more direct placement of waste rock in mined-out voids and reduced the manual iteration required to compare pit stages and waste-dump configurations.

Relative to the Reference Case, the DBS RIM Case reduced ex-pit dumping ground disturbance by 31 per cent and truck requirements by 14 per cent, while increasing AWT backfill by 105 per cent and BWT backfill by 202 per cent as shown in Figure 5 (Rio Tinto, 2026). The revised extraction sequence therefore reduced external disturbance and material movement while increasing direct backfill placement.

The methodology also enabled multiple sensitivities to be evaluated with less planner effort and made closure consequences visible during option selection. Earlier completion of mining areas created opportunities for progressive rehabilitation and reduced work deferred to the end of mine life.

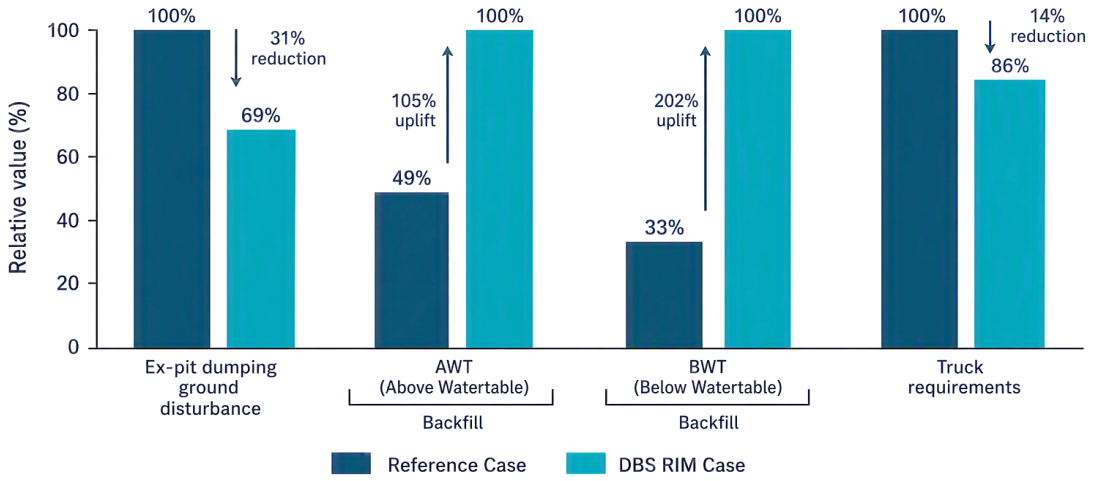


FIG 5 – Mine A relative outcomes for ex-pit dumping ground disturbance, combined AWT and BWT backfill, and truck requirements. Reference Case and DBS RIM Case values are normalised independently within each metric to a 0–100 per cent scale; absolute units differ.

Mine B

Mine B is a large-scale, multi-deposit operation assessing a new mining area within an existing operation, in which all backfill is placed Above Water table. The Reference Case and DBS RIM Case retained the same ore-feed, quality and mining-capacity requirements, while the DBS RIM Case introduced direct block scheduling and reduced-impact controls to test changes in external disturbance, backfill placement and truck requirements.

Relative to the Reference Case, the DBS RIM Case reduced ex-pit dumping ground disturbance by 79 per cent and truck requirements by 11 per cent, while increasing total AWT backfill by 63 per cent as shown in Figure 6 (Rio Tinto, 2026). The results indicate that earlier routing to in-pit destinations reduced new dump disturbance and the operating-hours proxy for truck requirements.

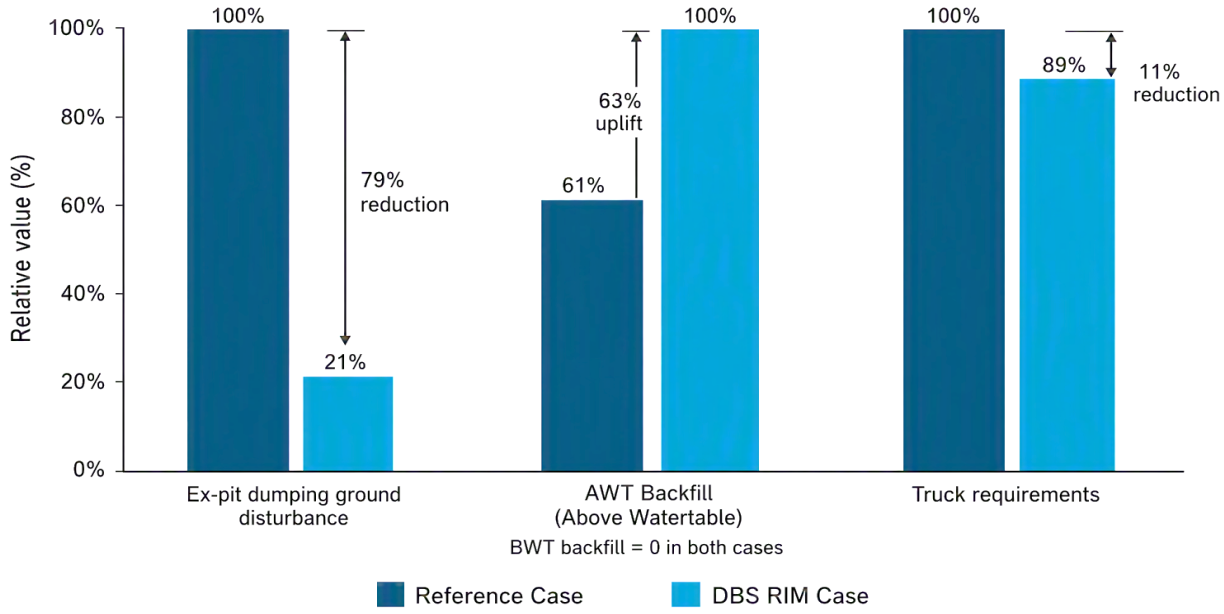


FIG 6 – Mine B relative outcomes for ex-pit dumping ground disturbance, AWT backfill and truck requirements. Reference Case and DBS RIM Case values are normalised independently within each metric to a 0–100 per cent scale; absolute units differ.

The concurrent reduction in ex-pit dumping and closure dump reshaping is consistent with waste being placed earlier in its intended final location.

Mine C

Mine C is a greenfield application in which alternative DBS-RIM schedules were tested before detailed pit-stage and waste-dump geometries were fixed. Ground-disturbance constraints on ex-pit waste dumps and financial penalties for land clearing were varied to assess how schedule-led stage guidance could shift waste from external dumps to in-pit backfill while design options remained open.

Across the tested scenarios, in-pit backfill represented approximately 83–92 per cent of total waste placement. Compared with the scenario using fully external waste placement, combined disturbance from mining, dumping and soil storage was approximately 42–47 per cent lower, while saleable ore production varied by no more than 3 per cent (Rio Tinto, 2026).

These results provide a quantitative basis for subsequent pre-feasibility and feasibility studies, including assessment of backfill sterilisation risk and the preservation of options for approvals and heritage discussions.

DISCUSSION

Table 2 consolidates the principal outcomes from the three case studies. For Mines A and B, DBS-RIM increased in-pit backfill while reducing ex-pit dumping ground disturbance and truck requirements. For Mine C, the tested scenarios achieved high backfill rates and lower combined disturbance with limited variation in saleable ore production. Although the comparison bases differ, the results support a common interpretation: reduced-impact outcomes depend on extraction sequence and the timely release of pit-void capacity before waste is committed to external destinations.

TABLE 2

Summary of reported case study outcomes relative to each baseline.

Case	Disturbance outcome	Backfill outcome	Planning and production outcome
Mine A	31% reduction in ex-pit dumping ground disturbance	105% increase in AWT backfill and 202% increase in BWT backfill	Saleable ore production varied by no more than 2%; 14% reduction in truck requirements; earlier progressive rehabilitation options
Mine B	79% reduction in ex-pit dumping ground disturbance	63% increase in AWT backfill	Saleable ore production varied by no more than 1%, 11% reduction in truck requirements; less closure dump reshaping
Mine C	42–47% lower combined disturbance from mining, dumping and soil storage	In-pit backfill rate of 83–92% of total waste placement	Saleable ore production varied by no more than 3%; outputs support subsequent design studies

Earlier completion of mining areas can reduce work deferred to closure and provide opportunities for progressive rehabilitation, provided the areas are physically complete and unlikely to be re-disturbed. Rehabilitation prescriptions, completion criteria, monitoring and adaptive management remain discipline-led activities rather than optimisation outputs. Uncertainty in water, material availability and post-mining land use should therefore be tested through scenarios (Australian Government, 2016b; ICMM, 2019).

The methodology improves decision traceability by reporting ground disturbance, closure obligations, residual liabilities, ore delivery and economic value within the same scenario comparison. This supports multidisciplinary review of alternatives and cumulative impacts during approvals and stakeholder engagement. Hard exclusions and consent conditions remain binding constraints, while schedule-led staging reduces iteration among optimisation, design, waste-landform and closure teams.

The principal limitations are uncertainty in geological, hydrogeological, environmental and closure inputs; the computational complexity of block-level, multi-destination mixed-integer linear

programming models; and the need to translate strategic patterns into safe, mineable designs. Closure cost estimates and disturbance penalties can create false confidence if treated as precise forecasts. Results should therefore be presented as sensitivities, translated through PBS using practical pit-stage and backfill geometries, reconciled with the strategic DBS outcomes, and reviewed by the relevant technical disciplines.

CONCLUSIONS

The pilot studies demonstrate that reduced-impact considerations can be integrated into strategic optimisation before pit stages and waste destinations are fixed. DBS-RIM evaluates mine/leave, timing and destination decisions together. Top Block and Bottom Block attributes link extraction to first disturbance and pit-void release, while AWT and BWT flow constraints provide an auditable representation of backfill eligibility and placement.

The methodology generates closure-aware strategic schedules that can be translated into practical pit-stage and backfill geometries through PBS, then reconciled with the DBS outcomes. Across the three cases, revised sequences increased in-pit backfill and reduced external disturbance or reliance on ex-pit dumping. Mines A and B also reduced truck requirements, while Mine C achieved substantial backfill and disturbance improvements with limited variation in saleable ore production. The results show that the outcome depends on releasing usable pit-void capacity before waste is committed to external destinations.

DBS-RIM is applicable to new mining areas within existing operations and to greenfield projects. Its principal value is a transparent comparison of alternatives that shows how extraction sequence and waste-routing decisions affect economic value, operability, ground disturbance, approvals, heritage conditions and closure obligations. The optimisation does not replace detailed design, consent processes or discipline-led assessment. Subject to practical translation, multidisciplinary review and appropriate governance, it provides a consistent basis for integrating reduced-impact and closure considerations into strategic mine planning.

FUTURE WORK

Future work will focus on standardising input preparation, calibrating closure-liability factors, strengthening the transfer from DBS to PBS, and testing the methodology across additional orebody and regulatory settings. Longer-term development will extend the objective and constraint set to include interactions among deposits and hubs. These developments should retain separate reporting of economic, operational, environmental and closure outcomes rather than combine them into a single score.

ACKNOWLEDGEMENTS

The authors acknowledge the Rio Tinto Iron Ore mine planning specialists especially Stefan Rohrmoser and the Rio Tinto Iron Ore Technology Delivery and Operations specialists who contributed to the pilot and proof-of-concept reviews across the Pilbara study areas.

Microsoft Copilot, using ChatGPT-based generative AI models available in September 2026, was used to assist with English-language refinement, structural editing, consistency review, caption and table wording, and preparation of the schematic illustrations and relative summary graphs in Figures 1–6. The core ideas, source data, quantitative analysis, technical judgements and interpretations were produced and verified by the authors. The authors reviewed all AI-assisted text and visual outputs, confirmed the reported values against the underlying study information, and retain full responsibility for the accuracy, originality, research integrity and publication rights of the paper. No AI tool is listed as an author or co-author. Confidential and proprietary information was handled only within the authors' organisation-approved Microsoft 365 environment, subject to the applicable service terms and data-protection controls.

REFERENCES

Australian Government, 2016a. Mine closure: Leading Practice Sustainable Development Program for the Mining Industry, Department of Industry, Innovation and Science, Canberra.

Australian Government, 2016b. Mine rehabilitation: Leading Practice Sustainable Development Program for the Mining Industry, Department of Industry, Innovation and Science, Canberra.

Franks, D M, Davis, R, Bebbington, A J, Ali, S H, Kemp, D and Scurrah, M, 2014. Conflict translates environmental and social risk into business costs, in *Proceedings of the National Academy of Sciences*, 111(21):7576–7581.

International Council on Mining and Metals (ICMM), 2019. *Integrated mine closure: Good practice guide*, 2nd edition (International Council on Mining and Metals: London).

Jenkins, H and Yakovleva, N, 2006. Corporate social responsibility in the mining industry: exploring trends in social and environmental disclosure, *Journal of Cleaner Production*, 14(3–4):271–284.

Lerchs, H and Grossmann, I F, 1965. Optimum design of open-pit mines, *Canadian Mining and Metallurgical Bulletin*, 58:47–54.

Newman, A M, Rubio, E, Caro, R, Weintraub, A and Eurek, K, 2010. A review of operations research in mine planning, *Interfaces*, 40(3):222–245.

Rio Tinto, 2026. Reduced Impact Mining and Direct Block Scheduling pilot and proof-of-concept studies, 2025–2026 internal technical reports and workshop materials, Perth.

United Nations, 2015. Transforming our world: the 2030 Agenda for Sustainable Development, United Nations, New York.

Integrating closure in the mine plan – practical ways to reduce risk and cost, and improve ESG outcomes

C Lison[1], R Irwin[2], R Ward[3] and J Heyes[4]

1. Senior Environmental Officer, Department of Mines, Petroleum and Exploration, Perth WA 6004. Email: christine.lison@dmpe.wa.gov.au
2. Team Leader Mine Closure and Technical Advice, Department of Mines, Petroleum and Exploration, Perth WA 6004. Email: rochelle.irwin@dmpe.wa.gov.au
3. General Manager Resource Strategy, Department of Mines, Petroleum and Exploration, Perth WA 6004. Email: rob.ward@dmpe.wa.gov.au
4. Strategic Closure Advisor, Clarity Mine Closure, Perth WA 6000. Email: jheyes@claritymineclosure.com.au

ABSTRACT

Mining is a temporary land use. This temporary nature introduces the challenge of being able to close mining operations in a manner that does not leave an unacceptable liability to the community and future land managers. The integration of closure activities within day-to-day operations and mine plans provides the potential to address the closure challenge early. Closure should not be seen as an extra environmental obligation, it is a planning, scheduling, design and cost-control issue that directly affects mine value, operability, risk and the success of closure transitions. Integrated closure has the strategic benefits of a reduction in environmental and financial risks throughout mining, and social and financial benefits that are more likely to align with closure objectives.

Existing regulatory frameworks expect integration of closure in the mine plan. This is seen as an essential strategy for responsible resource development in Western Australia. Integrated closure is an aspect of responsible resource development that is supported through recognition amongst investors that strong environmental, social and governance (ESG) outcomes align strongly with improved investment returns.

Mine planners make decisions every day that determine whether closure will be routine and cost-effective or fraught with rehandling, redesign and delayed relinquishment. This paper provides practical examples from operating mines in Western Australia that demonstrate effective integration of closure planning and implementation into the mine plan. The examples clearly show how early consideration of closure aspects (eg mined pits, waste streams, landform development, water management, stakeholder engagement, post-mining land use) can improve ESG outcomes through reduction of risk, operational complexity and future cost associated with closure liabilities.

INTRODUCTION

Why closure integration is a mine planning issue

An understandable key factor in the failure to integrate closure is the reluctance for mine operators to accept additional expenditure in the near-term to reduce future closure costs results (Shaw, Pedlar-Hobbs and Chubb, 2022). However, it is well established that the benefits associated with integrated mine closure planning and activities can be substantial. Shaw, Pedlar-Hobbs and Chubb (2022) argue that closure efforts can be made significantly more effective (and less expensive) if closure considerations are included in the day-to-day planning and operations of a mine. Where integration of closure activities is undertaken in a progressive manner, the value of a mine in Net Present Value (NPV), perhaps counterintuitively, can increase (Churr, McNulty and Posnik, 2014). Other immediate benefits include faster completion of closure landforms and reduced risk of acid rock drainage (ARD) (ICMM, 2025). Hattingh and Bothma (2013) discuss how upfront closure planning, guided by well-defined end-use planning, provides opportunities to implement cost-effective solutions for the post-mining landscape during operations, whilst creating lasting opportunities for communities. MacKenzie *et al* (2013) provides an example of how selection, segregation and strategic stockpiling of competent waste rock during operations, for use in later closure activities, can realise financial and operational efficiencies. Fawcett and Laurencont (2019)

claim that having clear closure objectives before operations begin can help prevent a mine becoming locked into perpetual maintenance to manage ongoing impacts post closure.

Closure integration completed too late creates rework and destroys optionality. It can lock mine plans into sub-optimal landform geometry, increase haulage and rehandle, create deficits in beneficial materials required for capping – all of which can result in deferred relinquishment and an associated extension of liability. Sanders, Slingerland and Murphy (2024), noted that closure cost estimates have often focused only on direct, quantifiable costs, overlooking uncertainties such as escalation, contingency, and risk. Without integrating closure considerations into mine planning and design maturity, these unknowns are not adequately managed, resulting in operations facing significant and often unexpected cost increases (Figure 1). The dependency of the closure success on integration of closure planning highlights that closure integration must be a focus of mine planning and not just a technical afterthought.

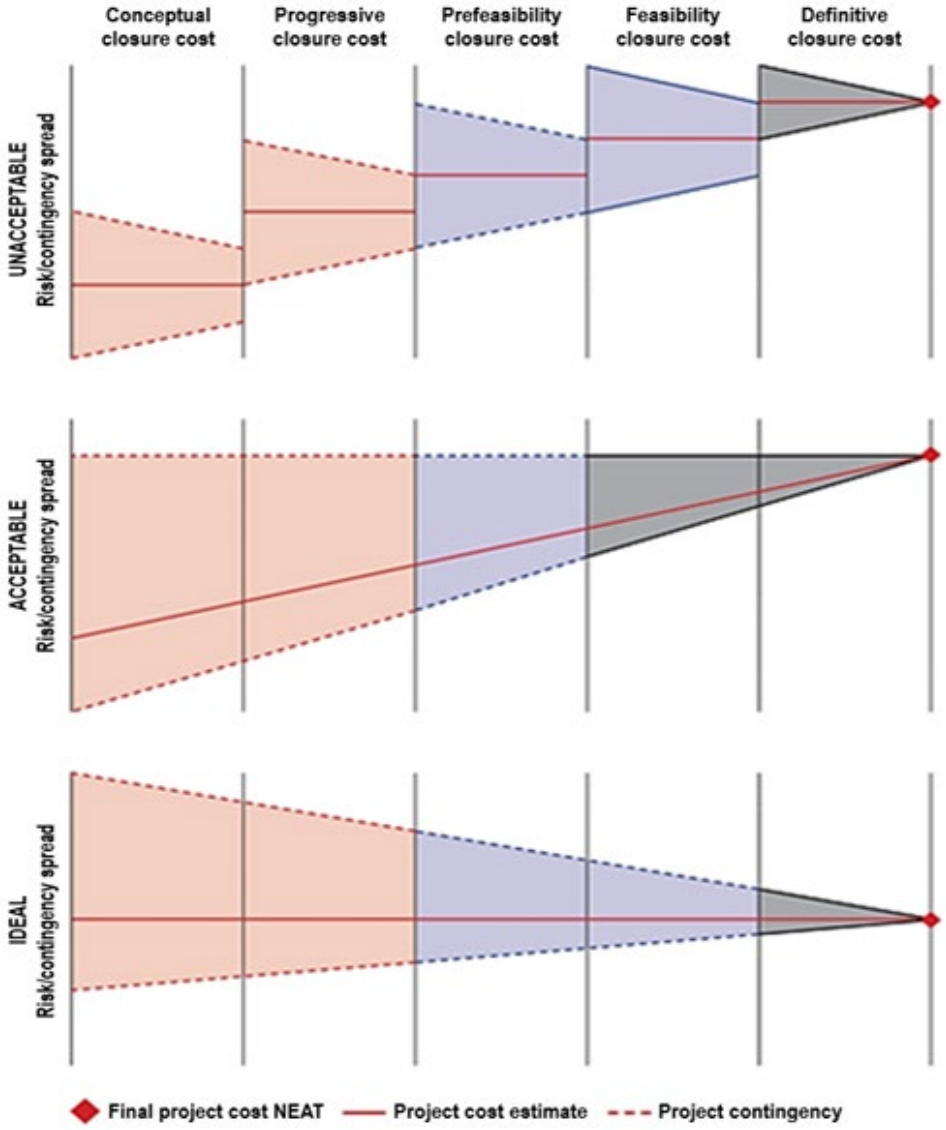


FIG 1 – Comparison of likely project cost history based on consideration of cost uncertainty (Source: Sanders, Slingerland and Murphy, 2024).

The issue of operational complexity often stems from the multi-disciplinary nature of closure activities and may make it harder to achieve integrated closure. For example, management of acid metalliferous drainage (AMD) requires consideration of waste geochemistry, landform design, water management, water quality, and environmental engineering. When these aspects are not considered at all, or recommendations to manage risk are not clearly communicated to mine planners, the opportunity to simultaneously manage AMD risk alongside waste scheduling during operations may

be missed. Delaying management of AMD will result in deterioration of water quality or land leading to contamination. This can significantly increase future closure costs where remediation, environmental monitoring and management may persist for years to decades beyond mine closure. Understandably, long-term management of contaminated land and water is a liability that land managers are unwilling to take back, potentially leaving the mine owner with no formal mechanism to relinquish (Sommerville and Ferguson, 2022).

The complexity of closure is also influenced by stakeholder expectations. Closure planning seeks to accommodate the needs of stakeholders, and these needs may change over time. Mine owners may seek to delay stakeholder engagement or agreements on post-mining land use (PMLU) until the end of mining to avoid complexities associated with changes in stakeholder expectations. However, deferral may increase the risk that the proposed PMLU is mis-aligned with stakeholder expectations. Further costs may be accrued at end of mining if the mine owner is required to seek alternative designs or closure strategies to address stakeholders needs. The misalignment may also lead to development of mistrust in stakeholders which may delay or stall future engagement and potentially cause harm to the company's reputation.

Regulatory and ESG drivers for integrated closure in Australia

The Department of Mines, Petroleum and Exploration (DMPE) considers early closure planning and progressive implementation (integrated closure) as an essential strategy for responsible resource development in Western Australia. DMPE's regulatory framework promotes an integrated closure approach to preparation of mining development and closure proposals and mine closure plans. These are required for a mine to receive environmental approval and remain compliant under the *Mining Act 1978*. Guidance prepared by the Department under the *Mining Act* identifies 'up front' consideration of closure (ie during the feasibility stage), stakeholder engagement, identification of closure risks and adaptive management as essential components of closure planning.

Whilst integrated mine closure isn't always explicitly discussed within ESG frameworks, it is fundamentally linked to all three factors. The environment and communities are most impacted by what is left behind after mining ceases, and decision-making for how mine closure eventuates is effectively a governance matter in managing risk and financial liability. Recent literature emphasises that decisions made by companies that impact on ESG factors can influence whether investors choose to invest in a company (Finucane and Bibby, 2023) and a businesses' overall success (Dowse and Woods, 2025).

The international mining industry has adopted numerous frameworks for reporting on sustainability and ESG factors (Finucane and Bibby, 2023). Australian industry bodies encourage integration of ESG into business (AMEC, 2023), or alignment with international frameworks such as the Global Reporting Initiative (GRI), the International Sustainability Standards Board (ISSB) Standards or International Council for Mining and Metals (ICMM, 2024) *Mining Principles.* The ISSB Standards are partially incorporated in the Australia jurisdiction through the Australian Accounting Standards Board (AASB) Australian Sustainability Reporting Standards S2 which requires climate-related disclosures to form part of an entities' annual reporting obligations (AASB, 2024).

The Global Industry Standard on Tailings Management (GISTM) also sets clear expectations for integrated closure, requiring operators to plan for closure and post-closure from the earliest stages of mine development (ICMM, 2020). Compliance with GISTM is increasingly seen as a benchmark for responsible mining and is referenced by regulators, investors, and industry bodies.

MINE PLANNING DECISIONS THAT SHAPE CLOSURE OUTCOMES

To understand opportunities to improve integration of mine closure into operational mine planning, the authors sought examples of mines that demonstrated commitment to optimisation of closure outcomes through early consideration of closure aspects, early stakeholder engagement, adaptive management and progressive rehabilitation. Interviews were held with representatives from three mining companies to determine where closure can be practically integrated in alignment with mine planning regimes and the drivers behind business decisions. Data included in mine closure plans, annual environmental reporting and sustainability reporting was used to support information provided in interviews. Informed prior consent was obtained from each mining representative for the use of

their information. The authors observations were evaluated and synthesised into five key opportunities that closure integration presented, which are summarised herein.

Reduce overall closure liability

All three mining companies had a clear mandate to integrate closure into day-to-day operations with the goal of reducing future costs associated with the overall closure liability and maintain good closure governance practices to ensure closure plans remain consistent with mine plans.

Completion of early closure studies including comprehensive materials characterisation allowed one company to identify mineral waste that had potential to generate AMD. The development of conceptual encapsulation/cover designs and early materials reconciliation showed that significant volumes of non-acid forming (NAF) material would be required for encapsulation of potentially acid forming (PAF) material. This assessment identified significant cost implications would arise if mining and haulage of NAF materials were required to continue for an extended period post-closure. To minimise this cost a business decision was made to fully integrate closure activities within mine operations planning including stockpiling of NAF material alongside landforms in preparation for rehandle at closure.

For another company, the requirement to integrate closure was underpinned by a recognition of the costs associated with the rehabilitation of large areas of disturbance. Completion of initial progressive rehabilitation was used to develop statistics on time frames and costs associated with completion of each hectare of rehabilitation. These statistics were used to refine the closure cost and provide justification for use of progressive rehabilitation to reduce the closure liability. Annual provisioning for rehabilitation and support from Senior Management helped to achieve prioritisation of progressive rehabilitation within the business. In addition, optimising mine void backfilling opportunities throughout operations has resulted in a reduction in an estimated 70M m^3 of mineral waste required to be rehandled post closure, equating to a cost saving of approximately AUD187M (based on a bottom up, first principles calculation method using NSW Government (2025) standardised rate of AUD2.68/m^3 for loading and hauling rock assuming large dozer (D10, D11) with haulage distance >2000 m).

Avoid future rehandling, redesign and operational inefficiency

The decision to stockpile run-of-mine NAF capping material during operations for subsequent encapsulation of PAF materials avoided significant future non-productive sourcing of appropriate materials. Were materials not available at closure, a mining and rock haulage fleet would need to be maintained post-cessation of operations for sourcing of NAF material from either the mine's underground operations or a new borrow pit. Early identification of this risk (insufficient quantities of NAF material for PAF encapsulation to prevent AMD) and early estimation of the overall closure liability allowed the mine company to recognise that future closure costs would be far greater than the additional operating costs associated with stockpiling.

Development of a stockpiling strategy for NAF materials also ensured PAF and NAF volumes could be continuously tracked and reconciled as operations progressed. Reconciliation of material volumes could then be verified against design specifications for conceptual closure designs such as cover design thickness, to ensure specifications were able to be practically achieved. Improved certainty in operational procedures to achieve design specifications helped to avoid the need for cover systems to be re-designed to accommodate for uncertainty in material volumes. This strategy also helped to increase stakeholder confidence that design controls to reduce environmental risks would be met. Consequently, operational performance was more likely to be driven by proactive reduction in risk and uncertainty as opposed to compliance-driven actions imposed by regulatory authorities.

One mining company identified that there was a high risk of re-work of final landforms if AMD and stability risks in mineral waste were not planned for and effectively managed during operations. A simple approach adopted was to utilise data from geological resource models, such as base of oxidation, lithology and sulfur values, to enable waste management classification within early stage mining models. Three basic classes were developed for high PAF risk, high erodibility risk, or beneficial materials. These three classes were than able to be included in mine planning

optimisations at the earliest stages of development. As improved data became available from infill drilling programs and blasthole data, the coding and mine plans were refined. But this early stage view of risk and management requirements, enabled risk based schedules and waste dump designs to be developed.

Reduced residual risk

A common position was that the integration of closure reduced environmental and social residual risks associated with degradation of land, water and biodiversity (by erosion and/or contamination). The reduction of residual risk, in particular for AMD, avoided significant costs, including the need to re-activate fleet services for remediation and reclamation works, significant cost overruns due to not being able to construct covers in scheduled time frames or treatment of water quality in perpetuity. These costs are in addition to the potential legal and reputational risk that may arise from mining companies not being able to meet their legal obligations and agreed closure outcomes.

Improved geological definition enabled stockpiling of beneficial mineral waste into the mine plan, which provided opportunities to verify design assumptions, improving confidence that landforms would meet closure outcomes and hence improve chances of timely relinquishment. Through progressive rehabilitation, techniques were refined over time to ensure that the final land use as agreed with stakeholders, was realistic and achievable.

For one mine site, stakeholders flagged access to land, maintenance of natural creek flows and visual amenity as being important to them at closure. Post-mining access to land was expedited through optimisation of the mine plan, to incorporate pit backfill opportunities. This enabled earlier availability of landforms for rehabilitation, with the added benefit of improvements in visual amenity.

Early integration of creek re-instatement into the mine plan provided the opportunity to ensure natural creek flows were well understood and able to be maintained at closure. The technique for creek re-instatement was optimised through trials to identify the parameters required to maintain regular flows. This required an investment in use of mine fleet for strategic identification and placement of rock for scour protection. The insight from the trials increased the amount of rehabilitation that could be formally confirmed to meet agreed closure outcomes, and maximised stakeholder access to land post-mining.

Improve confidence in the pathway to relinquishment

In Western Australia a mine's legal obligations may be relinquished where there is sufficient evidence to demonstrate that their environmental obligations have been met and the mine does not leave an unacceptable liability to the community or future land managers. One mining company gained formal sign-off for about 2000 hectares of rehabilitated land within ten years of construction completion by using a planned program of progressive rehabilitation. This work was supported by environmental performance targets and contractor payment milestones during construction, and by Key Performance Indicators (KPIs) during operations.

KPIs for short- and medium-term plans set minimum targets for the number of hectares to be rehabilitated and the amount of rock armour needed for landform rehabilitation. The mining manager was responsible for rehabilitation KPIs, and the geology manager was responsible for rock armour KPIs. These targets helped build a workplace culture that prioritised integrated closure. Combined with lessons from progressive rehabilitation, they also helped improve rehabilitation methods, support agreement on the PMLU, and build a strong understanding of material characterisation. Together, these outcomes created a clear pathway to relinquishment.

Greater confidence in the pathway to relinquishment helped the company meet the expectations of regulators, investors and stakeholders while keeping closure work cost-effective. Managing stakeholder expectations also reduced the risk of disagreement about closure outcomes as the mine neared the end of its life. For example, future discussions were more likely to focus on refining closure outcomes rather than changing them completely.

Create opportunities for stakeholders or future land users

For some mining companies, integration of closure was supported through early agreement with stakeholders on the PMLU. For example, refinements to final landform designs in mine plans enabled the creation of small recreational areas, for use by Traditional Owners to reconnect with country and use rehabilitated land for camping.

An unexpected benefit of early engagement is that it enabled mining companies to gain a better understanding of opportunities for leveraging stakeholder involvement and benefits. Opportunities to foster future economic activity within the surrounding community were identified through creation of local employment opportunities or retaining mining infrastructure post-mining for future economic activities. At one site, up to two thirds of progressive rehabilitation is completed by earthworks fleets belonging to Traditional Owners.

Improved geological definition allows for information on mineral waste properties and tracking to be shared with stakeholders or future land users. The availability of this knowledge empowers stakeholders to make informed decisions on future land use through lenses of risk and opportunity. This knowledge may also be used to support future decisions related to development projects based on residual value of mineral waste. The creation of new technologies to extract value from mined waste and tailings has become a greater focus in recent years with secondary prospectivity readily aligned with sustainability frameworks.

IMPLICATIONS FOR MINE PLANNERS

The potential for closure to be integrated in day-to-day operations was shown to be dependent on the completion of key activities prior to the start of operations; agreed PMLU, conceptual closure designs for landforms, waste scheduling and segregation requirements to mitigate closure risks. This requires input and integration into the mine plan from many disciplines to understand the technical and social risks and requirements for closure, underpinned by robust internal and external stakeholder engagement. Involvement of all key decision-makers in optimisation of the mine plan, helped to ensure that this information was made available to inform the development of a closure ready mine plan.

Whilst activities such as materials characterisation, haulage and stockpiling of cover materials poses an additional operational cost, the cost required to undertake these activities as closure approaches – when optionality is restricted – is far greater. The holistic financial value of integration of closure can only be understood where closure costs were estimated for the full life-of-mine and mine plans were advanced enough to identify the opportunities and constraints of business decisions.

Mine planning optimisations seek to deliver the best NPV for a project. The NPV is incomplete if it fails to adequately include closure costs. Comparative analysis of various mine planning scenarios is needed to achieve this, including a ‘do-nothing’ closure scenario, ie defer all closure activity until closure. This approach enables the Net Present Cost of proactive ‘closure in design’ decisions to be evaluated. Opportunities can then be advanced where there are financial and operational synergies (eg grade control) available to reduce closure costs.

Sommerville and Heyes (2009) described how by rethinking the pit mining strategy at a Pilbara iron ore mine, mine planners could unlock significant operational and closure benefits. For example, instead of mining separate circular pits and dumping waste externally, which would require additional land disturbance, the adoption of moving-face approach would allow for earlier backfilling. This method reduced haulage costs by 10 per cent, enabled sustainable closure by filling pits back to the water table, and lowered geotechnical and hydrological risks by minimising the strike length of openings.

Mine planning opportunities are not limited to long-term plans. There are many examples were progressive rehabilitation efforts, whilst well-intended, have failed due to inadequate planning. In the short and medium term, progressive rehabilitation can be optimised through the establishment of ‘gates’ to determine when and where progressive rehabilitation can proceed. Appropriate encapsulation of problematic materials, verification of resource sterilisation and quality control of final slope design and execution, are examples of key criteria or gates that can be used to establish

when progressive rehabilitation is ready to commence. This requires collaboration between mine planning, operations and rehabilitation teams.

BENEFITS OF CLOSURE INTEGRATION INTO MINE PLANS

The integration of closure activities within day-to-day operations provides the potential to address closure risks early, reducing the holistic cost of closure and improving confidence in the pathway to relinquishment. A reduction in environmental and social risks and improved governance through development of internal procedures (eg for waste characterisation, material handling and separation, waste scheduling and landform design/construction) is closely linked to achievement of ESG outcomes.

Table 1 presents an analysis of integrated closure activities that can be integrated with mine planning, the benefits associated with each and the potential ESG outcomes that could be achieved.

TABLE 1

Summary of closure activities that can be integrated within mine planning, the benefits associated with each and the potential ESG that can be achieved.

Integrated closure activity	Description	Benefit	ESG[1,2]
Progressive rehabilitation	Analyse multiple mine plan scenarios to identify areas available for progressive rehabilitation during operations. Rehabilitation can include re-profiling of landforms, application of surface covers, seeding or backfilling.	Reduced overall environmental impact/liability. Stakeholder and regulator expectations are met. Secured pathway to relinquishment.	Social/environmental: • Assess and systematically manage environmental and social risks[1]. • Achieve a minimum of no net loss or net gain of biodiversity by completion of closure[1]. Governance: • Establish systems to maintain compliance[1].
Optimisation of mine plan	Comparison of different mine planning scenarios to ensure opportunities for integration are not lost.	Reduced overall closure liability. Identification of ways to offset additional operational expenditure such as shorter haul routes.	Social/environmental: • Plan for closure[1]. • Improve energy efficiency and reduce greenhouse gas emissions[1]. • Address climate related risks and opportunities[2].
Improved geological definition	Increased assaying or characterisation of mineral waste during operations	Reduced overall environmental impact/liability. Support future decisions for land use. Ensure availability of materials for cover/encapsulation (or growth media)	Social/environmental: • Assess and systematically manage environmental and social risks[1]. Governance: • Integrate sustainability into decision-making[1].
Waste scheduling and segregation	Verification of assumptions of conceptual designs against reconciled volumes	Avoid future rehandling and redesign. Reduce residual risks for management of problematic materials.	Environmental: • Implement water stewardship practices[1]. • Prevent pollution[1].
Early agreement on PMLU	Secured alignment with stakeholder expectations	Improved closure outcomes. Identify opportunities to support local economy. Stakeholder and regulator expectations are met.	Social: • Contribute to community development[1]. • Support local economic opportunities[1]. • Conduct local stakeholder engagement[1].

1 When compared against ICMM (2024) *Mining Principles.*

2 When compared against AASB Australian Sustainability Reporting Standard S2 (2024) *Climate-related Disclosures*.

CONCLUSION

Mine planners are central to delivering successful mine closure. By integrating closure into planning from the start, planners help avoid costly rehandling, reduce long-term liabilities, and ensure a smooth transition to post-mining land use. Their decisions about material movement, landform design, waste scheduling, and stakeholder engagement create value that extends beyond compliance.

Holistic value means more than meeting regulatory and ESG investor requirements. It includes strengthening the mine's reputation, supporting community development, and creating opportunities for future land users. When planners embed closure thinking into everyday operations, they help build trust with regulators, investors, and communities, and contribute to outcomes that benefit both the business and society.

The future of mining depends on planners who see closure as a chance to add value through and beyond the operational horizon, manage risk, and deliver lasting benefits. With closure integration as a core part of planning, mine planners can lead the way in responsible resource development.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the contribution of IGO and Hancock Iron Ore who generously provided input into preparation of this paper.

REFERENCES

Australian Mining and Exploration Companies (AMEC), 2023. Guide for Mineral Explorers on ESG (Environment, Social, Governance) [online]. Available from: <https://amec.org.au/environment/> [Accessed: 24 May 2026].

Australian Sustainability Reporting Standard (AASB), 2024. S2 Climate-related Disclosures, Commonwealth of Australia.

Churr, A, McNulty, C and Posnik, S, 2014. Reducing closure costs and risks: a cost-benefit analysis of mining and rehabilitation methods, in *Proceedings of the Ninth International Conference on Mine Closure* (eds: I M Weiersbye, A B Fourie and M Tibbett), (University of the Witwatersrand: Johannesburg).

Department of Mines, Petroleum and Exploration, 1978. *Mining Act 1978,* December 1978.

Dowse, J and Woods, I, 2025. ESG Unlocked – How successful companies and investors build our natural, social, human and financial capital, 203 p (De Gruyter: Berlin).

Fawcett, M and Laurencont, T, 2019. Setting objectives: the key to successful mine closure, in *Proceedings of the 13th International Conference on Mine Closure 2019* (eds: M Tibbett and A B Fourie), pp 1063–1070 (Australian Centre for Geomechanics: Perth).

Finucane, S J and Bibby, R A, 2023. Closure disclosure – mine closure planning and reporting through an ESG lens, in *Proceedings of the 16th International Conference on Mine Closure 2023* (eds: B Abbasi, J Parshley, A B Fourie and M Tibbett), pp 1–16 (Australian Centre for Geomechanics: Perth).

Hattingh, R and Bothma, J, 2013. Taking the risk out of a risky business: a land use approach to closure planning, in *Proceedings of the Eighth International Conference on Mine Closure 2013* (eds: M Tibbett, A B Fourie and C Digby), pp 15–24 (Australian Centre for Geomechanics: Perth).

International Council on Mining and Metals (ICMM), 2020. Global Industry Standard on Tailings Management (GISTM) (International Council on Mining and Metals, London).

International Council on Mining and Metals (ICMM), 2024. Mining principles performance expectations (International Council on Mining and Metals, London).

International Council on Mining and Metals (ICMM), 2025. *Integrated Mine Closure Good Practice Guide*, 3rd edition (International Council on Mining and Metals, London).

MacKenzie, S, Beattie, T, Burne, N and Haymont, R, 2013. Case study – integrated planning to enhance closure outcomes for the Pardoo Mine in Western Australia, in *Proceedings of the Eighth International Conference on Mine Closure 2013* (eds: M Tibbett, A B Fourie and C Digby), pp 31–40 (Australian Centre for Geomechanics: Perth).

New South Wales (NSW) Government, 2025. Rehabilitation Cost Estimate Tool – Mining New South Wales, NSW Government.

Sanders, J, Slingerland, N and Murphy, D, 2024. Closing the gap: closure cost estimation trends and pathways to improved maturity, in *Proceedings of the 17th International Conference on Mine Closure 2024* (eds: A B Fourie, M Tibbett and G Boggs), pp 1149–1162 (Australian Centre for Geomechanics: Perth).

Shaw, J, Pedlar-Hobbs, R and Chubb, D, 2022. The role of key performance indicators throughout the mine life in achieving closure objectives, in *Proceedings of the 15th International Conference on Mine Closure 2022* (eds: M Tibbett, A B Fourie and G Boggs), pp 803–812 (Australian Centre for Geomechanics: Perth).

Sommerville, K and Ferguson, K, 2022. Let's reimagine our legacy of mining, in *Proceedings of the 15th International Conference on Mine Closure 2022* (eds: M Tibbett, A B Fourie and G Boggs), pp 3–18 (Australian Centre for Geomechanics: Perth).

Sommerville, K and Heyes, J, 2009. Be Open to Closure: It can save you money, in *Proceedings of the Iron Ore Conference 2009,* pp 317–322 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Commodity prices and capital allocation

Strategic capital allocation methods for commodity cycle resilience

E Deucker[1] and J Chung[2]

1. Director, AurumIQ, Perth WA 6015. Email: lizdeucker@gmail.com
2. Product Manager – Optimisation, Deswik, Perth WA 6000. Email: joyce.chung@deswik.com

ABSTRACT

Effective capital allocation remains one of the most consequential strategic challenges for mining companies. Capital constraints, long project lead times, commodity price volatility, and rising stakeholder expectations require investment decisions that balance financial performance with resilience and long-term strategic alignment. Often the focus centres on maximising net present value (NPV), a single-point estimate that collapses time, capital efficiency, and price sensitivity into one number. Portfolios constructed by ranking projects with an emphasis on NPV tend to cluster similar asset archetypes, front-load capital, and concentrate exposure to price cycles.

This paper explores a two-stage framework for building mine portfolios that deliver strong performance across the commodity price cycle. The first stage of the proposed framework employs a simulated multi-asset producer portfolio with development-stage growth options. It then applies Pareto-based multi-objective optimisation across NPV, free cash flow margin and capital commitment to generate an efficient frontier of portfolio candidates. The second stage evaluates this frontier using a regime-switching simulation engine. This engine models structural gold price archetypes, including extended supercycles and abrupt price dislocations. The framework is extended to incorporate grade-tonnage dynamic sensitivity analysis.

A case study portfolio constructed from published feasibility data is used to evaluate the framework against traditional NPV-led ranking approaches. Portfolio performance is assessed across simulated price regimes using measures of capital efficiency, cash flow stability, and downside resilience.

The proposed framework is computationally tractable and directly applicable to real investment decisions. It reframes mine portfolio construction as a resilience problem and gives executive teams a structured process for aligning capital deployment with long-term cycle exposure.

INTRODUCTION

Mining investment commits large capital over long lead times to assets whose financial outcomes will be shaped by evolving commodity prices, cost structures and market conditions. The gap between when a mining investment is committed and when it delivers its returns can span decades and multiple price cycles. As a result, investment decisions made under a single price assumption can materially shape corporate performance for decades. This time horizon transforms portfolio construction from a financial optimisation problem into a resilience problem; the primary objective is not to identify the highest-return projects available today, but to build a portfolio that performs well across a broad range of conditions.

Traditional rate-and-rank evaluation methods remain the most widely used in the capital allocation process in mining. These approaches commonly assess projects in isolation using financial metrics such as net present value (NPV), internal rate of return (IRR), payback period or capital intensity. While these metrics are useful for comparing project economics, they provide limited insight into how combinations of assets perform as a portfolio across different price cycles.

This limitation becomes particularly important in the cyclic commodity market. The consequences of this limitation have been discussed and documented in the previous literature, including Pidun and Stange (2017), Nagel *et al* (2014), Chandra and Hartley (2024), Espí and De La Torre (2013), Mareels, Moore and Vainberg (2020), Zadeh (2025) and Fisher and D'Alessandro (2021). Van der Bijl and Tholana (2020) identify a persistent pattern of pro-cyclical investment across major diversified miners, with capital concentrated at price peaks and withdrawn at downturns. They concluded that this pattern leads to poor returns, asset write-downs, and sustained share price underperformance. Static NPV-led ranking can reinforce this behaviour as it favours large, capital-

intensive projects that appear compelling under current or consensus price assumptions but whose performance across different price regimes is neither measured nor optimised.

To address this problem, this paper develops and tests a two-stage cycle-aware capital allocation framework that addresses these limitations. The proposed framework combines Pareto-based multi-objective optimisation (MOO) with regime-switching commodity price simulation that evaluates portfolio candidates across structural price environments rather than at a single assumed price. The analysis also examines dynamic asset-level price sensitivity, revealing that the deposit archetype systematically determines price asymmetry in ways that are invisible to static feasibility models.

The framework is applied to a case study portfolio of 14 development-stage gold projects drawn from published feasibility data and anchored to a base portfolio representing the existing operations of a major producer. The case study compares the proposed framework with traditional NPV-led ranking and assesses portfolio performance using measures of capital efficiency, cash flow stability, and downside resilience. The results show that projects that appear marginal under static NPV ranking may still provide material strategic value when considered as part of a cycle-resilient portfolio.

LITERATURE REVIEW AND RESEARCH GAP

Few industries face the combination of capital intensity, timeline exposure and price cycle uncertainty that characterises mining investment, yet the literature suggests these characteristics are rarely reflected in how allocation decisions are actually made.

Pidun and Stange (2017) argue that capital allocation is frequently treated as a financial planning task rather than a strategic capability, with firms allocating reactively based on past performance rather than directing investment toward future value creation and long-term cycle resilience. The consequence is that project-level metrics drive decisions that should be made at portfolio level.

Issues with capital allocation identified in the literature stem from:

- NPV as the dominant evaluation metric.
- project level focus creating imbalances at portfolio level.
- pro-cyclical investment behaviour.
- undervaluing optionality.
- price assumptions.
- cut-off grade theory limitations.

The dominance of NPV as a single decision metric

NPV remains the primary evaluation metric in mining capital allocation, despite well-documented limitations. NPV collapses variations from uncertainties in future project cash flows into a fixed expected project value today, and ignoring future uncertainties and contingencies can lead managers to make incorrect binary decisions (Chandra and Hartley, 2024). In the mining context, this is particularly problematic given the length of asset lives and the magnitude of price uncertainty over project horizons.

Portfolio construction at the project rather than portfolio level

The emphasis on NPV compounds into portfolio-level problems. Traditional rate-and-rank methods, though simple and widely understood, assess projects in isolation rather than taking a portfolio-level view. This can result in over-similarity among prioritised investments, which, when compounded within the portfolio, can lead to poor capital sequencing, underappreciated risk exposures, or dilution of strategy.

Nagel *et al* (2014) identify this problem specifically in mining, using the Lonmin case, noting that rate-and-rank methods produce project clustering, suboptimal capital sequencing, and difficulty in accounting for project interdependencies. The evolution toward portfolio optimisation based on efficient-frontier approaches has demonstrated measurable improvements in capital efficiency and decision transparency.

Pro-cyclical investment behaviour

Van der Bijl and Tholana (2020) demonstrate that firms concentrate capital expenditure during periods of elevated commodity prices, resulting in poor financial returns, asset write-downs, and value destruction when market conditions deteriorate. The behavioural drivers of pro-cyclicality are well established. Optimistic price assumptions during boom periods (Zadeh, 2025), competitive pressure to deploy capital when peers are active, and constrained investment capacity during downturns combine to create a structural bias toward cycle-peak investment. A PwC survey identified the single worst driver of returns in mining as over-optimistic commodity price predictions over the mine's life, cited by 41 per cent of respondents (Regueiro y González-Barros and Espí, 2019).

McKinsey industry analysis (Mareels, Moore and Vainberg, 2020) confirms the scale of this problem; market capitalisation and capital expenditure cycles are strongly correlated with commodity prices resulting in chronically low price-to-book ratios and volatile investor confidence. This dynamic creates a compounding problem: companies are unable to capitalise on high prices because they underinvested during the downcycle, and become overextended at cycle peaks through excessive expansion. The authors recommend through-cycle investment discipline, including saving capital during upcycles, continuing to invest during downturns, dynamically reallocating resources, and debiasing internal decision processes.

Failure to incorporate managerial flexibility and optionality

Surveys of mining industry practice indicate that methods mainly ignoring the real option nature of projects remain the most commonly used project valuation techniques; namely static discounted cash flow methods such as NPV and IRR (Savolainen, 2016; Chandra and Hartley, 2024). The consequence is that projects offering embedded flexibility through staged development, expandable infrastructure or geological extensions are systematically undervalued relative to fully committed single-pathway developments.

Price assumptions

Conventional mine planning processes with scenario price decks (often paired with Monte Carlo simulation) offer some insight into the strategic options available under assumed conditions but do not quantify how investments respond when conditions differ from assumptions. Price is used as a proxy for risk, but it is a poor one: a single assumed price embeds no information about the distribution of outcomes or the sensitivity of the investment to price variation. Conservative price decks as a risk management approach reduce project valuations and can demotivate quality investments (Espí and De La Torre, 2013, cited in Vilhelmsson, 2022); yet in industry surveys price optimism is among the leading self-reported causes of capital destruction in mining, motivating investments that fail to return capital (Regueiro y González-Barros and Espí, 2019; van der Bijl and Tholana, 2020).

The decision about which price scenario to follow remains an executive judgement call, unaided by any analytical basis for quantifying upside capture relative to downside risk. Conventional spider chart sensitivity analysis flexes price on a fixed mine plan and offers the illusion of rigour without resolving the underlying problem.

As Tobisova, Senova and Rozenberg (2022) discuss, simulation strengthens risk awareness in investment decisions. Fisher and D'Alessandro (2021) highlight that upside and downside risk both matter in assessing investments.

Cut-off grade theory

Traditional cut-off grade theory, following Lane's framework, assumes a continuous and near-linear relationship between commodity price and optimal cut-off. As price rises, breakeven grade falls and lower-grade material becomes economic. However, this approach conflates breakeven grade with optimal grade and fails to account for the non-linear relationship between grade, tonnage, margin, and capital in real orebodies.

This gap between theoretical cut-off grade optimality and static feasibility study practice represents an underexplored dimension of capital allocation risk. Mine configurations locked in at planning prices may systematically misallocate capital across the price cycle, regardless of portfolio construction method.

Current literature gap

Taken together, the literature identifies NPV-led ranking that compresses complex behaviour into one metric, project-level evaluation that misses portfolio interaction, structural pro-cyclicality, undervalued optionality, and price-sensitivity analysis that fails to reveal how deposit archetypes respond across the cycle – but no integrated framework that combines portfolio optimisation, price-regime testing, and asset-level dynamic price sensitivity in one method. This paper addresses the research gap of a practical capital allocation framework that combines portfolio optimisation, price-regime testing, and asset-level dynamic price sensitivity. To address this gap, the paper focuses on three questions:

1. What should we build?
2. Will it perform across the cycle?
3. When should we deviate from the plan?

METHODOLOGY

The research develops and tests a cycle-aware capital allocation framework using a combined literature review, cash flow modelling, and portfolio simulation. The approach was structured in three layers of increasing analytical complexity, each addressing a limitation identified in the prior layer: the literature review above, cash flow models built from published feasibility data for the 14 case study projects (benchmarked against a conventional rate and rank baseline), and an MOO plus simulation layer that tests portfolio resilience and asset-level dynamic sensitivity. Figure 1 summarises the full workflow; Table 1 lists the key parameters and calibration values used at each stage.

Stage 1 builds on prior MBA research (Deucker, 2025) and optimises the case study portfolio, anchored to a base portfolio representing a top 10 producer's existing operations, simultaneously across NPV, free cash flow margin, and capital commitment under the five-year capital-constrained horizon, generating a Pareto efficient frontier validated against five rate and rank portfolios. Stage 2 evaluates this frontier under the regime-switching price simulation in Table 1, combined with grade tonnage (GT) curve adjustment as a proxy for dynamic mine configuration response to price. Because adding objectives to a Pareto optimisation expands rather than filters the efficient frontier, a robust Pareto framework runs the MOO across six objective combinations and retains only portfolios that sit on most fronts simultaneously: a small set that is hard to beat regardless of how the problem is framed.

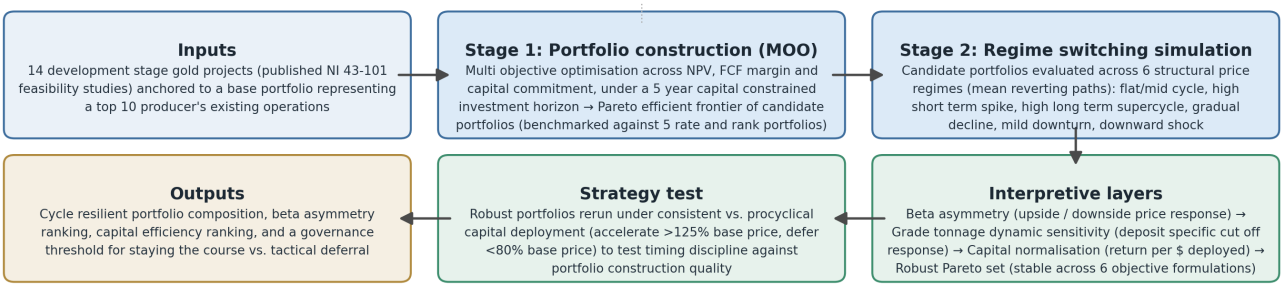


FIG 1 – Framework overview.

TABLE 1

Key model parameters and calibration.

Parameter	Value/definition used in this paper
Investment decision horizon	Five years
Case study assets	14 development-stage gold projects (published NI 43-101 feasibility data) plus a base portfolio representing a top-10 producer's existing operations
MOO objectives (Stage 1)	NPV, free cash flow margin, and capital commitment, optimised jointly under the five-year capital constraint
Rate-and-rank benchmark	Five portfolios built with conventional weighted scoring (value-, quality- and balance-weighted variants)
Price regimes (Stage 2)	Six mean-reverting structural paths: flat/mid-cycle (base), high short-term spike, high long-term supercycle, gradual decline, mild downturn (≈75 per cent of base), downward shock (≈50 per cent of base)
Beta asymmetry split point	Each portfolio's own median simulated price present value (not a fixed price point), so the threshold is consistent within each portfolio's outcome distribution
GT dynamic model cost response	Variable costs scale with ore throughput as cut-off grade shifts with price; fixed costs held constant, with the variable/fixed split calibrated to each deposit's published cost structure
Robust Pareto formulations	Six combinations of two- and three-objective frontiers tested; portfolios sitting on most fronts simultaneously are retained as 'robust'
Pro-cyclical acceleration trigger	Simulated price >125 per cent of base assumption → capital budget +30 per cent and the next queued project is pulled forward one year
Pro-cyclical deferral trigger	Simulated price <80 per cent of base assumption → all uncommitted project starts for that year deferred one year; a project pushed beyond the five year horizon is permanently dropped

Two interpretive layers extract further insight from the simulation. The first interpretive layer applies beta asymmetry analysis to the simulation results, quantifying how sensitive the portfolio NPV responds to price increases versus price decreases. Beta asymmetry measures how a portfolio responds differently to favourable and unfavourable key driver (commodity price) movements. Traditionally, in mining, sensitivity analysis often assumes a project's value moves in a broadly linear way as commodity prices change; if price falls, NPV decreases by a similar proportional amount, and *vice versa*. In practice, mining assets do not always behave symmetrically; their response depends on multiple key drivers such as grade distribution. Cost structure, commodity price, cut-off grade, market conditions and others. Following Fisher and D'Alessandro (2021), upside and downside beta are calculated separately across the simulation distribution. Their ratio, beta asymmetry, finds assets and portfolios that most capture price upside while limiting downside exposure. Capital normalisation further extends the analysis, revealing that the apparent price leverage of large-scale development programs largely reflects portfolio scale rather than capital efficiency.

Finally, a strategy test compares consistent versus procyclical investment timing on the robust portfolios. The test explores whether responding to observed price conditions within the investment decision horizon produces better or worse outcomes than maintaining a consistent capital deployment schedule. This directly addresses the pro-cyclical investment behaviour documented by van der Bijl and Tholana (2020). The analysis reveals a crossover effect: consistent investment outperforms pro-cyclical timing across most price scenarios, but tactical capital preservation improves per-dollar downside resilience in severe price dislocations. The threshold lies between a

moderate downturn, where intervention consistently destroys value, and a severe shock where deferral reduces capital committed into deteriorating conditions.

CASE STUDY, MODELLING AND RESULT ANALYSIS

While individual projects are typically evaluated on financial metrics such as NPV, capital allocation decisions ultimately shape enterprise performance at the portfolio level. The objective of portfolio construction is therefore not to maximise individual project value, but to assemble investments that produce a resilient and balanced financial profile across commodity price cycles.

This case study examines project selection under constrained capital decisions. Three strategies were tested in sequence, each addressing a limitation identified in the prior stage:

- First, rate-and-rank and MOO results are compared to establish whether portfolio construction method materially affects financial outcomes.
- Second, price deck-based selection is contrasted with beta asymmetry under regime-switching simulation, with the analysis extended to incorporate dynamic grade-tonnage sensitivity and capital efficiency measures.
- Third, consistent and pro-cyclical investment timing strategies are evaluated to test whether timing discipline improves cycle-adjusted performance.

Together, these tests build a progressively richer picture of what cycle-resilient capital allocation requires and what conventional evaluation practice systematically fails to provide.

Rate and rank project selection versus MOO selection – what should we build?

The first question is answered by comparing conventional rate-and-rank selection against MOO on the same case study data set. A rate-and-rank baseline was first tested, from strategic features spanning:

- Value (NPV, Free Cash Flow (FCF))
- Quality (FCF margin, All in Cost (AIC)/oz)
- Returns (Return On Invested Capital (ROIC), NPV/Capex (capital expenditure))
- Environmental Social and Governance (ESG) (CO_2e/oz)
- Geography (Fragile States Index (FSI) score)
- Scale (koz/a)
- Resilience (metal mix).

These were scored by min-max normalisation and re-weighted to reflect different strategic priorities.

Two economic project archetypes emerged from the data set: high-volume (HV) projects, typically open pit mines with high NPV and free cash flow but high capital needs and lower margins; and high-grade (HG) projects, typically underground, with the opposite profile (Figure 2). These archetypes are visually distinct in a plot of FCF margin versus NPV and form the basis for the portfolio diversification analysis that follows.

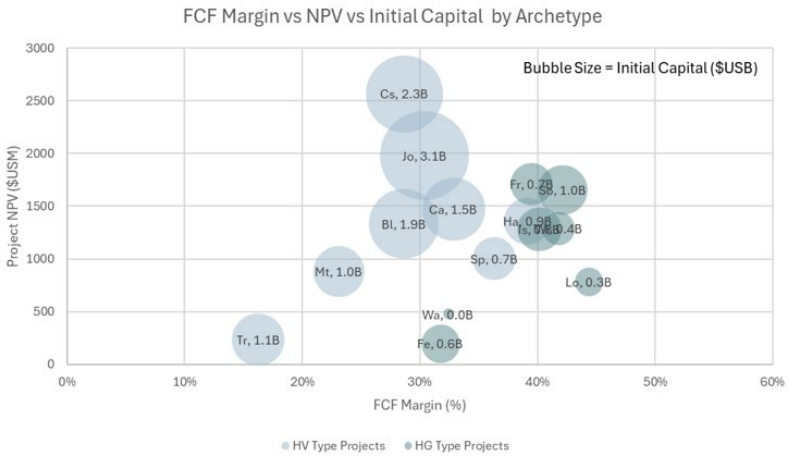


FIG 2 – FCF margin versus NPV versus initial capital by project archetype.

The projects were then ranked against different strategic preferences (margin driven, NPV driven, balanced etc) The method successfully prioritised projects aligned with each strategic weighting (Figure 3). Consistent with the literature, several limitations emerged. Results cluster by archetype, creating systematic bias as highlighted by Nagel *et al* (2014). Emphasising NPV favours large, capital-intensive projects with deferred returns; emphasising margin skews toward smaller, high-grade assets with shorter mine lives and higher geological risk. Concentrating either archetype in a portfolio compounds similar cash flow and risk profiles, undermining balance precisely where diversification is most needed. Metric sensitivity was also significant: substituting NPV for FCF, or production rate for mine life altered rankings materially, revealing the fragility of decisions to seemingly minor choices. Weight elicitation, typically reliant on stakeholder consultation, introduces further subjectivity (López-García, Liern and Pérez-Gladish, 2023).



FIG 3 – Rate and rank weighted scores – quality preference, value preference.

To test a portfolio-level approach, the cash flow model data set was imported into Python and an iterative algorithm was written to combine projects into portfolios under a five-year decision horizon and annualised capital constraint, reflecting the real-world condition that options do not remain available indefinitely. A deduplication filter was applied to remove permutations that yielded identical capital-constrained outcomes, reducing the candidate set from over six billion order-sensitive permutations to approximately 15 500 unique portfolios – a tractable yet representative sample achieved through a list-folding technique. Pareto-optimal portfolios were identified across three objectives: NPV, FCF margin, and capital commitment.

The MOO approach successfully identified portfolios where no alternative combination yields higher NPV or margin without a corresponding trade-off in capital. For validation, the rate-and-rank results were passed through the same algorithm (Figure 4); while one result was close, the Pareto approach consistently outperformed financially, confirming its efficiency in solving what is essentially a capital-allocation knapsack problem (Mathews, 1896).

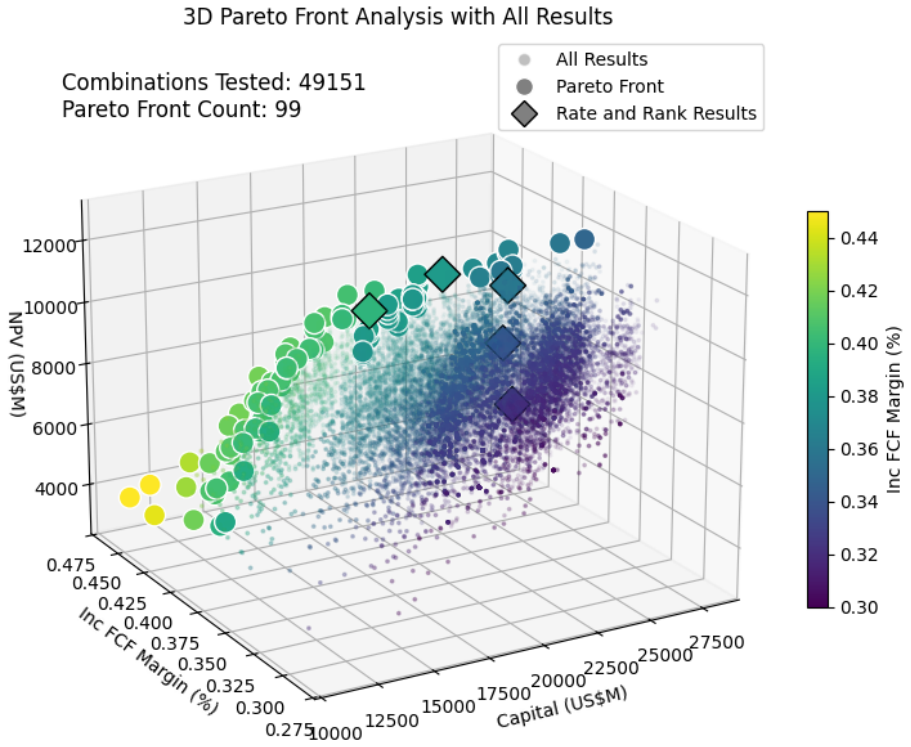


FIG 4 – Comparison of MOO and rate and rank.

Several limitations are apparent in MOO. The method produces a set of efficient portfolios rather than a single recommendation, requiring subsequent qualitative evaluation. Qualitative or non-optimisable parameters such as deposit type cannot be directly incorporated into the optimisation, though they can be applied as post-optimisation filters such as scale and jurisdictional risk of ESG considerations.

However, the advantages are substantial. MOO outperforms rate-and-rank methods for capital allocation by explicitly modelling trade-offs (Nagy, Mansour and Abdelmohsen, 2020), overcoming the clustering and sequencing inefficiencies inherent in rate-and-rank, and increasing project diversity while reducing sensitivity to individual metric choices.

Examining the composition of the highest-scoring portfolio (Figure 5) reveals an instructive pattern: the algorithm consistently selects projects from the HV/HG overlap first introducing assets that combine reasonable margin with meaningful scale, then alternates between high-margin and high-volume projects to balance the portfolio's cash flow and capital profile, before adding the largest capital-intensive projects that fit within the remaining constraint. This sequencing emerges from the optimisation without being prescribed, suggesting that the efficient frontier naturally gravitates toward diversified archetype composition rather than concentration.

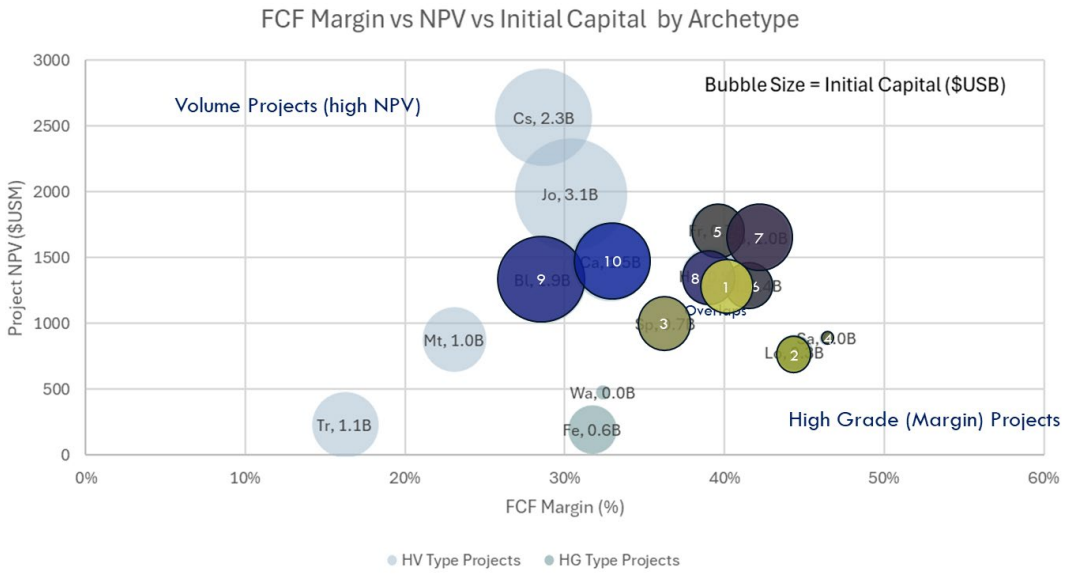


FIG 5 – Sequence of project investment in strategic best portfolio.

Failures of price assumptions and testing beta asymmetry as an alternative – will it perform across the cycle?

Navigating multi-decade investments in cyclical commodity assets requires more than price assumptions. As Zadeh (2025), following Poniewierski, observes, operations near their cost margin exhibit asymmetric economic responses to price movements, a property that is not revealed by conventional price deck planning. Rather than relying on high or low-priced decks, this research investigates measures that quantify both capacity for upside leverage in favourable conditions and downside resilience across simulated price regimes.

Regime-switching simulation was used to test portfolios under different future structural conditions such as downside shocks, upside volatility or stability conditions, with eventual mean reversion. The regimes included a variable time aspect – when the shocks hit in relation to capital spending and how sustained they were. Applying this simulation to the static physicals of the mine plans first exposed a problem rather than an answer: conventional feasibility study cash flow models, which scale revenue linearly with price while holding mine configuration fixed, produce essentially identical price asymmetry across all portfolios. This is a mathematical artefact of the linear model, every dollar of price increase produces a proportional revenue increase regardless of deposit type, and it is the same failure exhibited by conventional spider sensitivity charts, which hold the mine plan constant and report straight-line NPV-price relationships. It reflects a real information gap in capital allocation practice: decision-makers lack visibility into the asymmetric price-response properties that most differentiate investment quality across deposit types.

To measure that asymmetry properly, this paper applies beta asymmetry analysis, adapting Fisher and D'Alessandro's upside/downside beta decomposition approach to simulate price sensitivity rather than market return sensitivity. In plain terms, beta asymmetry measures how much price upside a portfolio captures for every unit of downside risk it carries: a score above 1 means the portfolio gains more value from a price increase than it loses from an equivalent decrease, and a score below 1 indicates the reverse. Formally, upside and downside beta (covariance/variance) are calculated separately by splitting the simulation distribution at the per-portfolio median price present value (PV), ensuring a consistent threshold for each portfolio across its own distribution of outcomes. Beta asymmetry offers a richer measure than simple upside/downside NPV comparison for two reasons. First, beta is calculated across the full distribution of simulated price paths rather than at two specific price points, avoiding sensitivity to the choice of comparison prices. Second, measuring covariance against the present value of each price path rather than its average level captures the interaction between price timing and portfolio capital structure – early price dislocations, when portfolios are still in capital drawdown, produce lower price PV values and are therefore weighted more heavily in the beta calculation than equivalent price movements later in mine life. This temporal

dimension, facilitated by the regime-switching simulation structure, is not captured in static sensitivity comparisons.

Because static mine plans produced no meaningful asymmetry differentiation, a grade-tonnage curve adjustment was introduced as a proxy for dynamic mine configuration response to price. Using published resource sensitivity tables, the model mapped gold price to an implied breakeven cut-off grade, adjusting ore tonnage and average grade accordingly, and recalculated annual cash flows using a variable/fixed cost split calibrated to deposit type. When ore throughput changes in response to cut-off grade adjustment, variable costs scale proportionally with throughput while fixed costs remain constant, producing a non-linear cost response that partially offsets revenue changes in both directions. This provides a dynamic model with a more realistic evaluation of orebody value across a spectrum of price conditions, improving estimation of the assets' value in response to price and revealing which will perform better in stress conditions while leveraging upside in high scenarios. This is shown in the static versus dynamic NPV and implied cut-off grade (COG) curves (Figure 6).

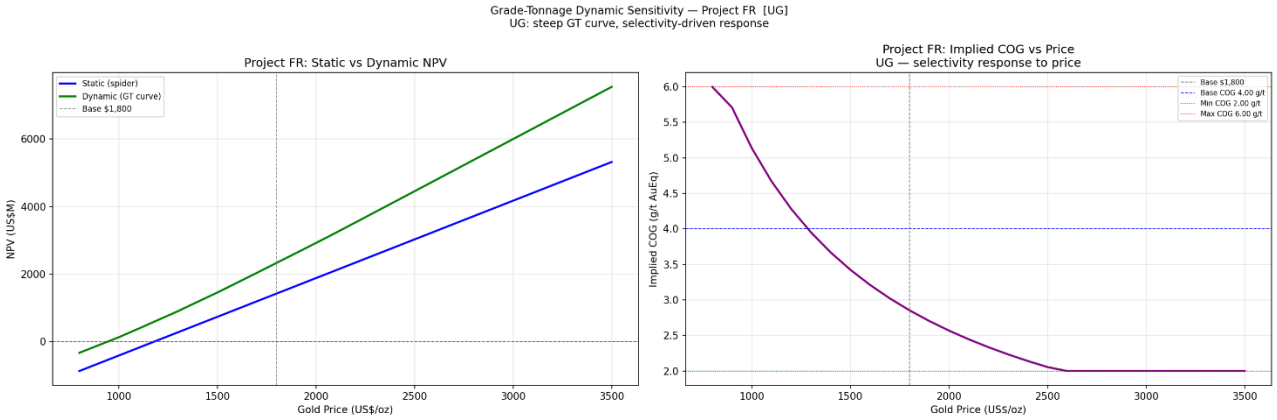


FIG 6 – Dynamic sensitivity.

The dynamic model produces substantially wider asymmetry differentiation across deposit types, with a GT asymmetry range as can be seen from the curved response of the Dynamic model in Figure 6 versus the linear Static line. High-grade underground deposits with steep grade-tonnage curves score highest as when prices fall, selective cut-off adjustment retains grade while sacrificing volume, partially protecting margin. Bulk porphyry open pits with flat grade-tonnage curves score lowest; their flat curves provide minimal grade recovery through cut-off adjustment, and pit shell constraints limit downside flexibility further. The block caving operation showed near-zero additional asymmetry as block caving is inherently non-selective and commits to extracting the full cave footprint regardless of price conditions. These results are conservative estimates of true dynamic asymmetry. The GT model captures the continuous selectivity effect of cut-off grade adjustment but not the discrete configuration changes that generate the most significant asymmetric value in practice, including new pit shells, underground level extensions and processing expansions.

Simply adding asymmetry as a fourth MOO objective at this point proved counterproductive. Adding dimensions to a Pareto optimisation expands rather than filters the efficient frontier; a consequence of the well-known curse of dimensionality in multi-objective problems, producing a larger and more diffuse set of non-dominated solutions rather than identifying the best. The MOO approach was therefore modified to a robust Pareto construction, testing six combinations of two- and three-dimensional objective frontiers and identifying the portfolios that most consistently mapped onto efficient frontiers across all formulations. This robust approach identifies solutions that are hard to beat regardless of price cycles, providing a more reliable basis for capital allocation decisions than any single objective formulation.

Meaningful comparison of risk and reward across deposit types requires normalising for capital deployed; large projects capture more NPV per unit of price risk, but with more exposure of capital, whereas the more sensitive small high margin mines capture more upside per unit of downside risk.

This is the capital efficiency argument for portfolio diversification: not simply that mixing archetypes reduces risk, but that the specific combination of high-volume upside leverage from high volume

assets and high-grade downside protection from more selective mines creates portfolio-level asymmetry that neither type achieves independently.

Dividing beta by initial capex normalises for portfolio scale, reframing the question from total NPV sensitivity to NPV sensitivity per dollar committed. The large absolute betas shrink dramatically when normalised because the portfolios are capital-intensive. This disconnect directly echoes the pro-cyclical investment pattern documented by van der Bijl and Tholana (2020): investment decisions driven by absolute NPV upside systematically overestimate the capital efficiency of large-scale development programs and underestimate the return per dollar available from capital-efficient high-grade alternatives.

Analysis of the top-scoring portfolios (Table 2) reveals a strikingly consistent archetype mix: a core of three to four high-grade underground assets (Fr, Wi, So, Lo) providing capital-efficient asymmetric downside protection, supplemented by one or two mid-scale open pit assets (Sp, Ha, Is) providing upside price leverage. The mega-porphyry assets (Cs, Jo, and Ca) are conspicuously absent despite their large absolute NPV, because their capital intensity and flat grade-tonnage curves generate downside beta that more than offsets their incremental upside contribution per dollar committed.

TABLE 2

Capital normalised selections.

Portfolio ID	Upside beta	Downside beta	Asymmetry	Upside/ Capex	Asymmetry * upside/ Capex	Projects
13187	1.294	1.190	1.087	0.317	0.344	Fr, Sp, So, Lo, Wi
884	1.177	1.083	1.087	0.312	0.339	Wi, Fr, Sp, So
498	1.054	0.973	1.084	0.313	0.339	Lo, Wi, Fr, So
13032	1.316	1.212	1.086	0.311	0.338	Fr, Ha, So, Lo, Wi
7032	1.321	1.221	1.082	0.313	0.338	Lo, Wi, So, Fr, Ha
20646	1.533	1.413	1.085	0.310	0.336	Sp, So, Fr, Ha, Lo, Wi
13128	1.284	1.189	1.080	0.311	0.336	Fr, Is, So, Lo, Wi
7033	1.362	1.260	1.081	0.310	0.335	Lo, Wi, So, Fr, Mt
20073	1.491	1.375	1.084	0.308	0.334	Is, Sp, So, Fr, Lo, Wi
7001	1.254	1.152	1.088	0.307	0.334	Lo, Wi, Sp, So, Fr

Comparing the price deck-based selection to the beta method, four strategies were compared:

1. NPV at base price – conventional NPV-maximising selection (what rate-and-rank does).
2. NPV at high price deck – optimistic price selection (what pro-cyclical behaviour looks like).
3. NPV at low price deck – conservative/stress-tested selection.
4. Robust Pareto – Asymmetry/MOO framework result.

Figure 7 compares how each strategy performs across all four price scenarios in the simulation.

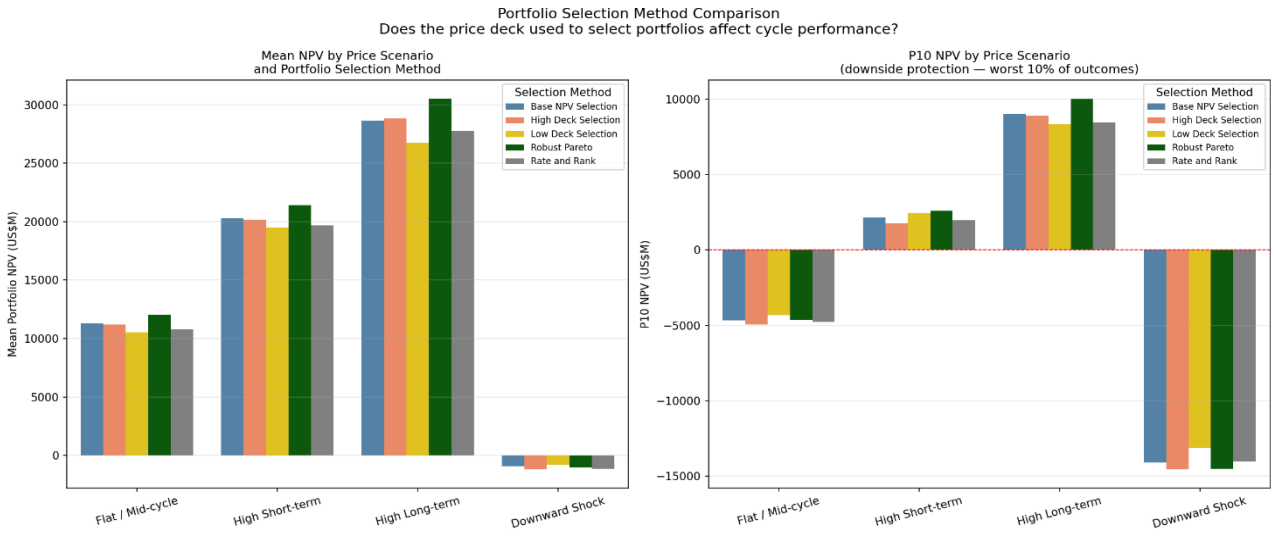


FIG 7 – Comparison of price deck selection versus asymmetry.

The Robust Pareto framework consistently outperforms price-deck-based selection and rate-and-rank in terms of mean NPV across most price scenarios, confirming that portfolio construction method materially affects expected performance. However, all strategies converge in the worst 10 per cent of simulated outcomes, suggesting that development project selection alone cannot protect against severe downside scenarios. Tail risk in a multi-asset producer portfolio is driven primarily by the existing base portfolio's price exposure rather than by development project selection. This points to a limitation of the framework as tested: cycle resilience at the portfolio level ultimately requires not just better development project selection, but active management of the base portfolio's cost structure and price sensitivity.

The low deck (conservative) selection performs surprisingly well on mean NPV in flat and high scenarios, almost matching base NPV selection. That may reflect that conservative price decks don't dramatically change which projects get selected, because good projects look good at any price. However, the Robust Pareto method selects a different portfolio composition, which shows up in its mean NPV outperformance in the upside scenarios, particularly long-term.

The framework successfully identifies a development pipeline for a robust portfolio that will perform throughout the cycle, contingent on good economic management of the base portfolio.

Consistent investment versus pro-cyclical – when should we deviate from the plan?

The third question carries the most direct governance consequence, and it is answered by performing a strategy test to examine whether responding to observed price conditions within the investment decision horizon produces better or worse outcomes than maintaining a consistent capital deployment schedule regardless of price. This distinction is central to the pro-cyclical investment behaviour documented by van der Bijl and Tholana (2020): companies expand capital programs when prices are high and contract when prices are low, producing systematically poor cycle-adjusted returns. Both rate-and-rank and robust Pareto portfolio sets were tested under each strategy, so the comparison can separate whether timing discipline and portfolio construction quality are independent sources of value, substitutes for each other, or complementary strategies whose combination produces outcomes superior to either alone.

The consistent strategy executes the original investment schedule without modification regardless of simulated price conditions. Both strategies' operational cash flows respond to price through revenue scaling and the GT modified models are used to represent operational management changes to cut-off in response to price. The pro-cyclical strategy modified the timing of investment in each portfolio in response to regime simulated price; in high price conditions the capital allowance was increased and investments allowed to move forward in the queue and in poor price conditions investment was deferred by a year.

The comparison model can be summarised as three distinct analytical layers that interact but remain conceptually separate.

Layer 1 – Operating mine response to price (GT dynamic model)

Mines already in production adjust their cut-off grade in response to each year's simulated price. Higher prices lower the cut-off, increasing ore tonnes at the cost of average grade and margin. Lower prices raise the cut-off, sacrificing volume to protect grade and unit economics. This happens continuously throughout the simulation regardless of investment decisions representing the operational flexibility that real mines exercise. This layer was developed separately in the grade-tonnage analysis and is incorporated here as the baseline cash flow response for all committed projects.

Layer 2 – Portfolio composition (pre-determined, not dynamic)

The second layer, which projects are in the pipeline, is fixed before the simulation runs. This is an intentional design choice. Portfolio composition was determined through the robust Pareto multi-objective optimisation, which selected for cycle resilience across the full distribution of price environments. Allowing portfolio composition to change dynamically in response to price would conflate the portfolio construction decision with the timing decision, making it impossible to isolate the effect of either. By holding composition fixed and varying only timing, the simulation cleanly separates two questions that in practice often get conflated: what should we build, and when should we build it.

Layer 3 – Annual investment decision (the timing strategy)

Each year within the five-year investment decision horizon, the strategy observes only the current year's simulated price, not the future path, and makes one of three decisions about uncommitted capital:

- If price is above 125 per cent of the base assumption, the capital budget expands by 30 per cent and the framework attempts to pull the next available project in the investment queue forward by one year. This models the optimistic capital acceleration that characterises pro-cyclical behaviour at price peaks: more cash available, board confidence high, pressure to deploy.
- If price falls below 80 per cent of the base assumption, all new project starts planned for that year are deferred to the following year. The entire queue shifts away one year and whatever project falls beyond the five-year horizon is permanently dropped. This models capital retrenchment, the cash flow squeeze and loss of investment confidence that leads companies to shelve projects at cycle troughs.
- If price sits between these thresholds, the investment schedule proceeds exactly as planned. This is also the consistent strategy's behaviour at every price level: it simply ignores the price signal entirely.
- Projects already committed in prior years continue in both strategies regardless of current price. Their capital is sunk, their construction continues, and their operational cash flows respond through Layer 1. Only uncommitted future starts are subject to the timing decision.

Figure 8 shows the case study outcomes from rate-and-rank versus robust Pareto techniques under different price outcomes. Two behaviours are generally apparent, that Robust methods outperform rate-and-rank; and that consistent through-cycle investment generally outperforms the pro-cyclical behaviour, with the exception of high short-term conditions.

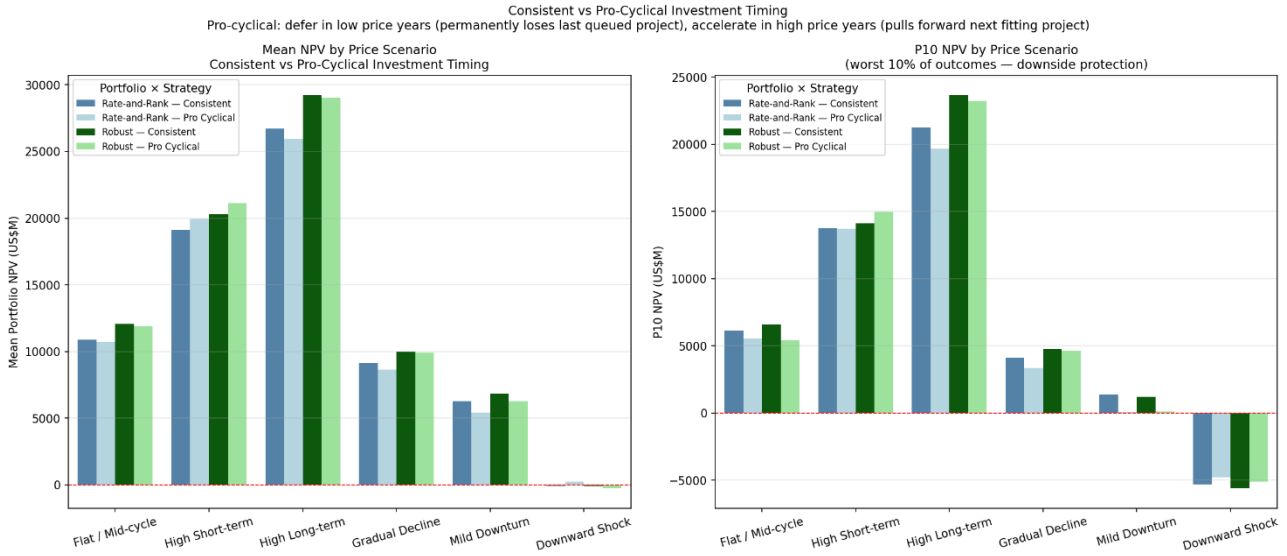


FIG 8 – Comparison of pro-cyclical and consistent investment behaviour.

Robust portfolios outperform rate-and-rank under consistent timing in every scenario (dark green consistently above dark blue on mean NPV), with the largest gap in the high long-term scenario (≈$29 000M versus ≈$26 500M mean NPV); pro-cyclical timing narrows this gap (light green and light blue are closer together), partially destroying the advantage that better portfolio construction creates.

Pro-cyclical acceleration outperforms consistent investment in the high short-term scenario. On further analysis, this is due to the way the high price acceleration was constructed; the capital constraint on acceleration used in the case study inadvertently filters toward smaller, faster-to-production projects that capture the price spike before it reverts. This is an informative finding; if a company wishes to capitalise on prevailing high price conditions, they should do so rapidly and in a manner targeting short payback and limited capital exposure. In the high long-term scenario, the benefit disappeared because a sustained supercycle allows both strategies to build all queued projects and capture elevated prices across the full portfolio regardless of timing. Procyclical behaviour also outperformed consistency in the extreme P_{10} (worst 10 per cent) of downward shock conditions, indicating that, in severe price dislocations, there is a threshold where changing the strategy and preserving capital by deferring uncommitted projects is prudent.

Whilst this case study finding should not be extrapolated to a generalised finding, the inference is that we can now also gain insight into when to stay the course and when to pivot from the strategy. The threshold between ‘intervene’ and ‘stay the course’ lies somewhere between a 25 per cent price decline (mild downturn scenario, price 75 per cent of base) and a 50 per cent decline (downward shock, price 50 per cent of base). The mechanism is structural: acceleration is bounded by project availability, while deferral is unbounded, and the permanent loss of queued investment during downturns consistently outweighs the marginal benefit of pulling projects forward during upswings. Companies that expand aggressively at price peaks and retrench at troughs destroy value through this asymmetry; the finding here distinguishes that systematic pro-cyclical behaviour, which is value-destructive, from tactical capital preservation in genuinely extreme conditions, which may be value-protective on a per-dollar basis.

This is the direct answer to the third question, and it resolves into a governance rule for executive teams. Pro-cyclical behaviour, as commonly practised, conflates two decisions that this framework shows should be kept separate: what portfolio to build, which should be largely independent of current price, and when to modify the build schedule, which should depend on price severity. In practice:

- Build for cycle resilience regardless of current price, using cycle-aware methods (MOO plus regime-switching simulation) rather than the prevailing price deck.

- Hold the investment schedule through normal cycle variation – down to roughly 20–25 per cent below trend, consistent investment outperforms at every outcome percentile in this case study.
- Apply tactical capital deferral only once price dislocation exceeds roughly 30–40 per cent below the long-run trend, where deferral improves per-dollar downside resilience in the tail.

Quantifying this threshold in a specific portfolio context is a direct output of the simulation framework, not a matter of executive judgement alone – though the crossover point identified here is specific to this data set and would need validation against a larger and more diverse project population before being applied as a governance trigger elsewhere. The method for locating it in a new portfolio, however, is directly transferable.

DISCUSSION – WHAT THE THREE ANSWERS MEAN TOGETHER

Conventional price deck evaluation (with or without Monte Carlo simulation) gives a view of project performance under assumed conditions, but not a measure of how investments respond when conditions differ from assumptions. The decision about which price scenario to follow remains an executive judgement call, unaided by any analytical framework for quantifying upside capture versus downside risk. Cash flow models that don't change mining selection dynamically with price are analytically blind to the value of economic asymmetry arising from the ability to reconfigure the mining operation in response to price conditions.

This blindness operates at two levels. At the asset level, spider sensitivity charts show linear symmetric price response because they hold the mine configuration constant. At the portfolio level, rate-and-rank frameworks using NPV as the primary metric favour large capital-intensive projects that appear compelling at current or optimistic prices but carry no embedded flexibility to protect capital in downturns.

Rate-and-rank methods offer clarity, transparency, and participatory appeal but their subjectivity and the effect of systematic archetype clustering compound risk at portfolio level rather than diversifying it. MOO resolves this by discovering portfolios that are efficient across multiple objectives simultaneously. it consistently favoured mixed-archetype portfolios. The practical benefits are tangible: capital spend is more evenly distributed across the investment horizon, free cash flow is more consistent as high-margin underground assets offset the front-loaded capital intensity of volume plays, and capital efficiency improves because the portfolio is not concentrated in the largest projects simply because they have the largest absolute NPV. The emergent pattern of diversified archetype portfolios arose naturally from the optimisation, suggesting that financial efficiency and archetype diversification are complementary in building robust portfolios.

The consistent versus pro-cyclical timing analysis extends the framework from portfolio construction to capital deployment discipline. The finding that consistent investment outperforms pro-cyclical timing across most price scenarios is consistent with van der Bijl and Tholana's empirical evidence; but this research goes further by identifying the mechanism and quantifying a threshold. The value destruction arises from the asymmetry in the decisions; projects can be deferred indefinitely but once investment opportunity is lost it can't be 'caught up' efficiently.

The exceptions to this are intriguing. Pro-cyclical acceleration outperforms consistent investment in the high short-term scenario but not in the high long-term scenario, where one might expect the benefit to be larger. The explanation lies in the mean-reverting structure of the simulation: high price regimes revert after three to eight years, and larger capital-intensive projects cannot achieve payback and upside capture within that window. Only smaller, faster-to-production assets were able to convert the price spike into realised value before conditions normalised. This points to a broader observation not captured in the model but familiar to industry: in sustained price environments it is not the price level alone that determines investment value, but the differential between price and cost inflation that must be maintained long enough for capital to be recovered. A decade of elevated prices absorbed by input cost escalation produces no more real return than a flat cycle, which further complicates the case for aggressive capital expansion at cycle peaks.

The finding that in extreme downside price dislocation tactical capital deferral improves per-dollar downside resilience is also intuitive; however, the ability to find the crossover point where tactical intervention becomes better than staying the course is has strong practical value. Cycles are

navigated from the inside, without knowledge of how deep or how long a dislocation will run, or when to reopen the investment window to be positioned for the recovery. In moderate downturns of 20 to 25 per cent below trend, consistent investment outperforms at every outcome percentile. In severe dislocations of 50 per cent or more, tactical deferral improves per-dollar resilience in the tail. The crossover lies somewhere between these conditions and will vary by portfolio, but the framework provides a method for estimating it in a specific context rather than relying on judgement alone.

The practical implication is that investment governance frameworks should distinguish between two questions that pro-cyclical behaviour conflates: what is the right portfolio to build, which is largely independent of current price, and under what conditions should the build schedule be modified, which is conditional on price severity. Systematic pro-cyclical behaviour gets both wrong. Tactical deferral in genuinely extreme conditions, applied to an already cycle-aware portfolio, is a different and potentially value-protective response.

Limitations and future directions

Several limitations bound the conclusions. The case study covers 14 development-stage gold projects, which is sufficient to demonstrate the framework but insufficient for statistically robust generalisation. The GT dynamic model captures continuous grade selectivity but not the discrete configuration changes – new pit shells, underground extensions, processing expansions – that generate the most significant asymmetric value in practice. The timing model treats the investment queue as fixed, whereas in reality portfolio composition itself responds to price through project approvals, acquisitions, and divestments. And the crossover threshold identified here – somewhere between a 25 per cent and 50 per cent sustained price decline – is specific to this data set and would need validation against a larger and more diverse project population before it could be applied as a governance trigger.

CONCLUSION

This paper set out to address a structural weakness in mining capital allocation practice: the dominance of static, NPV-led evaluation frameworks that assess projects individually, embed single price assumptions, and systematically produce pro-cyclical investment behaviour. The two-stage framework developed here – combining multi-objective portfolio optimisation with regime-switching price simulation – provides a computationally tractable and practically applicable alternative.

The core findings can be stated concisely. Multi-objective optimisation consistently identifies portfolios that outperform rate-and-rank selection on cycle-adjusted performance, not just static financial metrics. The emergent portfolio archetype (mixed high-grade underground and high-volume open pit assets, sequenced to balance capital and cash flow) is financially efficient and structurally diversified in ways that conventional selection methods do not achieve. Robust Pareto analysis, testing optimality across multiple objective formulations simultaneously, identifies a small set of portfolios that are hard to beat regardless of how the problem is framed.

Beta asymmetry, measuring upside capture relative to downside exposure across the full distribution of simulated price paths, provides a more informative selection criterion than price deck optimisation. Unlike a high or low price assumption, it characterises how an investment responds to price variation rather than how it performs at a single assumed level.

The timing analysis delivers a directly actionable finding. Consistent investment discipline outperforms pro-cyclical timing across most price environments. The structural asymmetry of the timing mechanism means pro-cyclical behaviour cannot capture as much upside as it loses on the downside. However, a crossover effect is observed: in severe price dislocations, tactical capital deferral improves per-dollar downside resilience relative to rigid consistency. This suggests that investment governance frameworks should distinguish between two questions: what to build and when to build it, applying cycle-aware portfolio construction to the first and price-severity-conditional discipline to the second. Quantifying the threshold between ‘stay the course’ and ‘stop the bleeding’ in a specific portfolio context is a direct output of the simulation framework.

Taken together, these findings reframe mine portfolio construction as a resilience problem rather than a NPV optimisation problem. The goal is not to find the portfolio with the highest NPV under the

current price assumption; that portfolio is well identified by conventional methods. The goal is to find the portfolio that performs well across the range of price conditions it will encounter over a multi-decade investment horizon. That is a fundamentally different problem, and it requires fundamentally different tools. This paper provides a working implementation of those tools, validated by case study against published feasibility data, offering an improved strategic framework for multi asset mining companies.

ACKNOWLEDGEMENT

This paper builds on an unpublished MBA capstone research project, Improving Strategic Capital Allocation Frameworks for Mining Project Investment in Multi-Asset Companies, completed at Torrens University in 2025. The prior work is available from the author on request.

REFERENCES

Chandra, A and Hartley, P R, 2024. Sequential investment decisions for mining projects using compound multiple volatility real options approach, *Resources Policy*, 97:105241. https://doi.org/10.1016/j.resourpol.2024.105241

Deucker, E, 2025. Improving strategic capital allocation frameworks for mining project investment in multi-asset companies, Unpublished MBA capstone, Torrens University.

Espí, J A and De La Torre, A, 2013. Factors influencing metal price selection in mining feasibility studies, *Mining Engineering*, 65(8):45–51.

Fisher, J D and D'Alessandro, J, 2021. Portfolio upside and downside risk – both matter, *The Journal of Portfolio Management Special Real Estate*, 47(10):158–171. Available from: <https://user.ncreif.org/globalassets/public-site/research/ncreif-insights/portfolio-upside-and-downside-risk----jpm.pdf>

López-García, A, Liern, V and Pérez-Gladish, B, 2023. Determining the underlying role of corporate sustainability criteria in a ranking problem using UW-TOPSIS, *Annals of Operations Research*, pp 1–24. https://doi.org/10.1007/s10479-023-05543-8

Mareels, S, Moore, A and Vainberg, G, 2020, July 8. *Through-cycle investment in mining,* McKinsey and Company. Available from: <https://www.mckinsey.com/industries/metals-and-mining/our-insights/through-cycle-investment-in-mining>

Mathews, G B, 1896. On the Partition of Numbers, *Proceedings of the London Mathematical Society*, s1–28(1):486–490. https://doi.org/10.1112/plms/s1-28.1.486

Nagel, V, Kloppers, B, Choolan, F, Kruger, L, Fraser, C and Kew, R, 2014. Lonmin's capital allocation process–the evolution from individual ranking to portfolio optimisation, in 6th International Platinum Conference (Platinum–metal for the future), October 2014. Available from: <https://www.platinum.org.za/Pt2014/Papers/309-Nagel.pdf>

Nagy, M, Mansour, Y and Abdelmohsen, S, 2020. Multi-Objective Optimization Methods as a Decision Making Strategy, *International Journal of Engineering and Technical Research,* 9:516–522. Available from: <https://www.researchgate.net/publication/340362482_Multi-Objective_Optimization_Methods_as_a_Decision_Making_Strategy>. https://doi.org/10.17577/ijertv9is030480

Pidun, U and Stange, S, 2017. The Art of Capital Allocation, Boston Consulting Group. Available from: <https://web-assets.bcg.com/img-src/BCG-The-Art-Capital-Allocation-Mar-2017_tcm9-149411.pdf>

Regueiro y González-Barros, M and Espí, J A, 2019. The returns on mining exploration investments, *Boletín Geológico Y Minero*, 130:161–180. Available from: <https://www.researchgate.net/publication/332582627_The_returns_on_mining_exploration_investments>. https://doi.org/10.21701/bolgeomin.130.1.010

Savolainen, J, 2016. Real options in metal mining project valuation: Review of literature, *Resources Policy,* pp 49–65. https://doi.org/10.1016/j.resourpol.2016.08.007

Tobisova, A, Senova, A and Rozenberg, R, 2022. Model for sustainable financial planning and investment financing using Monte Carlo method, *Sustainability*, 14(14):8785. Available from: <https://www.mdpi.com/2071-1050/14/14/8785>

van der Bijl, J and Tholana, T, 2020. Analysis of capital allocation by mining companies, in *Proceedings of the 27th International Symposium on Mine Planning and Equipment Selection, MPES 2018,* pp 315–323 (Springer). https://doi.org/10.1007/978-3-030-33954-8_39

Vilhelmsson, R, 2022. Simulation-based decision-making system for optimal mine production plan selection, *Mineral Economics*, 35:59–71. https://doi.org/10.1007/s13563-021-00297-w

Zadeh, J, 2025, April 4. How commodity prices really impact mining companies' performance, Discovery Alert. Available from: <https://discoveryalert.com.au/news/commodity-prices-impact-mining-company-performance-2025/>

Deterministic and stochastic metal price models and implications for strategic mine planning

S F Hoerger[1], T Brady[2] and K Dagdelen[3]

1. Principal Consultant, Peak View Mine Planning LLC, Englewood CO 80111, USA. Email: sfhoerger@gmail.com
2. Professor of Practice, Colorado School of Mines, Golden CO 80401, USA. Email: thomas.brady@mines.edu
3. Professor, Colorado School of Mines, Golden CO 80401, USA. Email: kdagdele@mines.edu

ABSTRACT

Mine engineers build plans with a focus on managing risk and maximising the present value of an operation based on revenues and costs discounted over time. Project teams invest a large amount of time and detail in estimating grades, grade distributions, recoveries, operating costs and capital costs as well as modelling mine, mill and market constraints. These detailed inputs subsequently feed sophisticated mine planning algorithms. In comparison, the project team is often given very simple, conservative metal prices as standardised assumptions to guide mine planning and to control financial evaluation and sensitivity analyses.

Using standardised, simple scenarios as metal price assumptions is a good corporate governance practice for ensuring comparability of financial indicators created by different teams across a corporate portfolio. This ensures corporate decision-makers have consistent metrics to compare brownfield expansions or extensions, greenfield projects and/or external merger, acquisition and divestiture opportunities.

Using constant metal prices supports many fundamental mine planning processes, such as pit limit analysis and breakeven cut-off grade calculations which require fixed price assumptions. Typically, large corporations use conservative constant prices for reserve reporting and other mine planning processes. However, a conservative set of metal price assumptions for mine planning may undervalue options for accelerating production or extending mine life to additional resources. Corporations may use a different set of metal price assumptions for financial evaluation and business planning processes, often starting at a high price and decreasing to more conservative long-term prices. Using this type of decreasing metal price assumption as an input to mine plan optimisation processes can lead to unexpected results such as higher mining rates and inflated cut-off grades.

Examples are used to demonstrate how mine plans and perceived value can change based on differing levels of metal price model complexity using deterministic and stochastic long-term mine schedule optimisation processes based on Mixed Integer Linear Programming. Using an idealised orebody model, example cases show how different price models drive mine schedule changes in terms of cut-off grades, mining rates and pit limits. These cases illustrate how more complex price models might lead to different mine plan recommendations despite having similar near-term or long-term price expectations. The authors discuss trade-offs between price model complexity, value identification and plan robustness and finally, discuss implications for aligning metal price assumptions across a company's mine planning and financial evaluation processes.

INTRODUCTION

Mine technical services teams invest huge effort in modelling geology, simulating grades, predicting metallurgical responses, estimating costs, measuring productivity, and other activities to define inputs to the mine planning process. Yet, the single most important economic parameter, metal price, is usually an input assumption provided by someone outside the technical services organisation such as a Chief Executive Officer, a corporate reserve governance person, a chief economist or an external consultant. Metal prices determine the profitability of planned production, but metal prices also guide the production plan, so understanding price-plan interactions is important for both the givers and receivers of metal price assumptions.

Mine production planning involves determining which volumes of a deposit will be mined when and which material will be sent to which destination (Clark and Dagdelen, 2023). An optimum production schedule will maximise an economic objective (often net present value) while satisfying the operation's capacity, blending and sequencing constraints. Many production schedule optimisation methods have been presented (Newman *et al,* 2010; Dimitrakopoulos and Lamghari, 2022) with varying underlying assumptions and different approaches to practically achieve solutions to this complex multi-dimensional problem.

Production scheduling plays a key role in maximising the value of constrained open pit mining complexes. Efficient production schedules defer stripping costs until needed and select mining locations and rates such that limited processing capacity is efficiently utilised by balancing the opportunity cost of wasting (or stockpiling) marginally profitable rock against the opportunity cost of deferring the processing of future higher-grade ores with resulting deferral of future profits (Lane, 1988; Dagdelen, 1992; Rendu, 2014). For long-life open pit mines, significant value might be created by deferring a stripping campaign or by adjusting cut-off grade policies and mining rates so that the lowest grade marginal mill feed is replaced by higher grade ore. Depending on metal prices, a mining complex might improve profitability per mill tonne by reprioritising underground mines, open pit laybacks and stockpiles. Metal price forecasts are critical inputs to identifying and evaluating these strategic trade-offs.

In recent years, as gold prices have increased, industry practice for large mining companies has trended toward conservative reserve planning prices that have been significantly below average yearly spot gold prices (Table 1). The reserve planning prices have also been below the three-year trailing average gold prices. Smaller companies often use less conservative pricing.

TABLE 1

End of year gold reserve reporting prices for select USD reporting gold producers versus yearly average spot prices (Sources: macrotrends.net and company reports via Google).

	2019	2020	2021	2022	2023	2024	2025
Full Year Average Au Price	1405	1785	1793	1799	1954	2405	3473
3-yr Trailing Average	1313	1485	1661	1792	1848	2052	2610
EOY Reserve Reporting Price							
Barrick	1200	1200	1200	1300	1300	1400	1500
Agnico Eagle	1200	1250	1250	1300	1400	1450	1600
Kinross	1200	1200	1200	1600	1600	1600	1600
AngloGold Ashanti	1100	1200	1200	1400	1400	1600	1700
Northern Star (Pogo)	1150	1350	1350	1350	1500	1725	1725
Gold Fields	1200	1300	1300	1400	1400	1500	2000
Newmont	1200	1200	1200	1400	1400	1700	2000

The objective of this study is to gain insights into the interactions between strategic mine plans and metal price assumptions. The study seeks to investigate:

- Does a significant disconnect between mine planning prices and actual prices put significant value at risk?
- Can different types of metal price models lead to more or less valuable plans?
- Can stochastic price models and stochastic optimisation lead to higher value plans or more robust plans compared to deterministic plans based on deterministic price models?
- Can stochastic price models benefit the planning process without requiring a full stochastic optimisation?

Note that this study is focused on developing a single net present value (NPV) maximising mine plan based on an initial stochastic understanding or deterministic version of orebody and price uncertainty. This single-mine-plan approach will provide a conservative NPV compared to a real

options approach which considers options to vary mine life, locations and rates as uncertainty is revealed over time (Elkington and Gould, 2012).

METAL PRICE MODELS FOR STRATEGIC MINE PLANNING

Multiple uses of metal price forecasts

Metal price forecasts are used as inputs for two separate processes. First, metal prices are key input assumptions for developing mine plans. Second, after a mine plan is created, different metal price assumptions may be used for financial evaluation of the resulting business plan.

Mine planning prices for plan development

Many traditional mine planning methods require a single constant price assumption for each saleable commodity. Metal prices are necessary static inputs for computing breakeven cut-off grades, optimised ultimate pit limits, or annual reserve statements. Other deterministic mine planning processes, such as production scheduling or cash flow analysis can incorporate metal prices that change over time.

The estimation and reporting of annual reserves requires a single metal price for each relevant metal. Additional prices may also be utilised for resource reporting and sensitivity analyses. During periods of rising prices, the reserve price is usually conservative to ensure that all reserves will generate cash flow, even if price gains are not sustained. Conservative pricing also reduces the risk that future reserve reporting will need to use a lower price resulting in a downward reserves revision. Typical information that guides setting of reserve metal prices will assess: (a) three-year trailing average prices; (b) recent low points in metal price charts; and (c) surveys of what metal prices other companies are using. These information sources are evaluated against spot prices, futures prices and the previous year's reporting price to finalise a reserve reporting price. In addition to metal prices, reserve reporting also requires assumptions for location-specific discount rates, key exchange rates, byproduct metal prices and input commodity costs such as diesel. Ideally, a corporation will use the same reserve prices and other assumptions for all its mines to ensure comparability. However, site to site differences in partnership agreements and planning calendars may prevent complete consistency across all mines and projects in the corporate portfolio.

Unfortunately, futures markets are not very accurate predictors of realised future prices. As shown for the 2008 to 2014 time period (Figure 1), gold futures prices did not accurately predict realised price increases nor decreases. Gold futures prices are highly correlated with spot prices with small trends driven by storage costs, borrowing costs and other factors. Adding to the challenge of anticipating price trends, metal futures markets typically have very limited trading activity for contracts beyond 12 months.

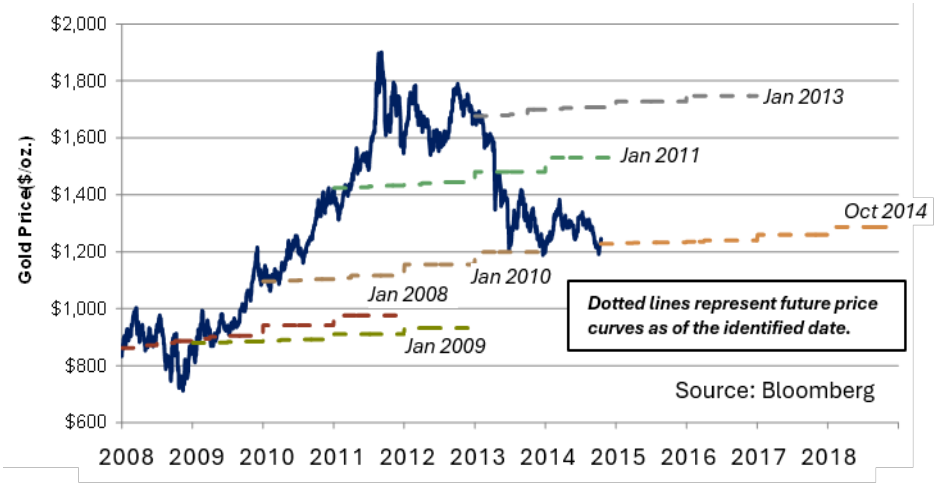


FIG 1 – Comparison of actual gold prices versus futures contracts available at different dates.

Business planning prices for financial evaluation

For strategic planning or project evaluation purposes, financial evaluation may focus on reserve pricing while also considering price forecasts for multiple upside and downside scenarios. For example, companies may use an externally published industry cost curve to define a potential floor price at which competitors may need to shutdown production. Also, financial evaluation with conservative pricing is often a requirement for bankers and other providers of investment capital to ensure that projected cash flows provide sufficient coverage of loan repayments. Assessment at more optimistic price forecasts will add perspective on the project's upside potential. More optimistic price forecasts are often guided by alternative macroeconomic scenarios provided by metals analysts (Figure 2).

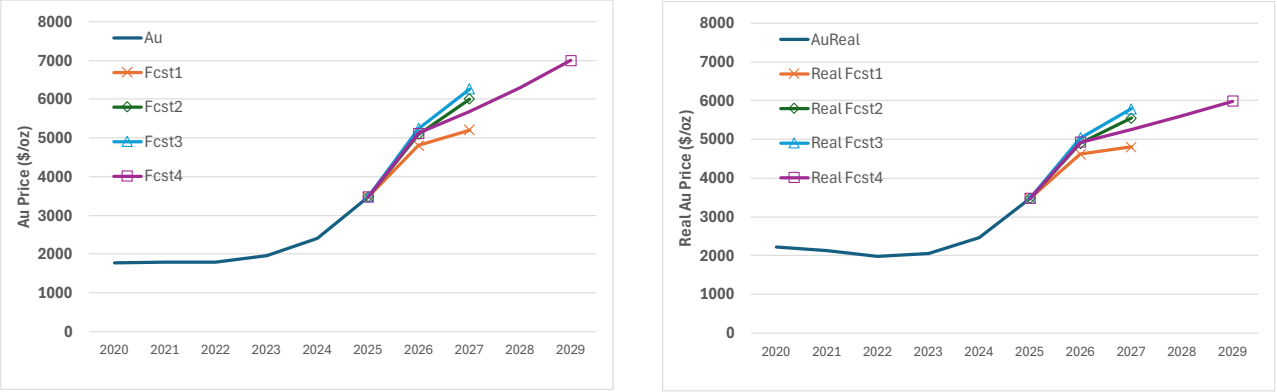


FIG 2 – Example of analyst gold price forecasts from May/June 2026 on a nominal (left) and real (right) basis (adjustment to 2025 basis: US consumer price index, future inflation of 4 per cent per annum).

For business planning purposes, less conservative approaches are needed to more accurately estimate revenue, costs (some of which correlate with metal prices) and cash flow (to accurately plan taxes, dividends and/or financing requirements), with heavy focus on the first few years of a given plan. Key concerns include understanding which metal prices could trigger loan covenant restrictions or inability to internally fund capital spending programs. Long-term metal prices are critical for accounting valuation and impairment calculations. Changes in long-term price assumptions will receive extra scrutiny (and inertia) because of the uncertainty in predicting long-term prices versus the potentially significant accounting valuation impacts.

A business planning metal price forecast may often show higher prices in the early years before transitioning to a lower, more conservative price in the outer years. Early year prices are set with an emphasis on spot prices to most accurately predict near term cash flows. Long-term prices for a base case scenario are often set conservatively to align with reserve pricing and/or to ensure business sustainability during potential future low points in a commodity price cycle.

In addition to a base case scenario, business plans or new investments are often evaluated against multiple price scenarios corresponding to different macroeconomic scenarios or simple sensitivities. Metal price forecasts from different banks are available through subscription services such as Bloomberg. The forecasts plotted in Figure 2 were created when the spot price of gold was averaging approximately $4500/oz. Key forecast drivers cited by the metals analysts included central banks shifting from US dollar reserves to gold reserves, individual investors increasing their portfolio gold allocation, inflation-adjusted interest rates and U.S. dollar strength partially offset by jewellery demand price sensitivity. Supply scenarios are more important for commodities other than gold and will attempt to factor in expected mine closings and openings and the potential actions of any country or company that controls a significant portion of the commodity's mining or processing capacity.

Types of metal price forecasts

Deterministic forecasts

Traditionally, the most common type of metal price model is a point forecast that does not change over time. For most mine planning methods, the point forecast works best if it is considered a constant (or real) price and can be compared against constant costs which also are considered invariant with time.

A more complex type of metal price model is an annual forecast. Key considerations in an annual forecast include ensuring appropriate correlation between all relevant metals and major input cost assumptions. If planning will work with real cost assumptions, the annual metal prices should also be on a real basis (ie nominal prices adjusted downwards for escalation; see Figure 2 example) and discount rates should also be on a real basis.

Point forecasts and annual forecasts usually define a base case and various alternative scenarios. At a corporate level, defining standardised base case and alternate scenario assumptions allows for different projects and investment alternatives across a corporate portfolio to be evaluated with consistent and comparable metrics. Conservative metal price scenarios may be used for internal go versus no-go versus wait decisions to minimise risk, while scenarios closer to current spot prices may be used for acquisition and divestiture decisions and negotiations.

Stochastic forecasts

Alternatively, while less commonly employed, companies may use stochastic price forecasts, either as a discrete set of probabilities attached to several scenarios or as simulated annual metal prices. The authors have seen some benefit in executive level decision-making confidence when a probability distribution of project NPV's is presented; however, the authors have seen more decision-making clarity from breakeven analysis where each project computes the metal price at which NPV equals zero. Computing a single NPV breakeven price requires the simplicity of a single primary metal price with byproduct metal prices treated either as constant or varying via a fixed relationship with the primary metal price.

A frequently used (Visnjic, 2018) stochastic method for modelling metal prices over time assumes geometric Brownian motion combined with mean reversion (Schwartz and Smith, 2000). This model can be implemented to simulate annual price forecasts, with the following equation (Gillespie, 1996):

$$P_{t+1} / P_t = \exp[(\ln M - \ln P_t)(1- e^{-k}) + \sigma N \text{sqrt}((1-e^{-2k})/2k)] \quad (1)$$

Where $P_{t=0}$ is the starting price, M is the long-term mean equilibrium price, k represents the annual mean reversion rate, σ is the standard deviation of price percentage returns, and N is a random number from a normal distribution with mean 0 and standard deviation 1.

DETERMINISTIC PRICE FORECASTS' IMPACT ON MINE PLANS

Mine production schedules maximise NPV when cut-off grades are defined based on constraints (Lane, 1988). When constrained by mining capacity, a mine should use a breakeven cut-off grade where revenue equals variable milling costs. The unconstrained mill tonnage is then defined by the number of tonnes above this mill-variable-cost breakeven cut-off grade. Under these mine-constrained conditions, the current period metal price defines the revenue and therefore defines the breakeven cut-off grade. When constrained by milling capacity, the marginal cut-off grade is set where enough revenue is provided to pay for variable milling costs, mill and site fixed costs, and the opportunity cost of delaying future profits. The revenue portion of this equation is defined by the current period metal price; however, the opportunity cost portion is defined by future metal prices in addition to future costs, constraints and orebody characteristics. So, when mill constrained, current and future metal price forecast assumptions have a large impact on identifying the best mine plan for an orebody.

When an operation is jointly constrained by both mining and milling capacity, the best cut-off grade is a balancing cut-off grade which results in full utilisation of both mining and milling capacity. The balancing cut-off grade should be the mine's operating cut-off grade if it is higher than the mine constrained cut-off and less than the mill constrained cut-off. Note that the balancing cut-off grade is

solely a function of the grade tonnage distribution and consequently is not impacted by metal prices unless metal prices move the mine constrained or mill constrained cut-off grades above or below the balancing cut-off grade. Although the balancing cut-off grade is not affected by metal prices, the financial incentive for increasing mining capacity will be affected.

MILP methodology for creating mine plans

The impacts of metal price forecasts in relation to NPV maximising mine plans will be demonstrated using a simple deposit model and a Mixed Integer Linear Programming (MILP) optimisation algorithm. The MILP algorithm chooses mining and processing tonnage flows that maximise the NPV of the orebody subject to mining system constraints. The model is based on a mine-layback-panel (equivalent to pit-phase-bench) level of mine geometry resolution with grade-tonnage resolution for each panel based on increments that have sufficiently narrow definitions that allow each increment's properties to be treated as constant throughout the deposit. Typically, increments would be defined by combinations of metallurgical types and grade bins.

The MILP formulation solves a 'two-stage stochastic programming with recourse' problem (Wolsey, 2021). The first stage decision variables, **W**, define a mine sequence in terms of where and when to mine. The goal is to optimise the mine sequence knowing the potential distribution of uncertainty but without knowing which specific prices or orebody realisation will actually occur. Once each period's metal price and orebody realisation are revealed, the second stage decision variables, **X**, define which increment tonnes are sent to which destinations (eg mill or waste dump) to maximise value while respecting process capacity constraints.

The algorithm does not rely on cut-off grades, but the solution effectively defines them by defining which increments are sent to the mill destination during each time period. This formulation requires that each mine is subdivided into predefined laybacks and panels (usually one panel per bench of a layback), but does not require that every panel is mined, so the solution result will effectively select which panels define the ultimate pit limit. The most critical equations in the MILP formulation are:

Solve for:

W(m,ℓ,p,t): tonnes mined from each panel of each layback in each time period

X(m,i,d,t,k): increment tonnes sent from each mine to each destination in each time period

F(d,t,k): binary variable 1 if destination is incurring fixed costs, 0 otherwise

With dimensions:

m = 1 to M mines, with each mine subdivided into ℓ = 1 to L laybacks times p = 1 to P panels

i = 1 to I increments

d = 1 to D destinations

t = 1 to T time periods

k = 1 to K realisations of orebody grade tonnage distributions × metal price profiles

To Maximise Net Present Value =

$$\sum_m \sum_\ell \sum_p \sum_t pv(t)*MineVal(m,\ell,p,t)*\mathbf{W}(m,\ell,p,t) +$$
$$1/K * \sum_k \sum_m \sum_i \sum_d \sum_t pv(t)*DestVal(m,i,d,t,k)*\mathbf{X}(m,i,d,t,k) +$$
$$1/K * \sum_k \sum_d \sum_t pv(t)*FixedCost(d,t,k)*\mathbf{F}(d,t,k) \quad (2)$$

Subject to:

Finite resources ∀ panels (m,ℓ,p):

$$\sum_t \mathbf{W}(m,\ell,p,t) \leq Tonnes(m,\ell,p) \quad (3)$$

Mine Continuity ∀ mine increments (m,i,t,k):

$$\sum_\ell \sum_p \mathbf{W}(m,\ell,p,t)*[tonnes(m,\ell,p,i,k)/Tonnes(m,\ell,p)] = \sum_d \mathbf{X}(m,i,d,t,k) \quad (4)$$

Mine Capacity ∀ mines (m,t):

$$\sum_{\ell} \sum_{p} \mathbf{W}(m,\ell,p,t) \leq \text{MineCapacity}(m,t) \quad (5)$$

Process Capacity ∀ destinations (d,t,k):

$$\sum_{m} \sum_{i} \mathbf{X}(m,i,d,t,k) \leq \text{DestCapacity}(d,t)*\mathbf{F}(d,t,k) \quad (6)$$

Where the notation $\sum_m$ represents a summation from m=1 to m=M and with the following constants:

MineVal(m,ℓ,p,t): $/t mining cost coefficient

DestVal(m,i,d,t,k): $/t value coefficient for cost and recovery, increment grade, realisation price

$pv(t) = 1 / (1 + d)^t$: present value factor using discount rate, d

tonnes(m,ℓ,p,i,k): Orebody tonnes for each mine, layback, panel, increment, realisation

$Tonnes(m,\ell,p) = 1/K * \sum_k \sum_i tonnes(m,\ell,p,i,k)$: Total tonnes for each panel

With deterministic orebody and metal price inputs, K equals 1, and the realisation dimension disappears to create a deterministic optimisation formulation. The full formulation (Hoerger and Dagdelen, 2025b) includes additional equations for mine sequence prerequisites and sinking rate constraints as well as equation extensions for modelling stockpiling.

Example case

To demonstrate the impacts of different metal price models, an example case assumes a simple gold mining complex consisting of a single open pit mine where mined rock can be sent to a mill or a waste dump destination. The deposit model approximates a dipping orebody consisting of 20 panels. Each panel has a constant volume of mineralised material and an overburden volume that grows as mining progresses deeper through the panels (Figure 3 left). The mineralised zone grade tonnage distribution for selective mining units assumes a lognormal distribution with mean of 1.0 grams per tonne (gpt) and standard deviation of 1.0 gpt (Figure 3 right). As the overburden ratio increases, the breakeven cost per ounce of each panel increases from $1150/oz to $2500/oz (Figure 3 centre blue line) based on the cost parameters in the following paragraph.

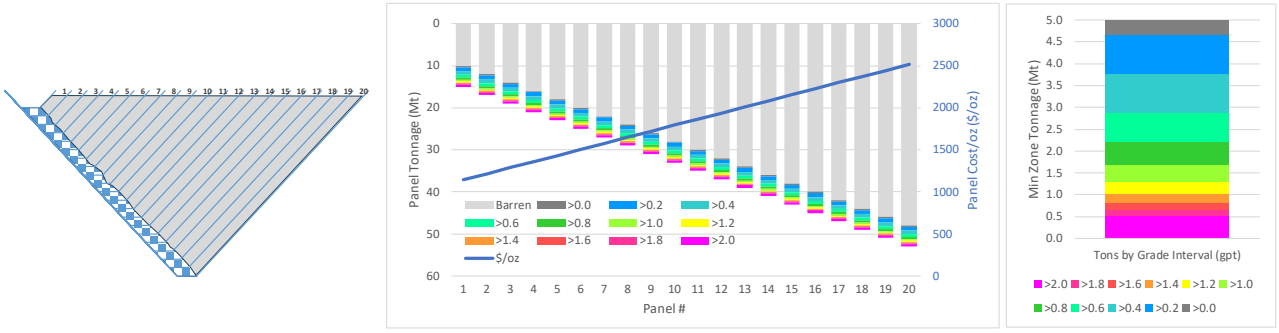


FIG 3 – Deposit model cross-section (left); 20 panels with increasing barren tonnage and constant mineralised zone tonnage in each panel (centre); lognormal grade distribution of selective mining units (mean = 1.0 gpt, standard deviation = 1.0 gpt) in mineralised zone (right).

The example case is evaluated using the following parameters. Annual mill throughput is constrained to 4 million tonnes (Mt) at a full capacity cost of $20/t of ore processed including $12/t variable costs; mill recovery is set to 90 per cent and refining costs are $20/oz. Present values are computed using an 8 per cent discount rate. A key detail of this example implementation is assigning a $32M/a fixed cost to mill operation which is incurred every year up to and including the final year of operation regardless of how many tonnes are milled each year. Mining capacity is unconstrained at a cost of $5/t of material mined which includes equipment leasing costs to enable flexible yearly mine tonnages allowing the examples to focus on a mill constrained system. The mineralised zone grade distribution is discretised into increments using 36 grade intervals to allow the algorithm adequate flexibility for varying cut-off grades in each period for each price forecast. Sequencing constraints are configured such that each panel must be completely mined out before starting the next panel.

Within a panel, the tonnes that are mined in a period are assumed to be mined in equal proportions from the overburden and all mineralised grade intervals.

This example case is evaluated under varying metal price assumptions to demonstrate the sensitivity of mine plan optimisation to metal price inputs.

Point forecasts

First, point value (or constant) metal price forecasts were considered for a mill-constrained case (ie unconstrained by mining). As shown in Figure 4, under higher assumed prices, mining starts at a faster rate and with a higher cut-off grade due to a higher opportunity cost of delaying future profits. In addition, mining finishes with a lower cut-off grade due to lower grades becoming profitable with higher metal prices. With higher metal prices, more of the panels becomes profitable and mine life is extended until prices exceed $2500/oz and the entire deposit is mined out. At $4000/oz and $5000/oz, the higher cut-off and faster mining early in the mine life result in depletion of all panels sooner than at $3000/oz. In all cases, although mining starts with an elevated cut-off, the mining rate still increases over time to keep up with the deposit's increasing strip ratio.

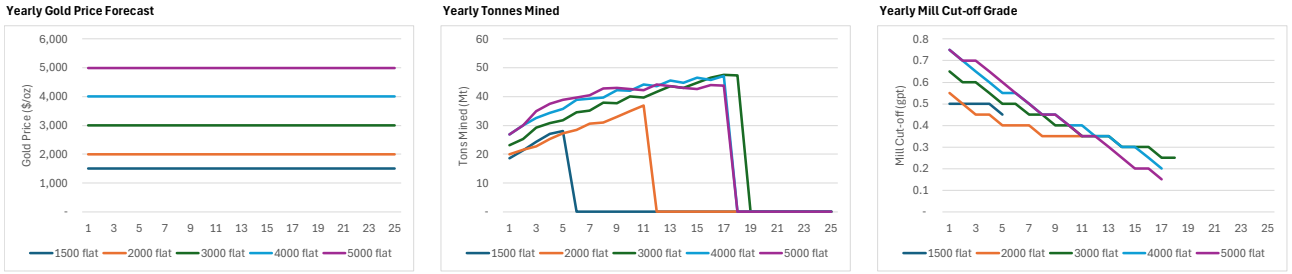


FIG 4 – Impact of constant metal prices on NPV maximising mine plans.

Constant growth/decline annual forecasts

Second, metal price annual forecasts (non-constant) were considered for the same mill-limited case (ie unconstrained by mining). As shown in Figure 5, increasing prices encourage slower mining and lower cut-offs than the constant price case because higher future prices decrease the opportunity cost of delaying future mining (compare Figure 5 to Figure 4). Conversely, decreasing metal prices increase the opportunity cost of delaying future mining, encouraging faster mining and higher cut-offs.

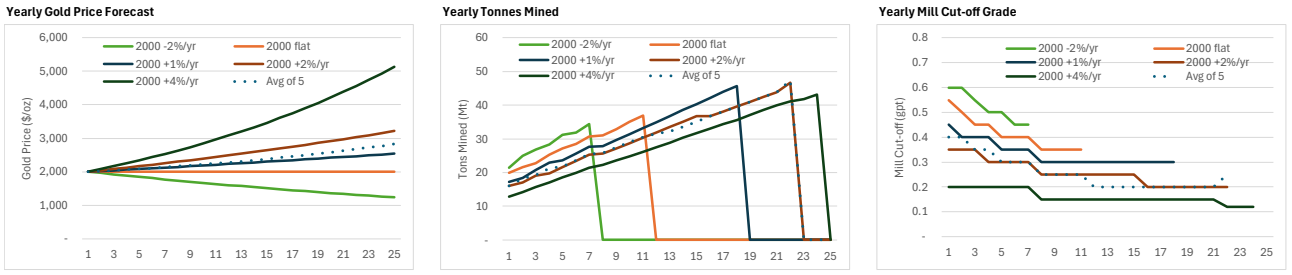


FIG 5 – Impact of increasing/decreasing metal prices on NPV maximising mine plans.

When this example was run with lower proportions of fixed costs per mill tonne, the increasing price scenarios created stop-start mine plans where mining and milling would stop for one or two years until the higher price created a larger margin and created higher NPV despite an extra year of fixed costs and an extra year of discounting future cash flows. This behaviour demonstrates some common challenges for strategic mine planning in a rising price environment. For a new project with modest returns, is it better to start the project today or wait for higher prices? Too often, projects are evaluated with just 'go' or 'no-go' options without considering a 'wait' option. For existing high-cost operations, is it better to incur fixed costs and preserve gold in the ground for future mining rather than produce the gold at a low profit margin today? Future price assumptions are critical inputs to answering these questions. When margins are small, a small percentage increase in price can still encourage delayed mining and offset a larger discount rate.

Mean reversion annual forecasts

Third, mean reversion types of metal price annual forecasts were evaluated with prices reverting to perceived long-term means. If the gold price is forecast to drop significantly in future years, optimised plans will mine faster and raise cut-offs to produce more gold today before prices drop tomorrow (Figure 6). Higher cut-offs are impacted by both the size and rate of expected future price decreases.

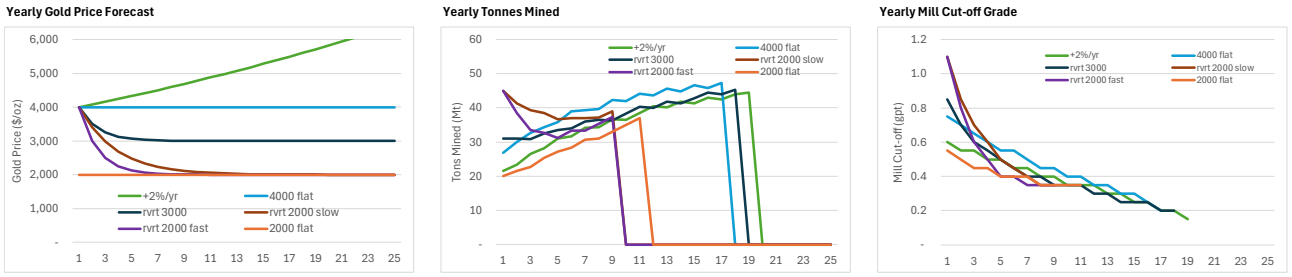


FIG 6 – Impact of mean reverting metal prices on NPV maximising mine plans.

A metal price annual forecast that has spot prices in the near-term while reverting to reserve prices in the long-term may be prudent for business planning but likely has unintended consequences if used for creating mine plans. A downward mean reverting price can encourage a faster mining rate and higher cut-off than either a flat high metal price or a flat low metal price (compare Figure 6 versus Figure 4). An upward mean reverting price can encourage slower (or stopped) mining while waiting for higher prices to improve margins.

Plan robustness – impact of mismatch between planned versus actual metal prices

To test robustness, each mine plan created using one assumed annual price forecast was evaluated against a range of alternative annual price forecasts (Table 2). The $4000/oz assumption plan uses a year 1 mill operating cut-off grade of 0.75 gpt and mining rate of 27 Mt to increase year 1 production by 17 per cent and $4000/oz cash flow by 16 per cent versus the $2000/oz assumption plan at 0.55 gpt and 20 Mt. If actual pricing remains constant at $4000/oz, the plan built using a $4000/oz flat pricing assumption will be 12 per cent better over the first five years compared to the plan built using a $2000/oz flat pricing assumption. The $4000/oz assumption plan would also provide 11 per cent and 13 per cent higher 5-year cumulative discounted cash flow respectively even if prices revert to $2000/oz or grow by 2 per cent/a. The plans created assuming $4000/oz flat or $4000/oz +2 per cent/a provide superior cash flows in years 1 to 3 compared to the plan created using the $2000/oz flat assumption regardless of which of the four price scenarios is actually realised.

TABLE 2

Impacts of planning with one set of price assumptions and realising a different set of prices.

Price Assumption for Creating Mine Plan	Price Assumption for Actual Cash Flow/NPV	NPV thru Year 5 $M	NPV thru Year 5 Δ vs 2000	NPV thru Year 9 $M	NPV thru Year 9 Δ vs 2000	NPV thru LOM $M	NPV thru LOM Δ vs 2000	5yr Avg Cost $/oz	5yr Avg Cost Δ vs 2000	Yearly Cash Flow 1	2	3	4	5
$2000 flat	$4000 flat	1,801	0%	2,600	0%	2,879	0%	1,303	0%	469	435	403	391	372
$4000 flat	"	2,027	12%	2,851	10%	3,591	25%	1,376	6%	545	501	460	424	390
$4000 reverting to $2000	"	2,045	13%	2,778	7%	2,778	-4%	1,456	12%	652	514	428	380	337
$4000 + 2%/yr	"	1,886	5%	2,703	4%	3,565	24%	1,321	1%	489	453	437	402	388
$2000 flat	$4000 reverting to $2000	1,051	0%	1,174	0%	1,188	0%	1,303	0%	469	279	178	128	95
$4000 flat	"	1,168	11%	1,248	6%	1,143	-4%	1,376	6%	545	317	195	127	84
$4000 reverting to $2000	"	1,227	17%	1,293	10%	1,293	9%	1,456	12%	652	317	173	106	67
$4000 + 2%/yr	"	1,092	4%	1,204	3%	1,116	-6%	1,321	1%	489	289	192	128	94
$2000 flat	$4000 + 2%/yr	1,899	0%	2,883	0%	3,271	0%	1,303	0%	469	448	427	427	421
$4000 flat	"	2,137	13%	3,165	10%	4,361	33%	1,376	6%	545	515	489	466	444
$4000 reverting to $2000	"	2,148	13%	3,065	6%	3,065	-6%	1,456	12%	652	530	456	418	384
$4000 + 2%/yr	"	1,989	5%	3,000	4%	4,411	35%	1,321	1%	489	466	463	441	439
$2000 flat	$2000 flat	473	0%	585	0%	597	0%	1,303	0%	143	122	103	90	77
$4000 flat	"	493	4%	560	-4%	455	-24%	1,376	6%	163	133	107	84	64
$4000 reverting to $2000	"	446	-6%	501	-14%	501	-16%	1,456	12%	171	119	88	67	49
$4000 + 2%/yr	"	486	3%	586	0%	498	-17%	1,321	1%	149	126	110	89	75

Note that these comparisons assume that the operation is not mine constrained, but the size of the NPV differences suggest that mine capacity could be profitably expanded if necessary. For simplicity, all examples presented have not included stockpiling. If stockpiling was considered, optimised mining rates and operating cut-offs would be elevated further due to the ability to process below-cut-off material later in the mine plan and thus reduce the opportunity cost of mining faster. Stockpiling also conserves low-grade potential resources and creates optionality.

Plan sensitivity to deposit model mean and selectivity

Finally, for perspective (and in preparation for a stochastic mine planning example), the mine plan impacts of changes in deposit mean grade and mine selectivity were evaluated by changing the mineralised zone's lognormal mean by ±10 per cent and relative standard deviation by ± 20 per cent. This sensitivity analysis was created using a $2000/oz point forecast.

This example showed that changing the mean grade increases NPV significantly and allows mining to extend to higher overburden ratios, with mining rates and cut-offs changed only slightly (Figure 7). For a given mean, lower relative standard deviation (ie lower selectivity) reduces the ultimate NPV (= the year 25 cumulative discounted cash flow in Figure 7) slightly. The relative standard deviation had a significant impact on the NPV maximising mining rates and cut-off grades. Also, for the elevated cut-offs of years 1 to 3, the yearly cash flow impact caused by relative standard deviation changes was similar to the impact caused by mean changes. These impacts highlight the importance of accurately estimating selectivity when planning an elevated cut-off strategy.

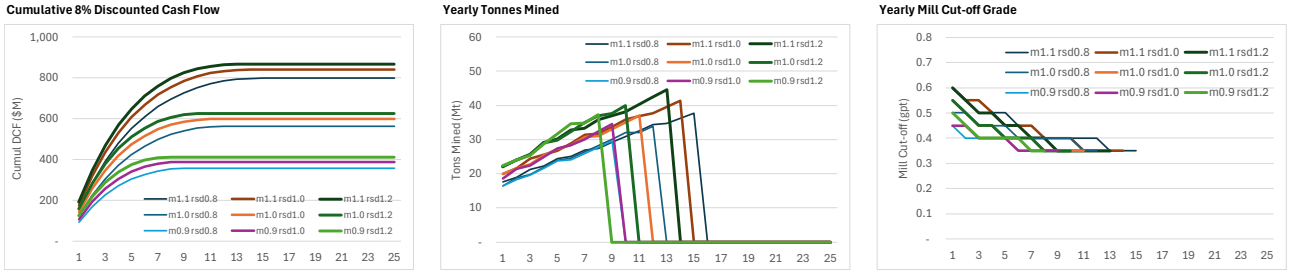


FIG 7 – Impact of different lognormal means and relative standard deviations on NPV maximising mine plans. All cases evaluated at $2000/oz.

Comparing Figure 7 against Figure 4, the cut-off grade increase and mine life extension of increasing the mean by 10 per cent and the relative standard deviation by 20 per cent were nearly equivalent to increasing the metal price forecast from $2000/oz to $3000/oz.

STOCHASTIC PRICE FORECASTS' IMPACT ON MINE PLANS

Stochastic optimisation allows using stochastic price forecast inputs as well as stochastic orebody models. For stochastic production schedule optimisation, a single mining sequence is determined to maximise NPV across a distribution of stochastic price and/or orebody realisations. For each realisation, as each period's metal price and grade/tonnage distribution are revealed, increment flows to each destinations adapt to maximise value subject to destination constraints.

Example case and methodology

Two examples will be used to demonstrate the stochastic optimisation process. The base orebody model, costs and constraints are the same as used in the previous deterministic section.

Ornstein-Uhlenbeck stochastic forecasts

For the first stochastic example, gold prices are simulated using an Ornstein-Uhlenbeck model starting at $2000/oz with a 2 per cent/a yearly increase in the real equilibrium price, a yearly standard deviation of 20 per cent and an annual mean reversion factor of 30 per cent. Using Equation 1, 25 realisations of 50-year price profiles are created (Figure 8 upper left). The MILP stochastic optimisation solves for mining rate and yearly cut-off grades to maximise NPV (Figure 8 upper right)

resulting in a 21-year mine plan. Output reports show the uncertainty in yearly cash flow and cumulative discounted cash flow (Figure 8 lower left and right).

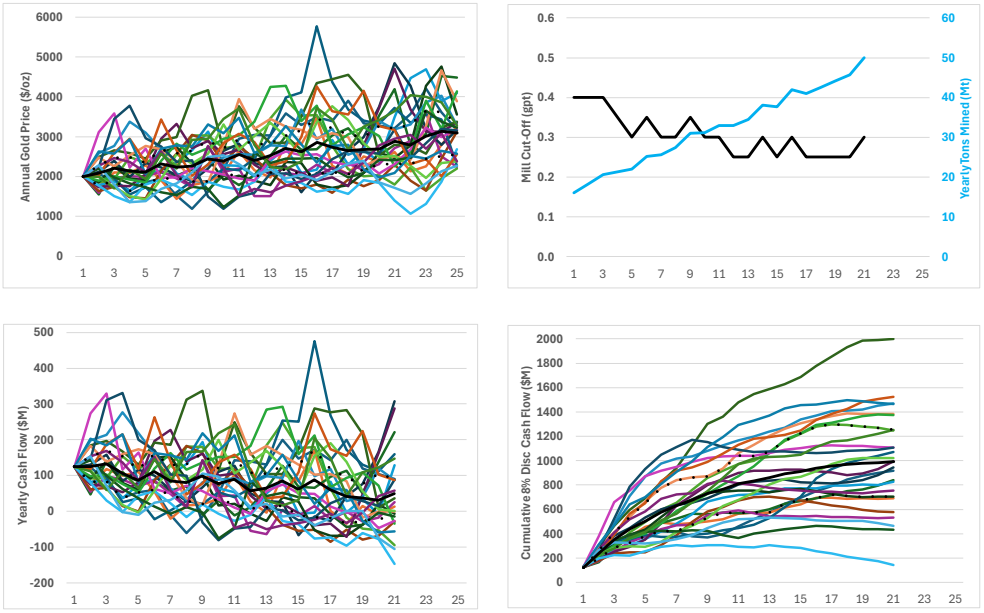


FIG 8 – Stochastic optimisation using 25 Ornstein-Uhlenbeck simulated annual price forecasts. Upper left – gold price forecast inputs (thick black line for average of 25 simulations; dotted black lines for 25th and 75th percentiles; thin coloured lines for 25 individual simulations); Upper right – solution for mining rate (right axis) and average cut-off grade (left axis); Lower left – annual cash flow uncertainty; Lower right – cumulative discounted cash flow uncertainty.

Although there is significant variability in yearly cash flow, there is almost zero variability in yearly cut-off grade and mill tonnes. In the upper right graph of Figure 8, the cut-off grade could be plotted as 25 lines, but the 25 lines would almost all plot exactly on top of each other. Of the 21 years of milling for 25 simulation realisations, only five of the 525 combinations resulted in a different yearly cut-off versus the most common yearly cut-off that provided 4 Mt of yearly mill feed. In hindsight, this makes sense because when defining a single mining sequence with constrained mill capacity the optimum cut-off will be a balancing cut-off which is based on the grade tonnage distribution which is also fixed once the yearly mining is fixed. The only time that a lower cut-off grade would be selected would be if the balancing cut-off falls below the mine constrained breakeven cut-off—these are the five data points where low simulated gold prices create a high mine constrained cut-off. For these five data points, yearly mill tonnage is slightly below 4.0 Mt.

Performing a deterministic optimisation using an annual metal price forecast equal to the average of the 25 realisations (thick black line in Figure 8 upper left) yielded an identical mining rate. The deterministic plan's reported NPV was <$1M lower than the stochastic optimisation as the deterministic plan did not realise the small benefit of the five out of 525 year × realisation combinations where cash flow is improved by excluding a small portion of low-grade tonnes from the mill feed.

Multi-scenario stochastic price forecasts with stochastic orebody uncertainty

For the final example, stochastic optimisation was performed with uncertainty simulated using 45 equally weighted permutations of five metal price scenarios and nine orebody realisations. The five metal prices all start at $2000/oz with varying constant yearly growth rates of +4 per cent, +2 per cent, +1 per cent, 0 per cent and -2 per cent (same five price scenarios as second deterministic example in Figure 5). The nine orebody realisations all started from the base orebody model but reflect mineralised zone mean grade and selectivity uncertainty in terms of the nine permutations of mean grades of 0.9, 1.0 and 1.1 gpt and relative standard deviations of 0.8, 1.0 and 1.2 (same nine orebody models as fourth deterministic example in Figure 7).

The stochastic solution solves for one expected NPV maximising mining sequence and chooses a cut-off grade for each year of each simulated realisation (Figure 9 top) resulting in a 22-year mine plan. As discussed in the previous example, once a mining sequence/rate is fixed, most of the cut-off grades are balancing cut-off grades which are a function of the grade tonnage distribution rather than metal price. Accordingly, many of the cut-off grade profiles in Figure 9 would plot exactly on top of each other, so the Figure 9 cut-off graph includes a small vertical offset depending on price to allow the lines to be seen.

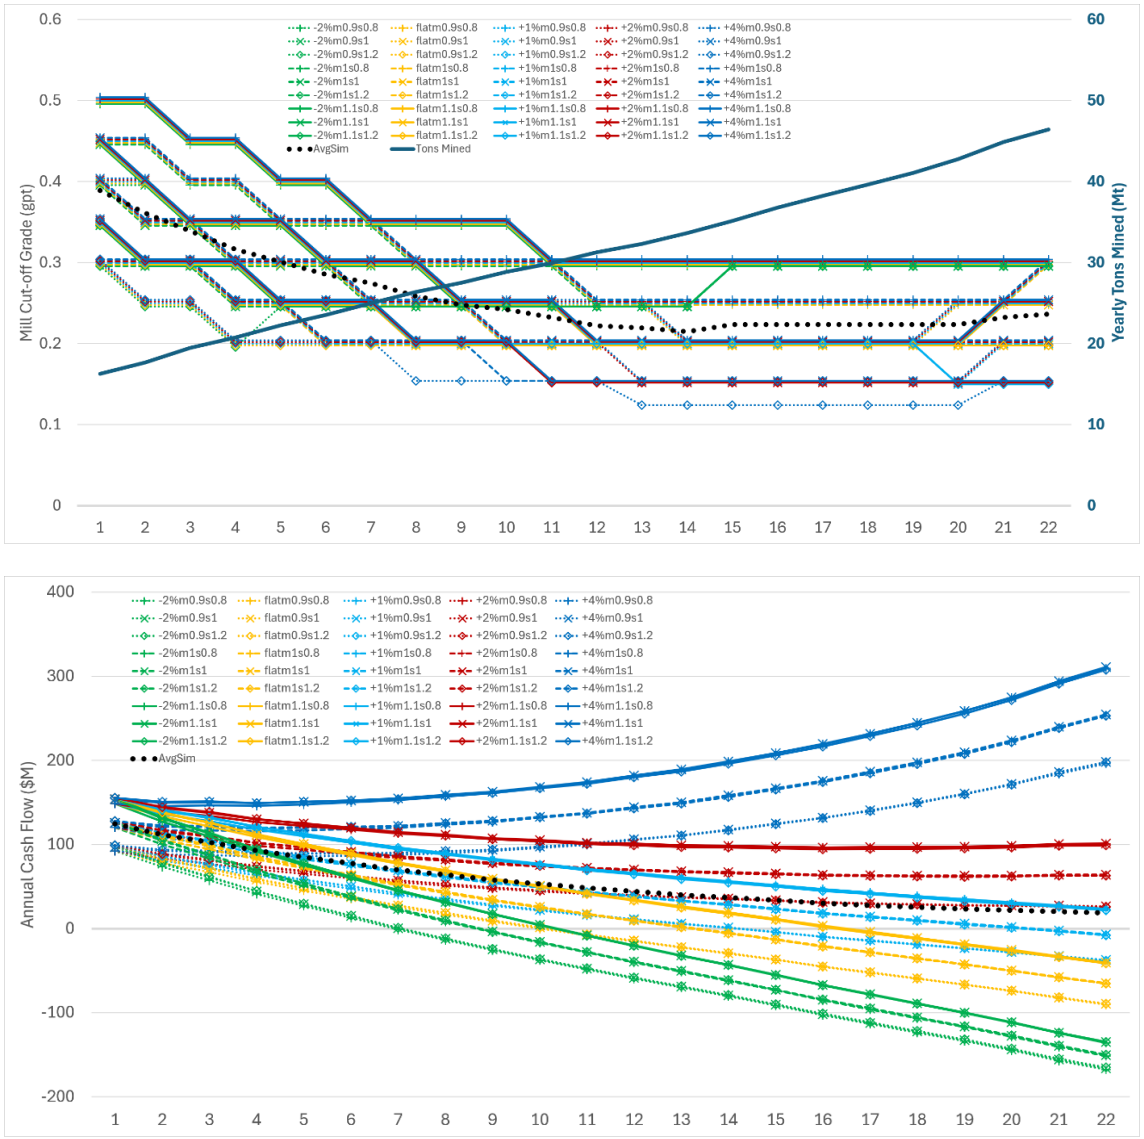


FIG 9 – Stochastic optimisation output under price and orebody uncertainty: mining rate and cut-off grade solution (top) and yearly cash flow uncertainty (bottom). 45 realisations with colour by price scenario, line style by mean mineralisation grade (gpt), symbol by relative standard deviation. Cut-off graph adds a small vertical offset by price scenario to reduce the number of lines that plot on top of each other.

For any metric of interest, 45 values can be evaluated for each year to show all combinations of price × orebody uncertainty. This type of data visualisation (Figure 9 bottom) might be helpful to explain how price and orebody uncertainty impact cash flow uncertainty over time. In year one, orebody mean grade uncertainty drives cash flow uncertainty with only a small amount of variability caused by uncertainty in relative standard deviation. Cash flow uncertainty grows over time as the metal price forecast scenarios diverge. In year 1, cash flow uncertainty ranges from $90M to $160M; by year 6, cash flow uncertainty ranges from $10M to $160M. All 45 realisations provide positive yearly cash flow through year 6, but by year 13, 40 per cent of the realisations result in negative annual

cash flow. This type of information could be used to prioritise the importance of reducing orebody uncertainty and/or building efficient plan optionality that can be turned on/off at higher/lower prices.

Because the stochastic optimisation process focuses on identifying the single best mine plan given the initial knowledge of uncertainty, it does not capture the full value that a real options analysis could reveal. With a real options analysis, the solution is not a single mine plan. Instead, the mine plan is allowed to change as more information becomes available as metal prices are realised and/or deposit model uncertainty decreases with additional drilling. With more information, an early shutdown option would be exercised to avoid years with negative annual cash flow.

Stochastic evaluation versus stochastic optimisation

For the previous two stochastic optimisation examples, a deterministic optimisation using the average of all price realisations and the average distribution of all orebody realisations provided a similar plan to the plan identified by stochastic optimisation. Accordingly, the stochastic optimisation did not provide insights beyond what could have been provided by a sensitivity analysis evaluating the deterministic plan result against all price and model scenarios.

A stochastic evaluation process (Hoerger and Dagdelen, 2025a) can test the deterministic plan (or any fixed mine plan) against the distribution of price and orebody realisations (either all permutations or a sample of the permutations). Each metal price/orebody realisation can be analysed independently with a relatively small MILP model (or other automated cut-off grade adjustment process). If there is a large gap between the average NPV from the stochastic evaluation versus the NPV from the deterministic plan, then stochastic optimisation may be warranted.

In the first stochastic optimisation example based on 25 Ornstein-Uhlenbeck price simulations, deterministic optimisation using the average price for each year yielded a plan with slightly lower NPV ($994.551M) than the stochastic optimisation plan ($994.565M). Note that this section of the paper uses an excessive amount of significant figures to demonstrate that the different optimisation processes generate nearly identical results for the considered examples. Stochastic evaluation of the deterministic optimisation plan showed exactly the same NPV as the stochastic optimisation plan because the deterministic optimisation yielded exactly the same mine sequence.

In the second stochastic optimisation example based on five price scenarios times nine orebody scenarios, deterministic optimisation with average yearly price and average grade/tonnage distribution showed a slightly different mine plan than the stochastic optimisation. The deterministic plan kept the mill full at 4.00 Mt in all 22 years and showed an NPV of $775.69M. Stochastic evaluation of the deterministic plan showed that in some year × orebody × price combinations where the breakeven cut-off was higher than the balancing cut-off, the mill would not be full (average of 3.95 Mt/a) resulting in an NPV of $770.75M which is only slightly lower than the stochastic optimisation NPV of $770.98M.

Computational notes

The MILP formulations for this paper were implemented in a C program named Peakfinder which creates a model file that is solved using IBM's CPLEX™ solver software. MILP solutions are read back into Peakfinder for summary reporting and exported to Excel for graphing and analysis. MILP solution tolerances were set to 0.01 per cent.

The largest MILP problem, with five price simulations and nine orebody simulations, created a stochastic optimisation MILP formulation with about 1000 binary variables, 167 000 linear variables, 90 000 equations and 3.5 million non-zero coefficients. Over 95 per cent of the linear variables and non-zero coefficients are associated with Mine to Destination variables for 37 increments × 2 destinations × 50 time periods × 45 simulations. Accordingly, the problem size would grow significantly by adding stockpiling or an additional destination or shrink by reducing the number of grade increments, time periods and/or simulations. Using a laptop computer (2.1 GHz, 32 GB, 16 cores), CPLEX™ solution times were less than 4 mins. All deterministic cases solved in less than 10 sec. The integer variables were related to very simple sequencing constraints and mill fixed costs, so it is difficult to project these solution times to more complex cases.

Note that a deterministic optimisation would not have the ×45 factor for number of Mine to Destination variables and number of non-zero coefficients leading to a much smaller MILP problem. Accordingly, because stochastic evaluation can analyse a mine plan against each realisation separately with a deterministic optimisation of the process flows, MILP problem size can be kept manageable and stochastic evaluation could easily loop through 45 realisations or even 4500 realisations.

CONCLUSIONS

Because metal price forecasts have several uses in managing a mining business, it is normal for a company to have multiple price scenarios for developing mine plans and additional scenarios for financial evaluation. Traditionally, price scenarios for mine planning might be based on static prices, while financial evaluation prices might show more variation between near-term and long-term pricing. Some price scenarios that make sense for financial evaluation and business planning purposes can have unintended consequences if used for mine planning purposes. For example, a significant future price decrease provides a conservative business planning view of future cash flows, but if that scenario is used for mine planning, it can drive significant increases in near-term mining rates and cut-off grades.

Developing mine plans using multiple simple real price scenarios can build an understanding of options and sensitivities without unintended consequences and/or excessive complexity. Static yearly metal prices provide the simplest scenarios, while smoothly increasing price scenarios can give insights into the potential benefits of postponing production from material which currently shows low or negative margins. Once mine plans are created, more variable annual price forecasts or stochastic price forecasts can be applied to evaluate key financial scenarios or assess plan uncertainty. Clear communication and use of standardised price scenarios helps to ensure comparable mine plans, financial plans, and strategic alignment across multiple sites and projects.

Standard industry practice for large gold companies is to create mine plans assuming a conservative flat metal price. The paper shows that conservative price assumptions for planning can prevent value from being realised. When mining rate is not constrained, using a higher metal price (closer to spot pricing) would tend to result in higher cut-offs and higher mining rates and could create both higher near-term cash flow and long-term value, even if prices eventually revert to lower levels. Although these results may be less applicable to other deposits and mining complexes, the study highlighted the potential benefits to be identified by developing mine plans using multiple metal price scenarios.

The paper also examined the use of stochastic metal price models. The examples considered showed negligible expected NPV benefits from stochastic mine plan optimisation incorporating price uncertainty. However, a stochastic model of prices can still be useful for characterising uncertainty in strategic plan outputs such as annual cash flows and NPV's. Most of this utility can be captured via stochastic evaluation and sensitivity analysis of a deterministic mine plan without the need for solving a full-sized stochastic mine plan optimisation.

REFERENCES

Clark, L and Dagdelen, K, 2023. Practical Mine Planning and Design, in *SME Surface Mining Handbook* (ed: P Darling), pp 29–67 (SME: Englewood).

Dagdelen, K, 1992. Cutoff grade optimization, in *23rd APCOM Proceedings* (ed: Y C Kim), pp 157–165 (SME: Littleton).

Dimitrakopoulos, R and Lamghari, A, 2022. Simultaneous stochastic optimization of mining complexes – mineral value chains: an overview of concepts, examples and comparisons, *International Journal of Mining, Reclamation and Environment*, 36(3):443–460.

Elkington, T and Gould, J, 2012. Real option value optimisation, in *Proceedings of the Project Evaluation Conference 2012*, pp 35–40 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Gillespie, D T, 1996. Exact numerical simulation of the Ornstein-Uhlenbeck process and its integral, *Physical Review, E*, 54(2):2084–2091.

Hoerger, S and Dagdelen, K, 2025a. Mine Plan Risk Assessment and Grade Uncertainty Characterization using Geostatistical Conditional Simulation: Gold Mine Case Study, in *Proceedings of the Minexchange 2025 SME Annual Conference*, preprint 25–057 (SME: Denver).

Hoerger, S and Dagdelen, K, 2025b. Comparison of stochastic vs deterministic open pit mine production schedule optimisation using mixed integer linear programming with conditional orebody simulations, in *Proceedings of the APCOM Conference 2025,* pp 663–683 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Lane, K, 1988. *The Economic Definition of Ore—Cut-off Grades in Theory and Practice* (Mining Journal Books: London).

Newman, A, Rubio, E, Caro, R, Weintraub, A and Eurek, K, 2010. A Review of Operations Research in Mine Planning, *Interfaces*, 40(3):222–245.

Rendu, J M, 2014. *An Introduction to Cut-off Grade Estimation*, 2nd edition (SME: Englewood).

Schwartz, E and Smith, J E, 2000. Short-Term Variations and Long-Term Dynamics in Commodity Prices, *Management Science*, (46)7:893–911.

Visnjic, M, 2018. Mineral Asset Valuation Under Price Uncertainty Using Real Options, MSc Thesis (Colorado School of Mines: Golden).

Wolsey, L, 2021. *Integer Programming*, 2nd edition (Wiley: Hoboken).

Cut-off grade and economic optimisation

A case study of blending-based cut-off selection in an open pit iron ore mine

C Almeida[1] and R Peroni[2]

1. Mine Planning Manager, Vallourec Inc, Belo Horizonte, Brazil. Email: caioalmeida.eng@gmail.com
2. Professor, Mining Engineering Department, Universidade Federal do Rio Grande do Sul, Porto Alegre RS, Brazil. Email: peroni@ufrgs.br

ABSTRACT

The determination of the cut-off grade and pit optimisation are fundamental steps in open pit mine planning, as they define the economic limits of the deposit and influence the mining scheduling. Traditional optimisation methods, such as the Lerchs–Grossmann algorithm, are based on the individual economic evaluation of blocks, considering the difference between revenues and costs. However, this approach presents limitations in iron ore deposits, in which the value of the material depends on the final product quality, often obtained through blending.

This paper presents a case study for defining the cut-off grade for deposits with these characteristics, incorporating blending into the pit optimisation process and, ultimately, into mine planning and scheduling. The approach considers the cut-off grade as the result of an integrated process involving geological, operational, and economic factors. The study presents a comparative analysis between the traditional method, which is defined as base case, and an approach based on blending scenarios, evaluating impacts on profitability, product quality, and mineral resources utilisation.

The methodology involves defining incremental cut-off grade scenarios between 38 per cent and 45 per cent Fe grade, followed by block valuation, pit optimisation, pushback definition, and mining scheduling. For each scenario, economic modelling was carried out with revenue estimates based on market indices (Platts) and the Freight on board (FOB) netback market pricing formula, in addition to cash flow and net present value (NPV) calculation. The study was applied to the Pau Branco Mine, in the Quadrilátero Ferrífero (MG), characterised by banded iron formations and high variability of grades.

Two alternatives were tested for each scenario: a capacity-constrained case, considering plant throughput limits, and a production-maximisation case that maintained the same life-of-mine while assuming no processing capacity constraints.

The results demonstrate that block profitability is not controlled exclusively by total Fe grade. In the base-case economic model, the lowest-grade profitable block contained approximately 38 per cent Fe, whereas some blocks with grades above 50 per cent Fe generated negative economic values. This behaviour was primarily associated with differences in particle-size distribution and its associated iron and contaminants grades. These findings highlight the necessity of incorporating grades according to the different particle-size distributions and processing responses into block valuation to determine the cut-off grade.

Compared with the traditional economic valuation approach (base case), the incorporation of blending constraints into cut-off grade optimisation resulted in improved project value. The 42 per cent Fe and 44 per cent Fe cut-off grade scenarios achieved the highest economic performance, increasing NPV by approximately 2 per cent and 4 per cent, respectively, when compared to the base case. In addition to a superior NPV result, both scenarios satisfied the product quality specification, maintaining silica grades in the +1 mm fraction lower or equal to the 7 per cent limit (7.0 per cent for the 42 per cent Fe scenario and 5.7 per cent for the 44 per cent Fe scenario). In contrast, the base-case scenario presented a silica grade of 7.4 per cent in the +1 mm fraction, exceeding the specification limit according the required product quality criteria.

Although the 44 per cent Fe scenario generated the highest NPV, the 42 per cent Fe scenario provided a more balanced outcome by preserving a larger plant-feed inventory and extending mine life while remaining compliant with product quality requirements.

Under the maximised-production strategy, the economic differences between scenarios were significantly reduced. The 40 per cent to 42 per cent Fe cut-off grade scenarios achieved NPV values within 1 per cent of the base case, while the 41 per cent Fe scenario generated the highest undiscounted cash flow, approximately 2 per cent above the base case.

It is concluded that the integration of blending into mine planning allows for a more integrated approach of value generation, contributing to decisions aligned with the requirements of the iron ore market.

INTRODUCTION

The optimisation and selection of final pits are fundamental steps in open pit mine planning, as they define the economic limits of mineral exploitation and a feasible sequence for deposit extraction. Traditionally, optimisation algorithms such as Lerchs–Grossmann (LG) are used, based on the individual economic feasibility analysis of each block. Valuation is performed as the difference between the estimated revenue, based on the block grade and selling price; and the mining and processing costs.

In the case of iron ore, this approach presents limitations, as the product is sold as a concentrate with defined specifications. Thus, the economic value of a block, if analysed individually, may not reflect its actual contribution to the project's value, since lower-grade materials, which individually would not be profitable, may only become viable if combined through blending with higher-grade blocks, while still generating products within specifications. Therefore, the analysis based on the economic cut-off grade of the individual block becomes limited. Feasibility then depends on combinations between blocks, increasing the complexity of optimisation. In this context, it becomes necessary to incorporate blending into block valuation and the definition of the cut-off grade. The adoption of the traditional method may lead to suboptimal decisions or jeopardise the ore reserves by excluding materials that could generate value when blended.

The proposal of this study consists of conducting a comparative analysis of mine scheduling based on the definition of different cut-off grades, in which the lowest grade studied was defined based on the processing plant technological limitations and the capacity to generate products according acceptable market specifications.

Stone *et al* (2018) developed the BLASOR system, an optimisation tool designed for strategic mine planning in blended iron ore operations involving multiple pits. The method employs mixed-integer linear programming (MILP) to determine optimal mining sequences and define ultimate pit limits and mining phases while simultaneously considering operational and geotechnical constraints, processing capacities, and product quality requirements. The authors demonstrated that integrating these variables into a global optimisation framework enables the maximisation of discounted cash flow (DCF) over the life-of-mine, reducing the reliance on empirical decision-making traditionally used in the development of mining phases.

The objective is to evaluate, through a case study, the definition of the cut-off grade based on blends resulting from mine scheduling and that considers this definition as part of an integrated process considering geological, operational and economic aspects. In addition, it lies in the application of a revenue model that explicitly considers the distribution of iron grades among particle-size fractions and its effect on product quality, block value and cut-off grade selection.

The general workflow consists of defining the scenarios, which are distinguished by the ranges of cut-off grades to be studied lying between the economic and technological limits. These scenarios will then be compared to a base-case scenario using the traditional method. Subsequently, planning variables are assigned to the model, pit optimisation is performed, pushbacks are defined, and mining scheduling is carried out. For each scenario, economic modelling and analysis are conducted, including revenue estimation, cash flow projection, and net present value (NPV) calculation. The objective is to find a cut-off grade that ensures the preservation of ore reserves while minimising project value losses (if any).

Iron ore – technical and economic aspect

Iron is one of the most abundant elements in the Earth's crust, occurring in various mineral forms. The main minerals of economic interest are oxides, due to their high-grade and ease of concentration, with hematite, magnetite, goethite, and siderite being the most relevant minerals (Carvalho *et al*, 2014).

Geologically, iron ore deposits are classified into five categories: bedded sedimentary, hydrothermal, volcanic, metamorphic/deformational, and supergene alteration deposits. Bedded sedimentary, resulting in banded iron formations (BIF), which host the largest global ore reserves. In Brazil, regions such as Quadrilátero Ferrífero, Carajás, and Corumbá are mainly composed of BIFs, referred to as itabirites and jaspilites, with grades ranging between 20 per cent and 55 per cent. Hematitic bodies also occur, with higher grades (above 64 per cent), as well as lateritic cover (*Canga*), with grades between 50 per cent and 65 per cent, but with a higher presence of contaminants (Carvalho *et al*, 2014).

From a metallurgical standpoint, products are classified as lump ore (with particle size fractions between 200 mm and 12.5 mm), sinter feed (12.5 to 0.15 mm), and pellet feed (below 0.15 mm). Quality is defined according the iron grade (with market reference grades of 58 per cent, 62 per cent, and 65 per cent) and by contaminant grades such as silica, alumina and phosphorus which affect performance in metallurgical processing stages.

Iron ore pricing is dynamic and based on the relationship between supply and demand, with strong influence from China, the world's largest steel producer. Its economic policies and industrial demand directly impact global prices. The pricing model has evolved from fixed annual contracts to a more flexible system based on the spot market and financial instruments. Due to volatility, price indices such as Platts IODEX, Argus, and Mysteel have become relevant, particularly for ore with 62 per cent Fe Cost and Freight (CFR) China.

The Platts IODEX is the main international benchmark index, including variations such as 58 per cent, 61 per cent, 62 per cent, and 65 per cent Fe, with the index selected according to the quality of the ore to be sold. The final price is adjusted by premiums or penalties according to deviations in specifications, considering elements such as Fe, SiO_2, Al_2O_3, P, and moisture content.

The Platts IODEX (P_{platts}) represents the price in dollars per ton delivered at the port of Qingdao, China, and is also used as a reference in negotiations in Brazil. In these cases, maritime freight costs must be deducted, with a reference value obtained from the C3 Freight market index, in addition to expenses related to port services and domestic logistics.

Equation 1 presents the formula known as FOB Netback.

$$Price = (Platts\ 62 + VIU) * (1 - H_2O\%) - C3\ Freight - Logistic\ Costs \quad (1)$$

where Price is the realised product price expressed in US$/wmt (wet metric tonne); Platts 62 is the Platts IODEX 62 per cent Fe reference price expressed in US$/dmt (dry metric tonne); VIU (value-in-use) represents the quality adjustment applied to the reference price and is expressed in US$/dmt; H_2O is the product moisture content; C3 Freight is the benchmark freight rate from Brazil to China expressed in US$/wmt; and Logistic Costs represent port handling, inland transportation and other logistics-related expenses expressed in US$/wmt.

Figure 1 schematically presents the pricing model for products marketed under the Ex Works-regime.

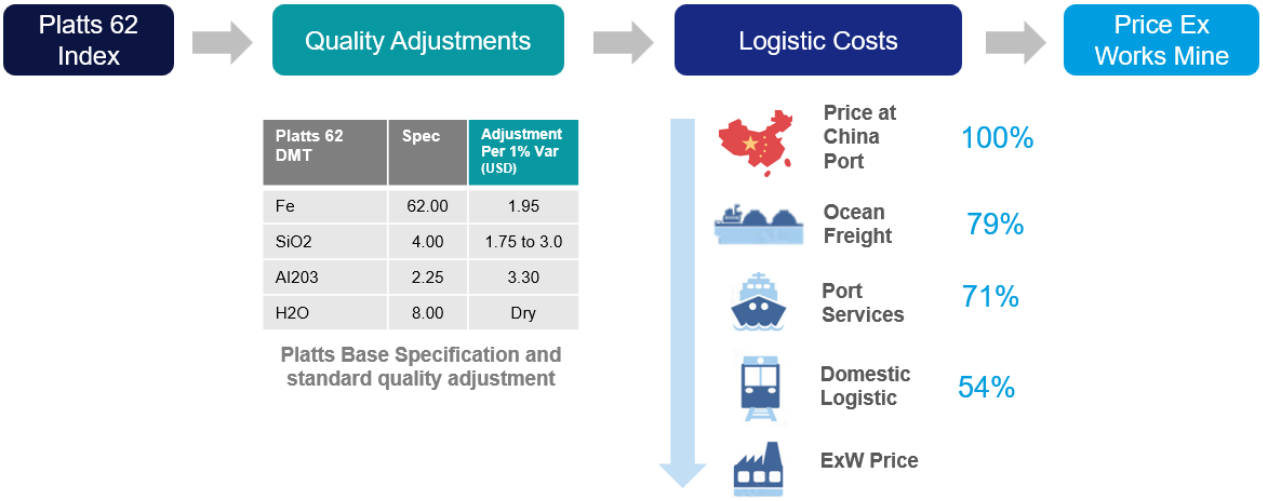


FIG 1 – Pricing model (ex works) for iron ore fines (sinter feed and pellet feed).

The value-in-use (VIU) calculates the actual value of the ore beyond the Platts index reference price, considering premiums and penalties associated with quality. For this methodology, the elements subject to adjustments are iron (Fe), alumina (Al_2O_3) and silica (SiO_2) and they are calculated differently, as presented in Equations 2, 3 and 4.

$$VIU\ Fe = \left(Fe - Fe_{Ref.}\right) \times\ Fe_{Dif} \tag{2}$$

$$VIU\ SiO_2 = \begin{cases} -0{,}5 \times SiO_{2_{Dif\,(3-4.5)}}, 3 < SiO_2 \leq 4.5 \\ -(0{,}5 \times SiO_{2_{Dif\,(3-4.5)}} + 2 \times SiO_{2_{Dif\,(4.5-6.5)}}), 4.5 < SiO_2 < 6.5 \\ -\left(0{,}5 \times SiO_{2_{Dif\,(4-4.5)}} + 2 \times SiO_{2_{Dif\,(4.5-6.5)}} + (SiO_2 - 6.5) \times SiO_{2_{Dif\,(6.5-9)}}\right), SiO_2 \geq 6.5 \end{cases} \tag{3}$$

$$VIU\ Al_2O_3 = \begin{cases} (Al_2O_{3_{Ref}} - Al_2O_3) \times Al_2O_{3_{Dif(1-2.5)}} \times 0{,}7, Al_2O_3 \leq 2.25 \\ (Al_2O_{3_{Ref}} - Al_2O_3) \times Al_2O_{3_{Dif(1-2.5)}}, 2.25 < Al_2O_3 < 2.5 \\ (Al_2O_{3_{Ref}} - Al_2O_3) \times Al_2O_{3_{Dif(2.5-4.0)}}, Al_2O_3 \geq 2.5 \end{cases} \tag{4}$$

In the equations, Fe, SiO_2 and Al_2O_3 represent the quality specifications of the saleable product. The terms identified as *Ref.* correspond to the benchmark specifications adopted for the Platts 62 per cent Index (Table 1). The terms identified as *Dif* represent the market-derived price differentials used in the VIU calculation (Figure 2), reflecting the premiums and penalties associated with deviations from the reference product specifications (Table 1).

TABLE 1

Reference product specifications of the Platts 62 per cent Fe index.

Plats 62%	Spec.
Fe	62%
SiO_2	4%
Al_2O_3	2%
H_2O	8%

Iron ore

Per 1% differentials (Range 60-63.5% Fe)

	Symbol	Within Min-Max	$/dmt	Change
Per 1% Fe	IOMGD00	60-63.5% Fe	1.75	0.00
Per 1% Alumina	IOADF10	1-2.5% Al2O3	2.80	0.00
	IOALE00	2.5-4% Al2O3	3.10	0.00
Per 1% Silica	IOALF00	3-4.5% SiO2	0.70	0.00
	IOPPS10	4.5-6.5% SiO2	1.40	0.00
	IOPPS20	6.5-9% SiO2	2.10	0.00

FIG 2 – Quality differentials used in the VIU calculation (source: S&P Global).

Cut-off grade definition

The cut-off grade definition constitutes an essential parameter in the economic evaluation of mineral deposits, as it establishes the limit from which the material is classified as ore or waste. According to Hustrulid, Kuchta and Martin (2013), the cut-off grade represents the 'minimum grade required for a block or portion of ore to be mined and processed profitably', being directly related to operating costs, commodity price, and beneficiation efficiency.

According to Hall (2014), the main objective in determining the cut-off grade is to maximise the economic outcome of the mining project. However, depending on the characteristics of the project, additional objectives may arise, such as extending the life-of-mine or complying with specific operational and strategic constraints.

Rendu (2013) highlights that the cut-off grade has a significant influence on several aspects of mining operations, such as the volume of material mined and processed, the average plant feed grade, the cash flow, and the longevity of the project. In addition to financial impacts, less measurable factors must also be considered, including environmental, social, ethical, and political aspects, which also affect the feasibility and sustainability of the project.

This evaluation must encompass all stakeholders involved, such as shareholders, financial institutions, analysts, workers, consumers, suppliers, local communities, governments at different levels, non-governmental organisations, and future generations. In this context, management must balance diverse interests, ensuring that planning decisions, such as the definition of cut-off grades, are aligned with the broader strategic objectives of the company.

The choice of cut-off grade therefore involves a trade-off between short and long-term outcomes. Higher cut-off grades tend to increase immediate profitability and reduce payback but may shorten the life-of-mine. Conversely, lower cut-off grades may extend operations, favouring more durable economic benefits, greater social stability, and more responsible use of mineral resources.

This concept is used in mine planning analyses, allowing adaptation to market fluctuations and contributing to maximising NPV, according to Lane (1988).

Lane (1964) introduced the concepts of breakeven cut-off grade (BCOG) and marginal cut-off grade (MCOG), which are widely used in the economic evaluation of mineral deposits. The BCOG represents the minimum grade required to cover mining, processing and administrative costs, as presented in Equation 2.

The breakeven cut-off grade is given according to Equation 5:

$$BCOG = \frac{m+c+a}{y(s-r)} \tag{5}$$

where m is the mining cost, c the processing cost, a the general and administrative cost, y the metallurgical recovery, s the selling price, and r represents the costs taken from the gross sales of metal product, such as the transportation, refining, and royalties costs.

When revenue covers processing and administrative costs but not the mining cost, the concept of marginal (or mill) cut-off grade applies, according to Equation 6:

$$MCOG = \frac{c+a}{y(s-r)} \tag{6}$$

Block model valuation

The block model spatially represents mineral deposits using prisms of compatible dimensions, organised in Cartesian coordinates within a regular 3D grid, and is widely used in grade estimation and subsequently in mine planning (Sinclair; Blackwell, 2002; Hustrulid, Kuchta and Martin, 2013). Figure 3 illustrates the block model displaying the spatial distribution of iron (Fe) grades throughout the deposit.

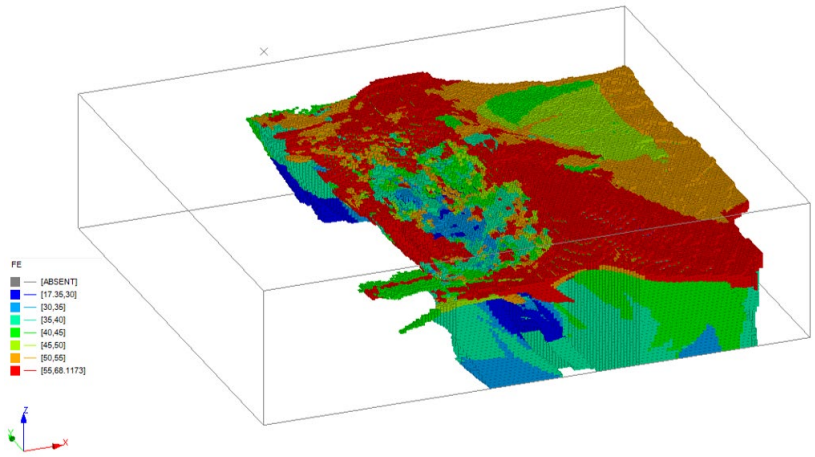


FIG 3 – Block model represented by ranges of iron grade.

Each block has attributes estimated and derived mostly from drilling data and geostatistical estimates, including geological, geotechnical, metallurgical, grades, density, rock types, recoveries, and economic attributes, such as costs, and selling price (Souza, 2013). Before determining the mine planning, an economic value is assigned to each block—based on the relationship between revenues and costs—allowing the evaluation of its feasibility and the optimisation of project returns (Souza, 2013).

Valuation corresponds to the assignment of monetary value to the block model, a step based on the construction of the profit function, which is fundamental for the optimisation stage and optimal pit selection. The calculation generates a net value (positive or negative), considering mining and processing revenues and costs, and is strongly influenced by the block grade (Diedrich, 2012). The concepts of Lane (1988) and Rendu (2013) are applied here.

Once the blocks are economically valued, it is necessary to define the delimitation of the economic pit. This process is carried out using an optimisation algorithm that aims to determine the final pit configuration considering technical, operational, and geological constraints, aiming maximum profitability with minimum cost and risk (Hustrulid, Kuchta and Martin, 2013). According to Caccetta and Hill (2003), these systems apply a parametric optimisation method by varying selling prices and/or costs to assess the project's sensitivity to these parameters and to generate a set of nested pits. The optimisation methodology typically uses the concept of NPV, such that the extraction sequence that generates the highest financial return is selected (Lerchs and Grossmann, 1965). According to Peroni (2002), the optimum is defined by algorithms that ensure maximum profitability, highest NPV, and best resource utilisation.

METHODOLOGY

The methodology addressed in this study considers an evaluation of mineral resources by assuming that the cut-off grade should not be defined in isolation, but rather as the result of an iterative process of evaluating blending scenarios.

Mine planning approaches available in literature incorporate processing decisions directly into the production scheduling optimisation. Johnson (1968) proposed one of the earliest formulations in which the decision of whether material should be processed or discarded is determined within the scheduling process itself. More recently, Caccetta and Hill (2003) and Boland *et al* (2009) developed mathematical programming formulations capable of simultaneously considering mining sequences,

processing capacities, blending requirements and block-level processing decisions. These approaches optimise extraction and processing decisions directly within the scheduling model.

The present work adopts a different approach. Rather than embedding cut-off grade selection within a large-scale production scheduling optimisation model, a scenario-based methodology was used in which cut-off grade are first defined and subsequently evaluated through pit optimisation, pushback design, mine scheduling and economic analysis. The objective of this study is therefore not to develop or compare scheduling algorithms, but to evaluate the impact of blending-based cut-off grade selection through an operational case study. The proposed workflow is presented in the following flow chart (Figure 4).

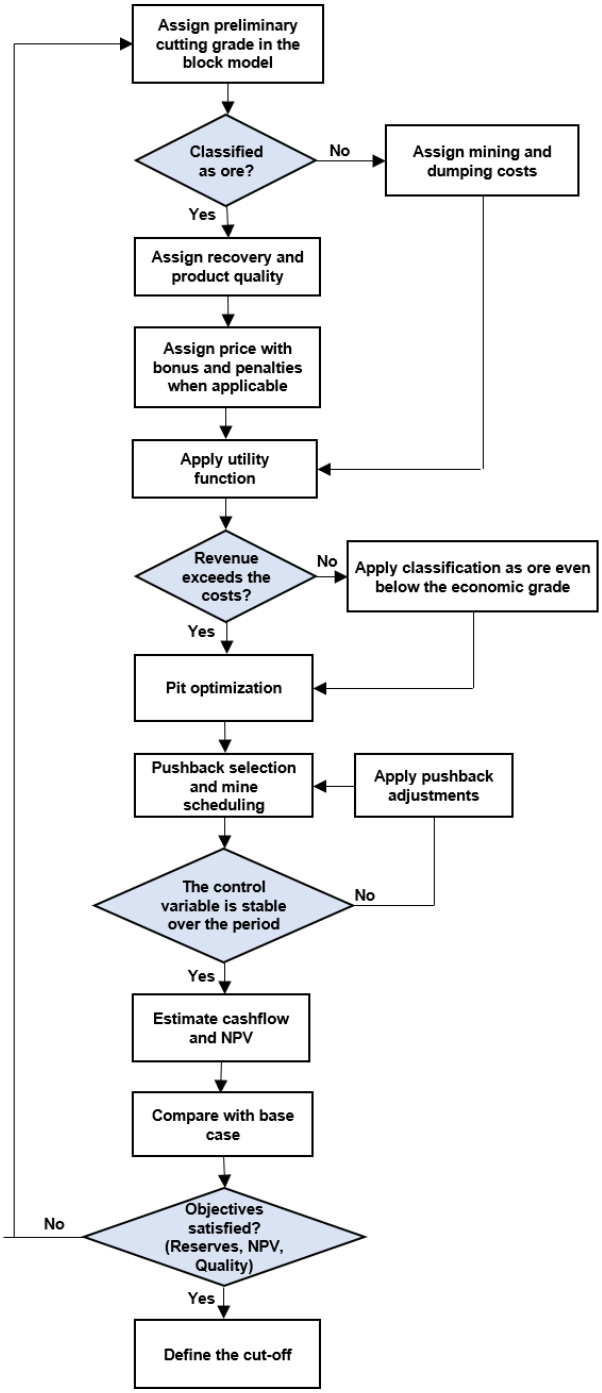


FIG 4 – Methodology flow chart.

For the present study, the cut-off grade was first defined, and the blocks of the model were preliminarily classified as ore or waste based on geological domains and technological limitations of the plant. It is important to note that at this stage the distinction between the scenarios to be analysed

is established. Subsequently, the profit function is structured, in which revenue and cost parameters are estimated, based on price, exchange rate, recovery, selling and operating costs.

Typically, according to the traditional approach, blocks with revenue lower than operating costs are classified as waste, while blocks with positive values are classified as ore. At the pit optimisation stage, all scenarios are optimised using the same pit selection criterion, corresponding to the pit shell with the highest NPV. This approach ensured a consistent comparison framework. After the optimisation and selection of the final pit, pushbacks were selected and the mine scheduling was carried out, typically aiming at stability of feed grades and the stripping ratio over time. If this condition is not met, adjustments are made to the pushbacks and scheduling to ensure this stability.

Finally, economic modelling is conducted based on the results of the scheduling and the average grades per period, resulting from the blending process, estimating revenues, costs, cash flow, and NPV for each scenario. The final choice considers financial, technical, and strategic criteria. The methodology applied to a case study in the Quadrilátero Ferrífero (MG), including comparative analysis of scenarios.

For the optimisation stage, two cases were considered: firstly (traditional scenario), only blocks with positive benefit are classified as ore; secondly, the ore classification criteria is applied regardless of benefit, such that all material previously classified as ore must be destined to the plant feed during optimisation.

For cases such as iron ore, the main difficulty at this stage lies in defining the revenue based on the product quality estimated for each individual block. In sales contracts, it is common for pricing to be established as a percentage of the Platts index (eg 50 per cent), associated with a product, chemical and physical specification. However, when using this model as a reference, situations are observed in the model where iron grades fall below or above the specified priced range, in which the revenues of this block may be underestimated or overestimated.

Case study

Block model and assignment of mine planning parameters

The adopted block model has dimensions of 15 × 15 × 15 m, with sub-blocks up to 5 × 5 × 5 m in geological contact areas. The total grades, particle size distribution, and grades by particles fraction were estimated by ordinary kriging for the attributes Fe, Si, Al, P, Mn, and Loss on Ignition (LOI).

Table 2 presents the block model attributes used to characterise particle-size distribution and iron grade. Variables W1 to W4 represent the mass proportions of each size fraction, whereas FE1 to FE4 correspond to the iron grades estimated within those fractions. The total iron grade is represented by the variable FE.

TABLE 2

Description of the estimated fields in the block model and example of description by fraction.

Attribute	Description
W1	≥ 6.3 mm
W2	≥ 1.0 and <6.3 mm
W3	≥ 0.15 and <1.0 mm
W4	<0.15 mm
FE1	Fe Grade within the W1 Fraction
FE2	Fe Grade within the W2 Fraction
FE3	Fe Grade within the W3 Fraction
FE4	Fe Grade within the W4 Fraction
FE	Fe Total

The particle-size fractions are mutually exclusive intervals and together represent 100 per cent of the block mass.

$$W1 + W2 + W3 + W4 = 100\% \quad (7)$$

In addition to the individual grades estimated for each particle-size fraction, the total iron grade was also estimated in resource model. However, it can also be calculated as the weighted average of the iron grades estimated within each particle-size fraction, according to Equation 8.

$$FE = \frac{W1 \times FE1 + W2 \times FE2 + W3 \times FE3 + W4 \times FE4}{W1 + W2 + W3 + W4} \quad (8)$$

For the simulation of products based on the block model, it is assumed that the sum of fractions W1 and W2 (denominated W12), corresponds to material with particle size greater than 1 mm, which represents the combination of lump products and sinter feed. The ore fraction with particle size below 1 mm is defined by the sum of fractions W3 and W4 (denominated W34) which are fed to the magnetic concentration process. The estimation of mass recoveries and product quality was carried out based on the combination of technological characterisation studies of the ore lithological groups and the correlation of historical data with the particle size fractions of the model.

The preliminary cut-off grade among the scenarios ranged from 38 per cent to 45 per cent Fe, and the ORETYPE attribute was defined, as shown in Table 3. Considering the minimum cut-off grade of 38 per cent was defined based on the processing plant's technological limitations and its capacity to produce saleable products. Conversely, the upper limit represents the economic threshold for low-grade itabirite classification, above which, the benefit function is consistently positive for this case study.

TABLE 3

Field for preliminary cut-off grade definition.

ORETYPE	Value
<38 per cent – Waste	0
>=38 per cent – Ore	1

In this case study, pricing using the FOB netback methodology is proposed for fine products (sinter feed and pellet feed), usually intended to be exported. For lump products, predominantly directed to the domestic market, prices were defined based on a percentage of the Platts index, adjusted according to iron grade ranges of the product.

The economic value assigned to each block was calculated as the difference between the estimated revenue and the mining, processing, general and administrative (G&A) and selling costs. Revenue was estimated by multiplying the mass recovery of each product by its respective selling price and summing the resulting values.

$$BV = \sum_{i=0}^{n} (Y_i(P_i - s_i)) - (m + c + a) \quad (9)$$

where BV is the economic value assigned to the block; Y_i is the mass yield of product i; P_i is the realised price of saleable product i (Brazilian Real/wet metic ton (BRL/wmt)); s_i is the selling cost associated with product i; m is the unit mining cost; c is the unit processing cost; and a is the general and administrative cost per ton of processed material.

The selling price was calculated using the FOB netback methodology for sinter feed and pellet feed products, which are predominantly sold in the export market.

For Natural Pellet Ore (NPO) and Hematitinha products, the FOB netback methodology was not applied, as these products are predominantly sold in the domestic market. Instead, pricing was based on realised domestic market transactions and estimated as a percentage of the corresponding Platts index, according to predefined iron grade ranges. The assumed values was derived from benchmarking market prices as Table 4.

TABLE 4

Price factors applied to natural pellet ore and hematitinha relative to the Platts 62 per cent Fe index.

Iron grade	Natural pellet ore (% of Platts 62)	Hematitinha (% of Platts 62)
50<	0	0 per cent
>50 and <=52	15 per cent	13 per cent
>52 and <=54	22 per cent	19 per cent
>54 and <=56	27 per cent	23 per cent
>56 and <=58	45 per cent	38 per cent
>58 and <=60	65 per cent	55 per cent
>60	70 per cent	60 per cent

The cost-per-ton assumptions were derived from current operational data obtained from the operation studied here. Due to confidentiality constraints, the actual values can not be disclosed. However, Table 5 presents the main cost lines applied. Important to state that the same cost per ton was applied across all evaluated scenarios.

TABLE 5

Cost components considered in the block economic valuation.

Description	Unit
Mining costs	BRL/t mined
Dumping costs	BRL/t waste
Processing costs	BRL/t processed
Tailings management costs	BRL/t tailings
Taxes	BRL/t sealable product
Handling costs	BRL/t sealable product
Selling costs	BRL/t sealable product
CAPEX sustaining	BRL/t sealable product

Recovery estimation is primarily based on the W12 and W34 particle-size fractions. The recovery of natural products is calculated from the W12 fraction (particle size >1 mm), after applying a penalty factor that accounts for the generation of fines during mining, crushing, handling, and processing operations. The fines coming from W12 fraction is transferred to the W34 fraction (particle size <1 mm), which is treated in the magnetic concentration plant.

$$Y = (W12 - d) + (W34 + d) \times R \tag{10}$$

Where *d* is the penalty factor and *R* is the magnetic concentrator recovery, which both are affected by the lithology. The first term of the equation is the recovery of natural products while the second term is the pellet feed recovery. To calculate the recovery of each natural product, the expected proportion of each product is applied in the first term of the equation.

Geotechnical parameters were defined for each geotechnical domain, with variations in face angle, bench height, and consequently the overall slope angle.

The physical mining constraints considered in this study were based on the areas allocated for the ore processing plant infrastructure and the current licensed limits, corresponding to an extension of approximately 200 m beyond the existing pit crest. To respect this constraint, an inclusion ‘Pit Limit’ polygon was established, and the pit optimisation process was kept within this polygon, as shown in Figure 5.



FIG 5 – Pit optimisation boundary after applying physical constraints.

Pit optimisation, pushback selection, and sequencing

Pit optimisation was performed independently for each scenario. The final pit was selected using the revenue-factor that resulted in the maximum NPV for each scenario.

For pushback definition, width adjustments were established with offsets at every 30 m from the current pit crest, and a sequence of pushbacks in ‘vertical strips’ was generated’ (Figure 6). This pushback control allows the combination of high and low-grades, as well as the proportion of ore types in the plant feed, since there is marked vertical variation from the lithologies.

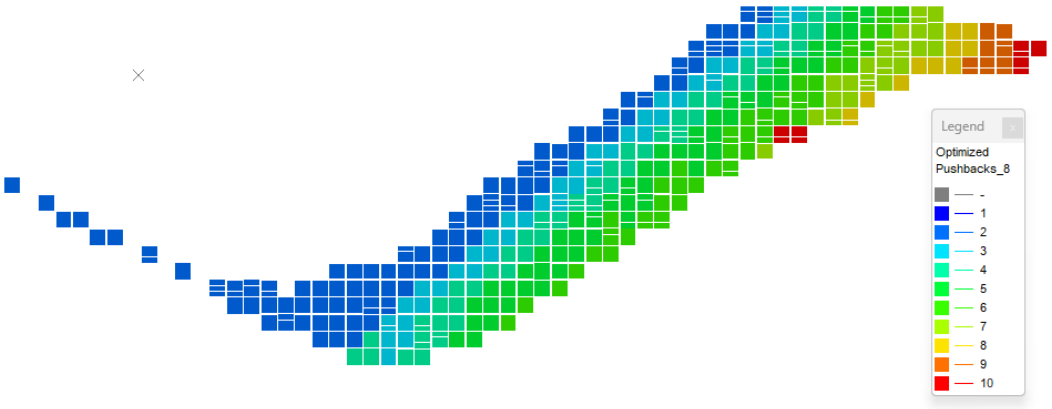


FIG 6 – Example section of pushbacks generated using limits defined by regular offsets.

From an operational perspective, the adoption of vertical pushbacks is viable given the size of the mining fleet and the geometric characteristics of the pit. Current mining operations utilise 50 t excavators and 8×4 haul trucks, with pit geometries defined by 15 m bench heights and 7.4 m berm widths and 12 m wide haul roads. The combination of relatively narrow pushbacks and a fully mechanical excavation process, the annual vertical advance reaches up to 15 benches per annum.

This ‘Vertical Strips’ approach enabled improved stabilisation of both grade distribution and lithotype availability for the mine scheduling stage. Without such control, the optimisation algorithm tends to prioritise early extraction in the upper portions of the deposit, where higher-grade material is typically concentrated, often associated with lateritic covers and supergene-enriched itabirites. As a result, mining sequences would initially target high-grade ore, followed by a gradual decline in grade as operations progress into intermediate zones composed of different types of intermediate and poor itabirites. Toward the final stages, higher grades are encountered again with the extraction of hematite bodies located at the bottom of the pit.

Scheduling targets were established with the objective of maintaining a stable feed grade and strip ratio throughout the life of the mine. For each scenario, the target values were based on the average grades and strip ratio of the selected ultimate pit. For example, if the average SiO_2 grade in the +1 mm fraction (SI12) was 7 per cent and the strip ratio was 1.2 t/t for the 42 per cent Fe cut-off grade scenario, these values were adopted as reference targets for each scheduling period. The scheduling process therefore aimed at minimising temporal variability while preserving the overall characteristics of the selected mining inventory.

Mine schedules were developed using annual production periods that were constrained by processing capacity of 7.5 Mt/a and mining capacity of 15 Mt/a. Economic evaluation was performed through discounted cash flow analysis using a discount rate of 14 per cent.

Ore stockpiling and reclaiming activities were not explicitly modelled in this study. This assumption is consistent with the operating conditions of the case study, where plant feed is predominantly supplied with ore coming directly from the mine. Furthermore, the available stockpile capacity is approximately 25 000 t, which is negligible considering the annual production rate. Consequently, stockpile management was considered to have limited influence on long-term scheduling decisions and was excluded from the analysis. While this simplification reduces the operational flexibility represented in the model, it is not expected to materially affect the comparative evaluation of the cut-off grade scenarios.

Subsequently, financial modelling was carried out using scheduling results to estimate revenues, costs, free cash flow, and NPV, considering average grades per sequencing period.

To complete the analysis, the same scenarios were tested, assuming a higher annual feed rate, considering the scheduled plant-feed tonnes tend to increase as the cut-off grade is reduced.

RESULTS AND DISCUSSION

Economic cut-off grade for base case

Considering the base-case scenario, evaluated using the traditional block economic valuation approach, the lowest-grade profitable block contains approximately 38 per cent Fe, whereas some blocks with iron grades as high as 53 per cent Fe present negative economic values. This result demonstrates that block profitability is not controlled exclusively by the total iron grade, but rather by the interaction between particle-size distribution, recovery, and product pricing. The example comparing these two blocks is presented in Table 6.

TABLE 6

Comparison between the lowest-grade profitable block and the highest-grade non-profitable block in the base-case scenario.

FE	FE	FE12 +1 mm	FE34	+1 mm Recovery	-1 mm Recovery	Mass Recovery	+1 mm Revenue	-1 mm Revenue	Profit
Lowest grade profitable block	38.0	42.6	34.3	35.3	30.6	65.9	0.0	249.0	51 976
Highest grade non-profitable block	52.7	47.9	55.9	40.8	19.3	60.1	0.0	299.2	-624

As shown in Table 6, both blocks present Fe grades below 50 per cent in the +1 mm fraction (FE12). According to the pricing assumptions adopted in this study, natural products with iron grades below 50 per cent are not considered as saleable ore and therefore generate no revenue from the natural products stream. Consequently, the economic viability of both blocks depends primarily on the value generated by the pellet feed production route.

Despite its substantially lower total iron grade, the profitable 38 per cent Fe block exhibits higher recovery, with natural products and pellet feed recoveries of 35.3 per cent and 30.6 per cent, respectively, resulting in a total mass recovery of 65.9 per cent. The pellet feed generated from this material produces sufficient revenue to offset mining and processing costs, yielding a positive profit of approximately BRL 51 976 per block.

In contrast, the non-profitable block, despite containing a significantly higher total iron grade of 52.7 per cent Fe, presents unfavourable processing characteristics. This material belongs to a colluvial lithology composed predominantly of clay-rich matrix and goethitic minerals, which exhibit low magnetic susceptibility. As a result, the pellet feed recovery is limited to only 19.3 per cent, substantially lower than that observed for the profitable block. Although the pellet feed product presents a higher iron grade (F_{34} = 55.9 per cent Fe) and consequently a higher unit revenue, the reduced recovery limits the total amount of saleable product generated, resulting in a negative economic value.

These results highlight the importance of incorporating particle-size-specific grades and metallurgical responses into block valuation. Two blocks with different total iron grades may exhibit opposite economic outcomes when recovery, product quality, and marketability are considered. This behaviour explains why the traditional economic classification of ore and waste cannot be based solely on total Fe grade and reinforces the importance of evaluating the distribution of iron among the different particle-size fractions when estimating block value and defining cut-off grade strategies.

Scenarios analysis

The results of the mine scheduling based on the cut-off grade range between 38 per cent and 45 per cent of Total Fe, are presented in Table 7. Among the analysis scenarios, the scenario based on cut-off grade defined by the difference between costs and revenues, referred as the traditional approach (or base case), is also included.

TABLE 7

Summary of results.

Description	Unit	COG 38	COG 39	COG 40	COG 41	COG 42	COG 43	COG 44	COG 45	Base case
Total material	Mt	233	248	242	249	249	249	248	248	249
Plant feed tonnes	Mt	145	143	131	123	114	108	106	104	116
Fe total	%	53.7	54.2	55.4	56.6	57.8	58.6	59.0	59.2	57.2
Fe +1 mm	%	57.7	58.3	59.4	60.4	61.3	61.8	62.1	62.2	61.1
Si +1 mm	%	12.6	11.7	10.0	8.5	7.0	6.1	5.7	5.5	7.4
Strip ratio	t/t	0.60	0.73	0.84	1.03	1.19	1.30	1.35	1.38	1.16
Sealable product	Mt	89	91	81	77	72	69	68	67	73
W12	%	41	43	40	40	41	41	41	41	41
Mass recovery	%	61.5	63.7	61.7	62.4	63.3	64.0	64.3	64.4	62.6
Tailings	Mt	56	52	50	46	42	39	38	37	43
Dump	Mt	144	157	161	172	177	180	180	181	177
Sales price	% of Platts 62	37%	40%	47%	50%	52%	52%	53%	54%	51%
LoM	Years	20	20	18	17	16	15	15	14	16
Undiscount cash flow	Relative to Base Case	-24%	-10%	1%	1%	0%	-3%	-3%	-3%	0%
Discount rate	%	14%	14%	14%	14%	14%	14%	14%	14%	14%
NPV	Relative to Base Case	-43%	-28%	-11%	-4%	2%	2%	4%	2%	0%

Table 7 summarises mine scheduling outputs and economic indicators obtained for the evaluated cut-off grade scenarios. The results demonstrate that cut-off grade selection has a significant influence on resource utilisation, product quality, waste and tailings generation, mine life, and overall project value.

As expected, reducing the cut-off grade increases the quantity of material classified as ore and sent to the processing plant. Plant feed increases from 104 Mt at the 45 per cent Fe scenario to 145 Mt at the 38 per cent Fe scenario, representing an increase of approximately 39 per cent, mainly due to the incorporation of lower-grade itabirite that would otherwise be classified as waste. Consequently, mine life increases from 14 years at the higher cut-off scenarios to 20 years at the lowest cut-off grades.

The larger ore inventory obtained at lower cut-off grades is accompanied by a progressive reduction in head grade quality. Average total Fe decreases from 59.2 per cent at the 45 per cent Fe scenario to 53.7 per cent at the 38 per cent Fe scenario, while silica in the +1 mm fraction increases from 5.5 per cent to 12.6 per cent, increasing the reliance on blending to maintain product quality specifications.

The commercial impact of this quality deterioration is reflected in the realised sales price, which increases from 37 per cent of the Platts 62 benchmark at the 38 per cent Fe scenario to 54 per cent at the 45 per cent Fe scenario. Although lower cut-off grades generate larger tonnages of saleable products, the reduction in product quality affects unit revenue through lower prices and higher quality penalties.

From the mining and sustainability perspective, lower cut-off grades improve resource utilisation and reduce the amount of material classified as waste. Dumped material decreases from 181 Mt at the

45 per cent Fe scenario to 144 Mt at the 38 per cent Fe scenario, corresponding to a reduction of approximately 20 per cent. Similarly, the strip ratio decreases from 1.38 t/t to 0.60 t/t, indicating that a greater proportion of the material within the final pit is classified as ore. However, increased processing of low-grade ores results in higher tailings generation, which rises from 37 Mt to 56 Mt across the same range of scenarios.

The economic indicators reveal a trade-off between resource preservation and economic performance. The lowest cut-off scenarios maximise plant feed tonnage and mine life but present the lowest project value. Relative to the base case, the 38 per cent Fe scenario reduces undiscounted cash flow by 24 per cent and NPV by 43 per cent, while the 39 per cent Fe scenario reduces NPV by 28 per cent, indicating that the additional ore recovered is insufficient to offset the impact of lower grades and reduced product quality.

Economic performance improves progressively as cut-off grade increases. The 40 per cent and 41 per cent Fe scenarios generate marginal improvements in undiscounted cash flow relative to the base case, while NPV becomes positive from the 42 per cent Fe scenario onwards. The highest NPV is obtained for the 44 per cent Fe scenario, approximately 4 per cent above the base case.

The comparison between the base case and the blending scenarios is particularly relevant. The base case presents the same total mined material as the 41 per cent, 42 per cent, and 43 per cent Fe scenarios, indicating that the differences observed arise primarily from ore classification rather than from changes in ultimate pit limits. In this context, the 42 per cent Fe scenario increases average Fe grade from 57.2 per cent to 57.8 per cent, reduces silica in the +1 mm fraction from 7.4 per cent to 7.0 per cent, maintains the same waste dumping and life-of-mine, and improves NPV by approximately 2 per cent. These results demonstrate that, in this case study, incorporating blending considerations into cut-off grade selection can improve both product quality and project value.

Cut-off grade selection

The selection of the preferred cut-off grade was not based exclusively on the maximisation of a single economic indicator. Instead, it was treated as a constrained mine planning decision in which economic performance, plant-feed tonnes, and product quality compliance were evaluated simultaneously.

The results indicate that the highest NPV was obtained for the 44 per cent Fe cut-off grade scenario, with an increase of 4 per cent relative to the base case, while the 40 per cent and 41 per cent Fe scenarios presented the highest undiscounted free cash flow, approximately 1 per cent above the base case.

In iron ore operations, product quality requirements represent a fundamental planning constraint because ore value is ultimately determined by the characteristics of the final commercial products. For the deposit analysed in this study, silica grade in the material above 1 mm was identified as the critical quality parameter, with a specification limit of 7 per cent.

As the cut-off grade decreases, additional lower-grade material can be incorporated into the mine inventory, but maintaining product quality through blending becomes progressively more difficult. Scenarios below the 42 per cent Fe cut-off showed non-compliance with quality requirements, with silica grades varying from 8.5 per cent to 12.6 per cent in the W12 fraction, whereas the 42 per cent Fe scenario exactly meets the specification limit.

The 44 per cent Fe cut-off scenario presented a 4 per cent higher NPV than the base case; however, total plant-feed tonnes were reduced by approximately 7 per cent compared to the 42 per cent Fe scenario (106 Mt versus 114 Mt). Similarly, the 45 per cent Fe scenario achieved the highest realised sales price (54 per cent of Platts 62) but resulted in the lowest plant-feed tonnage (104 Mt) and the shortest mine life (14 years).

Considering economic performance, quality compliance, and resource preservation, the 42 per cent Fe cut-off grade was selected as the preferred scenario. Relative to the 44 per cent Fe scenario, it increases plant-feed tonnes by approximately 8 Mt (7.5 per cent), saleable product by approximately 4 Mt (5.9 per cent), and mine life by one year while maintaining compliance with product quality

requirements. Although the NPV is approximately two percentage points lower than the maximum observed value, the scenario provides a more balanced utilisation of mineral resources.

Evaluation of scenarios with increased production

For all scenarios, the same annual feed rate to the beneficiation plant was initially adopted. However, as the cut-off grade is reduced, the available ore naturally increases, allowing reconfiguration of the production strategy. Therefore, scenarios incorporating higher production rates were developed for lower cut-off grades, while maintaining a constant mine life, to assess the impact of cash flow acceleration and its effect on NPV.

It is important to state that in the maximised-production scenarios no additional expansion capital or sustaining capital associated with capacity increases was included in the economic evaluation as well as the same cost per ton was assumed. Therefore, the results should be interpreted as a comparative assessment of the potential value associated with accelerated extraction rather than as a expansion study.

The hypothesis assumes that, although lower grades reduce unit profitability, increased production may be partially offset by earlier revenue realisation (Table 8).

TABLE 8

Summary of results for maximised production scenarios.

Description	Unit	COG 38	COG 39	COG 40	COG 41	COG 42	COG 43	COG 44	COG 45	Base case
Total material	Mt	233	248	242	249	249	249	248	248	249
Plant feed tonnes	Mt	145	143	131	123	114	108	106	104	116
Fe total	%	53.7	54.2	55.4	56.6	57.8	58.6	59.0	59.2	57.2
Fe +1 mm	%	57.7	58.3	59.4	60.4	61.2	61.8	62.1	62.2	61.1
Si +1 mm	%	12.6	11.7	10.0	8.5	7.0	6.1	5.7	5.5	7.4
Strip ratio	t/t	0.60	0.73	0.84	1.03	1.19	1.30	1.35	1.38	1.16
Sealable product	Mt	89	91	81	77	72	69	68	67	73
W12	%	41	43	40	40	41	41	41	41	41
Mass recovery	%	61.5	63.7	61.7	62.4	63.3	64.0	64.3	64.4	62.6
Tailings	Mt	56	52	50	46	42	39	38	37	43
Dump	Mt	144	157	161	172	177	180	180	181	177
Sales price	% of Platts 62	38%	41%	47%	50%	52%	52%	53%	54%	52%
LoM	Years	15	15	15	15	15	15	15	14	15
Undiscount cash flow	Relative to Base Case	-22%	-9%	0%	2%	1%	-3%	-3%	-3%	0%
Discount rate	%	14%	14%	14%	14%	14%	14%	14%	14%	14%
NPV	Relative to Base Case	-28%	-15%	-1%	0%	0%	-2%	1%	-1%	0%

From an economic perspective, the increased production strategy reduces the differences observed between scenarios. Compared to the fixed-production cases, the impact on NPV becomes less pronounced, indicating that part of the economic penalty associated with lower cut-off grades can be mitigated through earlier cash generation. Nevertheless, the lowest cut-off grade scenarios still present inferior economic performance, with NPV reductions of 28 per cent and 15 per cent for the 38 per cent Fe and 39 per cent Fe scenarios, respectively.

The best undiscounted cash flow performance was obtained for the 41 per cent Fe scenario, which generated approximately 2 per cent higher cash flow than the base case while maintaining a mine life of 15 years. Similarly, the 40 per cent, 41 per cent, and 42 per cent Fe scenarios presented NPV results very close to the base case, ranging between -1 per cent and 0 per cent.

Overall, the maximised-production analysis indicates that cut-off grades between 40 per cent and 42 per cent Fe provide the best balance between accelerated cash generation, product quality, and project value. On the other hand, there are operational challenges, such as expanding processing capacity, managing increased mine haulage, and possible logistical constraints, in addition to greater exposure to market risks.

Study limitations

Although the results obtained are consistent and relevant, some limitations must be considered. The modelling of the blending process was carried out in a simplified approach, not fully incorporating operational constraints such as stockpiling and reclaiming capacity, as well as logistical aspects. In addition, geological uncertainties associated with the block model were not explicitly considered, which may impact the predictability of the results.

Treatment of production-capacity increases in the maximised-production scenarios is also a limitation. Although higher plant feed rates may require additional investments in mining equipment, processing facilities, materials handling systems, infrastructure, or other supporting assets, these capital requirements were not modelled. Consequently, the economic benefits associated with accelerated production may be overestimated and should be validated through detailed engineering and capital cost studies.

Another important limitation refers to the dependence on the adopted scheduling, since different strategies may result in distinct combinations of materials and, consequently, different economic outcomes. Finally, the economic analyses were based on fixed market assumptions, not incorporating uncertainty scenarios or stochastic approaches, as suggested by Dimitrakopoulos (2018). These simplifications made, indicate clear opportunities for future work.

CONCLUSIONS

This study proposed and evaluated an alternative methodology for defining the cut-off grade in iron ore deposits, incorporating blending logic into mine scheduling in comparison with scenario selection using the traditional cut-off grade definition method based on NPV maximisation. The results indicate that the approach more comprehensively captures operational and economic aspects not considered by traditional methods based on individual block evaluation.

The comparison between the traditional approach based on the individual economic value of blocks and the proposed approach reveals substantial differences in results. In the classical approach, based on algorithms such as Lerchs–Grossmann (Lerchs and Grossmann, 1965), blocks are assigned economic values before optimisation and their value is generally evaluated independently of the final blend composition.

The scenario with a 42 per cent Fe cut-off grade presented the best balance between quality, operational stability, and value generation, reinforcing that the cut-off grade must integrate technical, economic, and market variables. However, the study presents limitations, such as simplifications in blending, absence of geological uncertainties, and deterministic assumptions, indicating opportunities for improvement.

In addition, it was concluded that block value is strongly influenced by the distribution of iron grades among particle size fractions, particularly in the material above 1 mm, which exerts greater influence on revenue generation. Consequently, total Fe grade alone is not sufficient to explain block profitability in this deposit.

It is recommended to incorporate blending into iron ore mine planning, treat the cut-off grade as a dynamic variable, include more detailed operational constraints, adopt stochastic approaches, and integrate more robust geometallurgical models. It also highlighted the need for studies on the

utilisation of lower-grade ores, focusing on more efficient processing routes to economically enable these resources and contribute to the sustainability and profitability of the project.

ACKNOWLEDGEMENTS

The lead author would like to express sincere gratitude to their wife and family for their continuous support, encouragement, and understanding throughout this professional journey. Their daily support has been fundamental to the author's personal and professional development.

The lead author is especially thankful to their advisor, Rodrigo Peroni, for his guidance, availability, and valuable support in the development of this work. His insights and dedication were essential for the completion of this study.

The lead author also extends their appreciation to their managers at Vallourec, who provided them with the time, resources, and support necessary to carry out this work. Their trust and encouragement made this achievement possible.

REFERENCES

Boland, N, Dumitrescu, I, Froyland, G and Gleixner, A M, 2009. Lp-based disaggregation approaches to solving the open pit mining production scheduling problem with block processing selectivity, *Computers and Operations Research*, 36(4):1064–1089.

Caccetta, L and Hill, S P, 2003. An application of branch and cut to open pit mine scheduling, *Journal of Global Optimization*, 27:349–365.

Carvalho, P S L, Silva, M M, Rocio, M A R and Moszkowicz, J, 2014. Minério de ferro, *BNDES Setorial*, 39:197–234.

Diedrich, J, 2012. Avaliação econômica de blocos em modelos de reservas minerais, MSc Dissertation, Universidade Federal do Rio Grande do Sul, Porto Alegre.

Dimitrakopoulos, R, 2018. *Advances in applied strategic mine planning* (Springer: Cham).

Hall, B, 2014. Cut-off grades and optimising the strategic mine plan, *Cut-off Grades and Optimising the Strategic Mine Plan,* Spectrum Series 20 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Hustrulid, W, Kuchta, M and Martin, R, 2013. *Open pit mine planning and design*, 2nd edn (CRC Press: Boca Raton).

Johnson, T B, 1968. Optimum open-pit mine production scheduling, University of California, Berkeley.

Lane, K, 1964. Choosing the optimum cut-off grade, *Colorado School of Mines Quarterly*, 59(4):811–829.

Lane, K, 1988. *The economic definition of ore: cut-off grades in theory and practice* (Mining Journal Books: London).

Lerchs, H and Grossmann, I F, 1965. Optimum design of open pit mines, *Mining Engineering*, 17:47–54.

Peroni, R L, 2002. Análise da sensibilidade do sequenciamento de lavra em função da incerteza do modelo geológico, PhD thesis, Universidade Federal do Rio Grande do Sul, Porto Alegre.

Rendu, J, 2013. *An introduction to cutoff grade estimation*, 2nd ed (Society for Mining, Metallurgy and Exploration (SME): Colorado).

Sinclair, A J and Blackwell, G H, 2002. *Applied mineral inventory estimation* (Cambridge University Press: Cambridge).

Souza, M I, 2013. Metodologia para otimização da cava final considerando a heterogeneidade do depósito mineral, PhD Thesis, Universidade de São Paulo, São Paulo.

Stone, P, Froyland, G, Menabde, M, Law, B, Pasyar, R and Monkhouse, P H L, 2018. Blasor: Blended Iron Ore Mine Planning Optimisation at Yandi, Western Australia, in *Advances in Applied Strategic Mine Planning* (ed: R Dimitrakopoulos), pp 39–46 (Springer: Cham).

Leveraging underutilised statistical techniques and low-cost data to drive optimisation and generate value in strategic life-of-mine planning – a case study

L Lawrence[1] and K Whitchurch[2]

1. Technical Services Specialist, Whittle Consulting, Melbourne Vic 3127, Australia. Email: leigh@whittleconsulting.com.au
2. Partner, Whittle Consulting, Jakarta DKI, Indonesia. Email: keith@whittleconsulting.com.au

INTRODUCTION

Early project development requires capital allocation decisions to be made under conditions of limited data and compressed timelines. In such environments, the ability to rapidly extract meaningful insights from low-cost, routinely collected data is essential to identify value drivers, quantify upside potential, and guide targeted follow-up investment.

Grindability is a first-order driver of processing performance and project economics, yet it is often overlooked and modelled with less rigour than metal grade. In strategic mine planning, the objective is to identify the primary system bottleneck and optimise value flow through it by modelling the drivers that constrain or enhance performance. For comminution-intensive deposits, this bottleneck is typically plant power rather than tonnage throughput, making grindability and its variability throughout the deposit critical inputs for economic assessment and optimisation.

While comminution test work performed on composites of representative samples remains a routine and essential practice for flow sheet design, such approaches can average out critical spatial variability and mask meaningful heterogeneities in ore hardness. Where direct non-composite comminution data are limited, correlated multivariate data sets may provide useful proxies for hardness and processing response at a level suitable for enhancing orebody knowledge and strategic planning.

Multivariate statistical approaches to geometallurgical domaining have been applied previously (Keeney and Walters, 2011; Sepúlveda, Dowd and Xu, 2018; Rajabinasab and Asghari, 2019), though uptake in routine practice remains limited. This case study applies principal component analysis (PCA) and k-means clustering, two underutilised statistical methods, to a multivariate data set from a porphyry Cu-Au-Mo deposit. In doing so, it demonstrates how routinely collected geological and geotechnical data can be leveraged to define meaningful grindability domains and model deposit-scale variability in early-stage strategic mine planning.

METHODS AND OUTCOMES

Routine point load index (PLI) measurements were converted to uniaxial compressive strength (UCS)-equivalent values. These data were integrated with downhole geological logging (lithology, alteration, texture etc) and assay results to form a multivariate data set. Multi-component lithology and texture codes were transformed into binary indicator variables, and all inputs were standardised prior to analysis.

Principal component analysis (PCA) was applied to reduce dimensionality and manage multicollinearity. K-means clustering was then applied to the principal component scores to identify internally consistent UCS populations (Figure 1). Cluster solutions from k=3 to k=5 were evaluated, with k=5 selected based on the elbow method. In this case, the combined method identified controls on rock strength that were not evident from conventional exploratory review, downhole UCS review or composite metallurgical test work alone.

Cluster separation was interpreted against geological logging, with lithology emerging as the dominant control on hardness. This provided a defensible basis for modelling rock strength across the deposit using downhole UCS values, domained by lithology. In other deposits, the main control may be alteration intensity, mineralogy, structural fabric or another logged attribute. The transferable element is the use of multivariate statistics to identify the controls on economically meaningful variability before translating them into domains that can be modelled and tested during optimisation.

However, it should be noted that domaining rock strength using k-means cluster groupings is an equally valid approach where correlation with a predefined geological domain is absent.

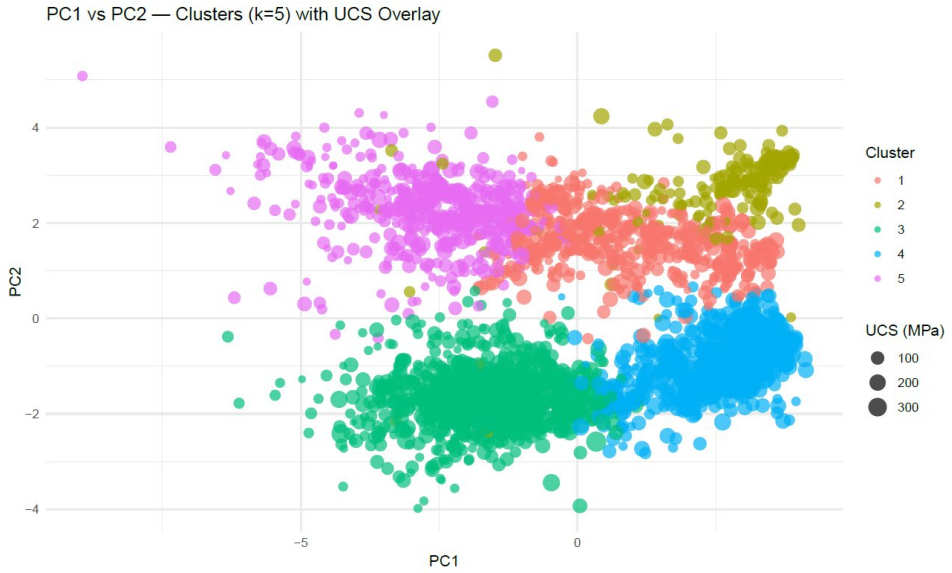


FIG 1 – PC1 versus PC2 with k-means clusters (k=5) and UCS overlay, illustrating correlation of hardness populations in multivariate space. The five clusters were subsequently consolidated into three grindability domains by merging clusters with similar mean Bond Work Index (BWi) estimates.

Bond Work Index (BWi) values were estimated from rock-strength proxies where direct comminution testing was limited. The UCS-BWi relationship proposed by Haffez (2012) was used as the starting point for converting UCS measurements into energy-based parameters. The calculated BWi values were benchmarked against measured composite data to confirm they were reasonable in the context of the available metallurgical test work. This preserved the relative ranking between domains while reducing the risk of overstating material softness.

Using lithology to domain and model UCS globally produced three distinct grindability populations for each zone (Table 1), with mean values subsequently implemented in Prober-F using an energy-constrained plant power model. These replaced the baseline assumption of a uniform BWi of 17 kWh/t for all material. Introducing variable BWi at the pit and phase design stage allowed the optimiser to differentiate material by energy intensity and adjust the life-of-mine schedule accordingly. The objective was not simply to increase tonnes processed, but to allocate limited plant energy to the material that generated the most value – with the expectation that the optimiser would seek this out preferentially.

TABLE 1

Grindability domains and associated UCS and BWi values used for optimisation.

Zone	UCS population	Description	Mean UCS (MPa)	BWi (kWh/t)
North	1	Weak/soft	57.9	14.2
North	2	Moderate	101	17.5
North	3	Strong/hard	145	19.6
South	1	Weak/soft	39.6	12.0
South	2	Moderate	69.5	15.3
South	3	Strong/hard	134	19.2

Relative to the uniform-BWi baseline of the master composite, modelling grindability indicated a potential net present value uplift of approximately 8–15 per cent. This reflected the system-level response of pit design, scheduling, energy-adjusted throughput and metal timing, rather than a single isolated parameter change. Spatial grindability modelling does not create value by itself; rather, it reveals the economic structure of the deposit more clearly. In some deposits that structure may support value uplift. In others, it may expose harder domains, lower realistic throughput, higher energy requirements or the need for additional capital.

CONCLUSIONS

This case study demonstrates that, even in early-stage projects where detailed comminution data are limited, low-cost rock-strength data analysed using mature statistical methods can meaningfully enhance orebody knowledge, preserve economically important variability, and support risk-informed strategic planning. By identifying the physical controls on hardness and creating an optimisation-ready representation of variability, these results can further guide future drilling, data acquisition and effective capital allocation. Although applied here to grindability, the same workflow can be used to target, understand and model any deposit-specific value driver that materially influences project economics and strategic planning outcomes.

REFERENCES

Haffez, G S, 2012. Correlation between Bond work index and mechanical properties of some Saudi ores, *Journal of Engineering Sciences*, 40(1):271–280.

Keeney, L and Walters, S G, 2011. A methodology for geometallurgical mapping and orebody modelling, in *Proceedings of the First International Geometallurgy Conference, GeoMet 2011,* pp 217–225 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Rajabinasab, B and Asghari, O, 2019. Geometallurgical Domaining by Cluster Analysis: Iron Ore Deposit Case Study, *Natural Resources Research*, 28(3):665–684.

Sepúlveda, E, Dowd, P A and Xu, C, 2018. Fuzzy clustering with spatial correction and its application to geometallurgical domaining, *Mathematical Geosciences*, 50(8):895–928.

Cut-off grade as a variable – an emergent property of multi-constraint mine plan optimisation at Evolution Mining's Red Lake operations

B Talbot[1], J Carswell[2] and L Jordan[3]

1. Principal Planning Engineer – Group Long Term Planning, Evolution Mining Limited, Sydney NSW 2000, Australia. Email: blake.talbot@evolutionmining.com
2. Principal Planning Engineer – Group Long Term Planning, Evolution Mining Limited, Red Lake, Balmertown, ON P0V 1C0, Canada. Email: jack.carswell@evolutionmining.com
3. Group Manager Long Term Planning (North America), Evolution Mining Limited, Sydney NSW 2000, Australia. Email: luke.jordan@evolutionmining.com

ABSTRACT

In the 18 months preceding Evolution Mining's (EVNs) FY27 Life-of-Mine (LOM) planning cycle, the Long Term Planning (LTP) team completed a comprehensive Strategic Optimisation process for Red Lake Operations (RLO) in Ontario, Canada. The process evaluated a broad range of strategic alternatives, including ore source selection, capital investment opportunities, processing infrastructure configurations, and development sequencing strategies, with the objective of identifying value-accretive pathways for the operation.

The Strategic Optimisation process included the ability to optimise and trade-off alternative inventory selections, including cut-off grade by orebody. To validate the inventory options selected in the study and translate them into detailed mine plan scenarios, EVN developed and implemented a Variable Cut-off Grade (vCOG) optimisation methodology using the Deswik® APEX (APEX) platform as part of the FY27 LOM optimisation process.

Unlike conventional planning workflows, where cut-off grade is treated as a fixed planning input, the vCOG methodology allows multiple inventory alternatives to be presented to the optimiser for each stope location. The optimiser then selects the inventory option that maximises value within the context of the broader mining system, simultaneously considering schedule timing, operating constraints, capital decisions, and competing opportunities across the operation.

The optimisation was applied across 26 underground block models and approximately 80 000 scheduling activities. More than 60 scenario combinations were evaluated to test the sensitivity of project value to key operational and strategic assumptions. The resulting optimisation outputs provided validation of the preferred strategic direction, generated orebody-level planning guidance, and delivered insight into the relative importance of major operational and capital constraints.

This work demonstrates a practical framework for integrating cut-off grade determination within the scheduling process and highlights how modern optimisation tools can be used to strengthen the link between strategic mine planning, inventory selection, and Life-of-Mine plan development.

INTRODUCTION

Cut-off grade (COG) is one of the most influential parameters in mine planning because it defines the inventory available for extraction and processing. Despite its importance, COG is traditionally determined before schedule optimisation takes place and is subsequently treated as a fixed planning assumption. Under this process, the inventory is established first and the schedule is then optimised within the boundaries of that inventory.

This approach creates a fundamental disconnect. The economic value of a mining operation depends on the interaction between inventory selection, schedule timing, operating constraints, and capital investment decisions. Yet the inventory itself is often defined without considering the economic outcomes of the schedule that it ultimately supports.

Evolution Mining's Red Lake Operations in Ontario, Canada provided an opportunity to explore an alternative approach. In the 18 months preceding the FY27 LOM planning cycle, the LTP team completed a Strategic Optimisation process to evaluate the operation's full strategic opportunity set. This work considered ore source selection, infrastructure investments, development prioritisation,

processing configurations, and sequencing alternatives to identify value-accretive pathways for the operation.

An outcome of the Strategic Optimisation process was the recognition that inventory selection could not be considered independently of the broader mining system. Different operating constraints, capital decisions, and sequencing strategies produced different economically preferred inventories, highlighting the limitations of applying a single static cut-off grade across multiple orebodies and planning periods.

To validate these findings and support their implementation within the FY27 LOM plan, the LTP team developed a Variable Cut-off Grade (vCOG) optimisation methodology, leveraging and extending the capabilities of the APEX platform to evaluate underground inventory selection, capital deployment and processing optionality simultaneously. Deswik® APEX Version 37.7.0 owned by Deswik® is the software used for this body of work. APEX is a schedule optimisation platform that applies mathematical optimisation techniques to maximise project value while satisfying operational constraints. The vCOG feature extends this capability by allowing multiple inventory alternatives to be evaluated for each stope location, enabling cut-off grade selection to occur within the optimisation process rather than before it.

The objective of the work presented in this paper was not to identify a single optimal cut-off grade for RLO. Instead, the objective was to evaluate inventory selection as part of the scheduling process itself, allowing cut-off grade decisions to respond dynamically to schedule timing, operating constraints, and competing opportunities across the operation.

The approach builds upon the theoretical framework described by Hall (2014), which recognised that economically optimal cut-off grades are dependent on the prevailing constraints within the mining system and may vary through time. Advances in optimisation software, cloud computing, and mine planning workflows now allow these concepts to be applied at operational scale.

This paper describes the methodology, implementation, and outcomes of applying vCOG within the FY27 LOM optimisation process at RLO. Particular emphasis is placed on the role of optimisation as a validation and translation tool, bridging the gap between strategic planning outcomes and practical Life-of-Mine planning decisions.

CUT-OFF GRADE THEORY

Break-even and Mortimer

The most widely used approach to cut-off grade (COG) determination is the break-even method, in which the COG is defined as the grade required to recover the direct cost of mining and processing a unit of material:

$$COG_{(break-even)} = \frac{Mining\ Cost + Processing\ Cost}{Commodity\ Price + Recovery}$$

While simple and intuitive, the break-even approach assumes all material above the threshold contributes equally to project value. It ignores the opportunity cost of constrained resources such as hoisting, milling and backfilling. Under these conditions, processing marginal material may displace higher-value ore, resulting in a COG that is lower than the economic optimum.

Mortimer (1950) addressed this limitation by recognising processing capacity as a constrained resource and introducing the concept of a minimum profitable grade. However, the approach remained fundamentally static, assuming a fixed operating state and constraint regime.

Lane's limiting constraint framework

Lane (1988) further advanced cut-off optimisation by recognising that the economically optimal COG depends on the limiting constraint within the mining system. He identified three primary constraint regimes: the mine, the mill and the market.

Under Lane's framework, the optimal COG varies as the binding constraint changes. When mining capacity is limiting, lower-grade material may be processed because spare milling capacity exists.

Conversely, when milling capacity is fully utilised, higher-grade material is favoured to maximise value generated per unit of constrained processing capacity.

Lane's work established that cut-off grade is a function of the operating system rather than a fixed property of the orebody.

Hall's time-dependent cut-off framework

Hall (2014) extended Lane's work by placing cut-off grade optimisation within an NPV framework. In this view, cut-off grade is not simply determined by the current limiting constraint; it is also determined by the timing of cash flows and the opportunity cost of using constrained capacity.

A central insight from Hall's framework is that the optimal cut-off grade is time dependent. Even where costs, recoveries and production capacities remain unchanged, the optimal cut-off will generally decline through the mine life. This occurs because early cash flows have greater value than later cash flows, and because the opportunity cost of processing lower-grade material decreases as the remaining inventory and life-of-mine reduce.

The practical implication is that higher-value material should generally be prioritised earlier, when processing capacity and capital are most valuable. Lower-grade material may become economic later, not because its physical characteristics have changed, but because its position within the remaining mine plan has changed.

For modern underground operations, this effect is amplified by differences between orebodies in grade distribution, operating cost, development access, infrastructure availability, backfill capacity and capital timing. As these factors change through the mine life, so too does the economically preferred cut-off grade. Cut-off grade should therefore be treated as a dynamic planning variable, rather than a fixed economic input.

The gap between theory and practice

Despite the strong theoretical foundation established by Lane and Hall, practical implementation has historically been limited by computational and organisational constraints.

In most planning processes, cut-off grade is established upstream as a fixed input. Schedules are then optimised within the boundaries of a predefined inventory. This creates a circular dependency whereby cut-off grade defines the inventory, the inventory defines the schedule, and the schedule determines project value.

To address this limitation, the RLO strategic optimisation study sought to determine cut-off grade within the optimisation process itself. Building upon the work of Lane and Hall, the RLO life-of-mine Planning team developed a vCOG optimisation methodology that leveraged and extended the capabilities of the APEX platform, allowing inventory selection, scheduling, capital deployment and processing decisions to be optimised simultaneously.

RED LAKE OPERATIONS

Site overview

RLO is a high-grade underground gold mine located in Northwestern Ontario, Canada, and operated by EVN The mine is part of one of North America's highest-grade gold camps.

RLO comprises four underground mines which are processed through two milling circuits to produce gold doré. As of 2025, the Mineral Resource stands at 61 Mt at 3.63 g/t Au for 7.1 Moz of contained gold, inclusive of tailings, while the Ore Reserve totals 18 Mt at 3.35 g/t Au for 2.0 Moz of contained gold. The reported Mineral Resources and Ore Reserves are based on the company's 2025 statement (Evolution Mining Limited, 2025).

RLO is characterised by significant geological and operational complexity comprising high-grade narrow-vein mineralisation, generally steeply dipping orebodies, and disseminated sulfide mineralisation distributed across multiple mining zones. Each orebody exhibits distinct grade distributions, mining geometries and depths (ranging near surface to 2.5 km), access constraints, geotechnical conditions, infrastructure requirements, and operating cost structures. These

characteristics challenge long-term planning and optimal inventory selection as a single cut-off grade applied uniformly across all orebodies is unlikely to reflect the economic characteristics of each system throughout the mine life.

The operation is supported by two paste backfill plants which provide paste fill to facilitate continuous extraction sequencing across two of the four mines. Development activities are distributed across multiple active mining areas and are characterised capital-intensive development investment to production, typical of narrow-vein underground stoping operations.

METHODOLOGY

Deswik® APEX and MILP optimisation

APEX is a cloud-based mine schedule optimisation platform designed to maximise the value of a mining operation while respecting operational, physical, and financial constraints. For this study, APEX was selected for two key reasons. First, it could evaluate the large and complex scheduling problem presented by RLOs 26 underground block models and associated infrastructure. Second, it incorporated a Variable Cut-off Grade (vCOG) capability, allowing cut-off grade to be optimised within the scheduling process rather than treated as a fixed planning assumption.

At its core, APEX uses a Mixed Integer Linear Programming (MILP) optimisation framework. In practical terms, MILP evaluates a very large number of potential scheduling combinations and identifies the sequence of mining activities and capital investment decisions that delivers the highest Net Present Value (NPV) while satisfying all specified constraints. The optimiser effectively performs the same trade-off decisions that a planning engineer would make manually, but at a scale and speed that would be impossible using conventional scheduling approaches.

Each scenario considered approximately 80 000 individual scheduling activities across 26 block models. These activities included stope extraction, development activities, capital projects, and supporting infrastructure decisions. The optimisation was subject to a broad range of constraints, including:

- underground development sequencing and access requirements
- hoisting and haulage capacity limits
- processing plant throughput constraints and debottlenecking options
- stope sequencing and geotechnical constraints
- paste backfill plant capacity
- capital expenditure timing and budget limitations
- interdependencies between mining areas and infrastructure.

The objective of each optimisation run was to maximise life-of-mine NPV while ensuring that all operational constraints were respected and where appropriate, tested against debottlenecking options.

As with most large-scale optimisation problems, the theoretical optimum cannot always be solved exactly within practical computation times. APEX therefore uses advanced solution techniques to converge on near-optimal solutions and reports solution quality through an optimality gap. This metric indicates the maximum potential improvement that could still exist between the reported solution and the theoretical optimum, providing confidence in the quality of the resulting schedule.

APEX model configuration

The APEX model was configured to represent key physical, operational, and economic components of RLOs. Inputs were compiled from the FY27 LOM planning process and included underground stope designs, development infrastructure, productivities and constraints, operating costs (including fixed and variable splits), capital including investment options, and material movement pathways.

Stope inventories were generated using Deswik®.SO across multiple cut-off grade bins for each orebody. Capital development designs were informed by the existing LOM infrastructure and sized to accommodate the broadest inventory footprint represented by the vCOG alternatives. Operating development was designed for the lowest COG bin stopes. A bespoke dependency rule set was then implemented to ensure the designs and sequences supported all inventory choices. Economic inputs included commodity price assumptions, fixed and variable operating costs, and capital expenditure estimates associated with infrastructure and development options.

The model incorporated the operation's material flow network, including processing, stockpiling, and alternative ore source pathways (example in Figure 1).

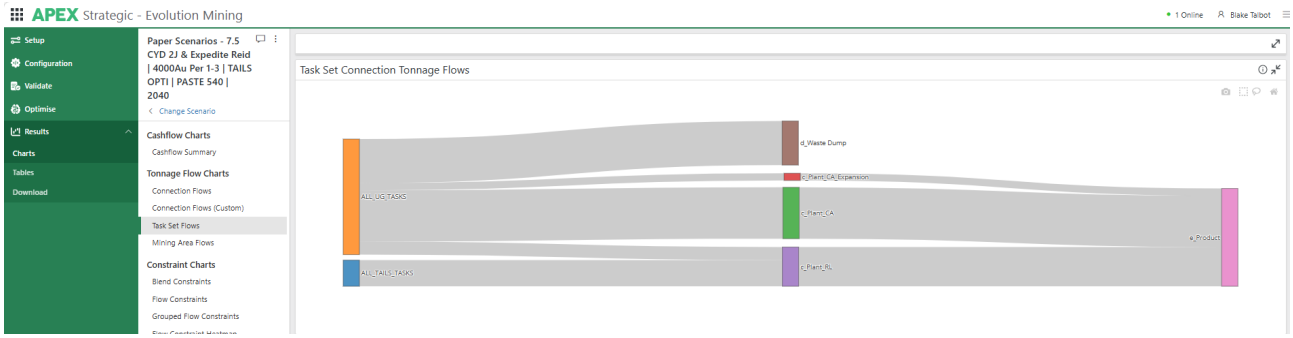


FIG 1 – Example APEX material flow network.

Operational constraints, including development capacity, material movement limits, backfill capacity and processing throughput, were represented within the optimisation together with capital triggers and operational decision options. While these constraints bounded the solution, they did not constitute a detailed resource-levelled schedule. Instead, the optimiser was constrained by annual capacity limits, making the outputs suitable for strategic planning and option evaluation rather than operational execution. This configuration enabled APEX to evaluate inventory selection, production sequencing, capital timing, and operational trade-offs within a single optimisation framework.

The vCOG feature

The vCOG feature extends the optimisation process by allowing cut-off grade to become a decision variable rather than a fixed input.

Traditionally, a single stope design is generated for each mining shape using a predetermined cut-off grade. The scheduler is then required to work within the inventory defined by that grade assumption. In contrast, vCOG allows multiple inventory options to be available to the optimiser, enabling it to determine which inventory selection creates the greatest value within the context of the overall schedule.

For LOM optimisation, a suite of nested Deswik®.SO stope shapes were generated for each stope location at cut-off grades ranging from 2 g/t Au to 8 g/t Au in 1 g/t increments, optimised for grade. Each shape represented a different inventory selection, with associated tonnes, grade, and economic value (see Figure 2).

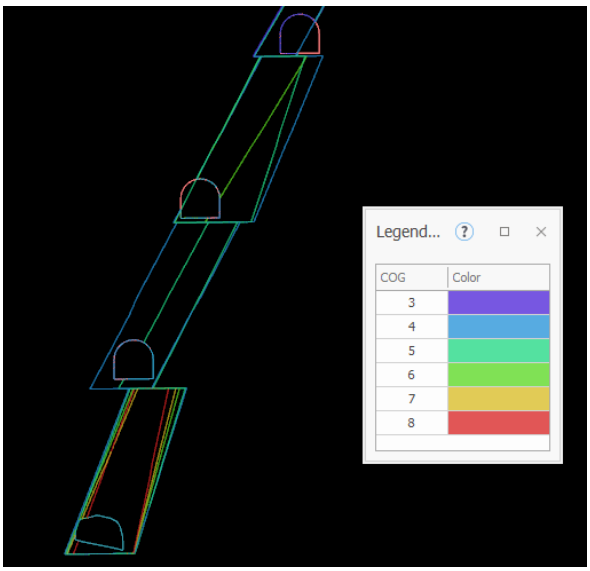


FIG 2 – Section view: SO's representing seven COG bins.

The shapes were nested, meaning that lower cut-off grades included all material captured by higher cut-off grades plus additional lower-grade material. As the cut-off grade increased, tonnes decreased and grade generally increased. This created a series of alternative inventory options for each stope location.

Within APEX, these alternatives were presented as mutually exclusive choices. The optimiser could select one cut-off grade option for a stope or elect not to mine the stope at all. Importantly, this decision was made simultaneously with all other scheduling decisions, including extraction sequence, development timing, processing allocation, and capital investment timing.

As a result, the selected cut-off grade reflected not only the economics of an individual stope, but also its contribution to value across the entire mining system. This approach aligns with the framework proposed by Hall (2014), in which cut-off grade is determined by the state of the system and its constraints rather than by a standalone economic calculation.

The vCOG feature can be applied at varying levels of resolution. At the most detailed level, each stope can be assigned an independent cut-off grade. Alternatively, common cut-off grades can be applied across larger groupings such as mining levels, orebodies, or planning domains.

For RLO, optimisation was performed at the stope level to maximise flexibility within the optimisation process. The resulting cut-off grade selections were subsequently analysed and aggregated by orebody to develop practical planning parameters suitable for shorter term planning execution.

Scenario design and constraint framework

The LOM Optimisation evaluated 63 scenario combinations designed to test the sensitivity of value to key operational and strategic assumptions (Table 1).

Nine base scenarios were developed by varying three primary factors:

1. Gold price assumptions.
2. Availability of an alternative surface ore source.
3. Paste backfill plant capacity.

The gold price scenarios considered both a constant long-term gold price and a hybrid pricing approach that incorporated elevated near-term prices to reflect prevailing market conditions. Surface ore source scenarios evaluated the inclusion or exclusion of an alternative processing feed, while paste plant capacity varied between the existing operating configuration and expanded capacity options.

TABLE 1

EVN's APEX optimisation scenario matrix.

Dev constraint options	Static Au Price | UG OPTI | PASTE Rate 2	Hybrid Au Price | UG OPTI | PASTE Rate 2	Hybrid Au Price | UG OPTI | PASTE Rate 1	Hybrid Au Price | UG OPTI | PASTE Rate 3	Hybrid Au Price | UG OPTI | PASTE Rate 2 | No RL Mill	Static Au Price | UG + Surface OPTI | PASTE Rate 2	Hybrid Au Price | UG + Surface | PASTE Rate 2	Hybrid Au Price | UG + Surface | PASTE Rate 1	Hybrid Au Price | UG + Surface | PASTE Rate 3
1. Base Scenario	1.1	2.1	3.1	4.1	5.1	6.1	7.1	8.1	9.1
2. Mining Zone 1: 1 Dev Resource	1.2	2.2	3.2	4.2	5.2	6.2	7.2	8.2	9.2
3. Mining Zone 1: 2 Dev Resource	1.3	2.3	3.3	4.3	5.3	6.3	7.3	8.3	9.3
4. Expedite Mining Zone 2: -1 year	1.4	2.4	3.4	4.4	5.4	6.4	7.4	8.4	9.4
5. Combine scenarios 3–4	1.5	2.5	3.5	4.5	5.5	6.5	7.5	8.5	9.5
6. Expedite Mining Zone 3: -1 year	1.6	2.6	3.6	4.6	5.6	6.6	7.6	8.6	9.6
7. Combine scenarios 1–6	1.7	2.7	3.7	4.7	5.7	6.7	7.7	8.7	9.7

TABLE 2

Optimisation scenario matrix relative value heatmap.

	Relative NPV (variance to mean)								
Dev constraint options	Static Au Price | UG OPTI | PASTE Rate 2	Hybrid Au Price | UG OPTI | PASTE Rate 2	Hybrid Au Price | UG OPTI | PASTE Rate 1	Hybrid Au Price | UG OPTI | PASTE Rate 3	Hybrid Au Price | UG OPTI | PASTE Rate 2 | No RL Mill	Static Au Price | UG + Surface OPTI | PASTE Rate 2	Hybrid Au Price | UG + Surface | PASTE Rate 2	Hybrid Au Price | UG + Surface | PASTE Rate 1	Hybrid Au Price | UG + Surface | PASTE Rate 3
1. Base Scenario	-3%	1%	-28%	6%	-26%	0%	2%	-13%	1%
2. Mining Zone 1: 1 Dev Resource	-1%	1%	-25%	6%		-2%	4%	-19%	10%
3. Mining Zone 1: 2 Dev Resource	-5%	3%	-24%	7%		4%	4%	-12%	3%
4. Expedite Mining Zone 2: -1 year	-1%	6%		12%		1%	6%	-10%	12%
5. Combine scenarios 3–4	-4%	7%		13%		7%	9%	-9%	7%
6. Expedite Mining Zone 3: -1 year	0%	3%		7%		0%	1%	-13%	10%
7. Combine scenarios 1–6	5%	8%		15%		4%	10%	-9%	10%

Across all scenarios, the optimisation was provided with a common processing infrastructure framework comprising existing operating capacity and a range of optional capital investment decisions. These options were available to the optimiser where economically justified and formed part of the overall decision space.

This enabled the optimisation to assess the interaction between inventory selection, scheduling constraints, processing capacity and capital deployment within a single value-based framework.

Across the nine base scenarios, a further seven development constraint options were evaluated to test the impact of alternative capital investment strategies. These scenarios included varying levels of development acceleration across key mining areas, additional development resourcing, targeted infrastructure investments, and an unconstrained case representing the combined implementation of all major debottlenecking initiatives. These scenarios were tested and are presented in order of ease of application as determined with the operation.

Each scenario (summarised in Table 1) represented a realistic and costed operating configuration, allowing value to be compared on a consistent basis.

The vCOG feature was active in all scenarios. Consequently, both the schedule and the inventory were allowed to respond to the unique conditions of each scenario. This differs from conventional planning approaches, where a fixed cut-off grade defines the inventory before scheduling begins, often on the basis of premediated capital investment decisions. Under the vCOG methodology, inventory selection became part of the optimisation process itself, allowing different combinations of constraints, capital investments, ore source availability, and economic assumptions to generate different optimal inventory outcomes.

This capability provided a direct mechanism for validating the inventory selections implied by the strategic optimisation study and translating those outcomes into practical planning parameters for the LOM plan.

Several clear and consistent trends emerged from the optimisation results and are presented in Table 2.

RESULTS

The primary purpose of the LOM optimisation was twofold: to validate the strategic direction established during the preceding Strategic Optimisation process and to generate practical planning guidance for implementation within the FY27 LOM plan. The optimisation results provided confidence in the relative value of key investment and development decisions while generating orebody-specific cut-off grade guidance. The outcome also provided the foundation for future planning processes that could allow for cut-off grade optimisation to occur as part of the standard approach to LOM planning.

NPV sensitivity and strategic validation

The scenario matrix (Table 1) provided a quantitative basis for evaluating the relative contribution of key operational constraints and capital investment opportunities. Results were expressed as percentage variance from the mean NPV across all solved scenarios, enabling direct comparison of the value generated by each strategic lever.

To enable like-for-like comparison between static and hybrid metal price scenarios, hybrid scenario outputs were normalised after optimisation. Revenue was recalculated using the static long-term gold price assumption while retaining the schedule, inventory selection and capital decisions generated under the hybrid price deck. The resulting value difference was then applied to the equivalent static price scenario, where all other parameters were held constant. This isolated the value associated with optimisation behaviour rather than the direct revenue uplift from higher near-term metal prices.

Paste backfill capacity

The optimisation demonstrated a high sensitivity to paste backfill capacity. Scenarios incorporating the highest paste plant throughput consistently outperformed the base configuration, delivering NPV improvements of up to 13 per cent.

This result reinforces the importance of the backfill cycle as a controlling constraint within the underground mining system at RLO. Increasing fill capacity reduced time lags between successive stoping activities, allowing more efficient utilisation of development access and production infrastructure. Across the optimisation suite, investment in backfill capacity was consistently associated with improved economic outcomes.

Development resourcing prioritisation

The optimisation framework also provided insight into the relative value of alternative development resourcing priorities.

Across multiple scenario configurations, different development acceleration opportunities exhibited materially different value outcomes. Under the hybrid pricing assumptions, the highest-performing development acceleration scenarios achieved NPV outcomes approximately 12–15 per cent above the scenario mean, compared with 6–10 per cent for alternative development acceleration strategies.

The consistency of these results across multiple optimisation runs demonstrated the ability of the framework to discriminate between competing capital allocation options and identify those that generated the greatest value under the prevailing constraint environment and conversely, the scenarios which delivered little value that could be excluded.

Alternative ore sources and capacity expansion

Scenarios incorporating an alternative ore source consistently generated additional value and provided a useful test of the optimisation framework's ability to evaluate competing material streams within a constrained production system.

The optimisation was provided with a range of capital investment decisions that could be selected where economically justified. These opportunities were subject to timing constraints and interacted with inventory selection, scheduling and processing capacity throughout the planning horizon.

Results demonstrated that ore source selection, cut-off behaviour and capital investment decisions are highly interdependent. Consequently, evaluating these decisions within a single optimisation framework provided greater insight than assessing each component independently.

Practical cut-off grade outputs by orebody

A key objective of the LOM optimisation was to translate detailed optimisation outputs into practical planning parameters that could be readily understood and applied by the operation.

While the optimisation selected inventory at the stope level, the practical implementation of such inventory selection decisions is impractical at RLO. To address this, optimisation outputs were analysed by orebody across the 26 underground block models and aggregated.

The optimisation does not directly output a cut-off grade. Instead, it determines which stopes are economically selected and when they are scheduled. Consequently, the distribution of mined stope grades (Figure 3) provides insight into the effective cut-off behaviour exhibited by the optimisation.

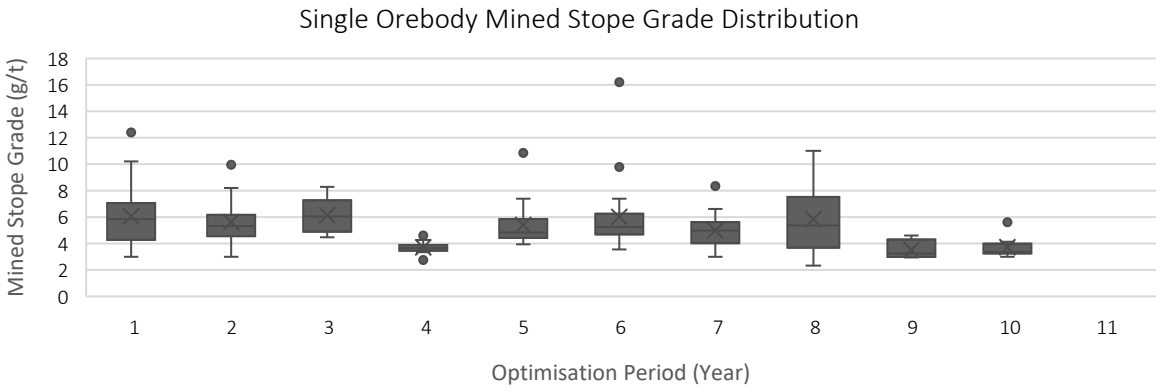


FIG 3 – Single orebody mined stope grade distribution.

For each orebody, the first quartile (P25) of mined stope grades scheduled throughout the life-of-mine was adopted as the Planning COG. Review of the optimisation outputs indicated that the P25 mined grade provided a practical approximation of the effective orebody cut-off, recognising that mined grades typically exceed their underlying economic cut-off grades. An Incremental (Marginal) COG was also derived to represent the lower bound of economic inclusion observed within the optimised schedule. This allowed the LTP team to assess lower grade production shapes within the context of the timing and constraints in the detailed mine schedule.

Material differences were observed between orebodies, reflecting variations in grade distribution, operating cost, mining geometry and position within the production system. For longer-duration orebodies, effective cut-off behaviour also varied through time as constraints, capacity utilisation and market conditions evolved. Consequently, the Planning COG should be viewed as a summary metric and may be supplemented by annual or phase-based analysis where greater precision is required. This also highlighted the value of implementing a cut-off grade optimisation process that could be replicated on a more routine basis than your standard cut-off grade optimisation studies.

Together, the Planning COG and Incremental COG provided a simple and consistent framework for communicating optimisation outcomes to operational planning teams and supporting economically informed decision-making.

Impact of gold price assumptions

To better understand the impact of vCOG on inventory selection under varying metal price environment, a comparison was drawn between two representative optimisation scenarios, shown in Figure 4.

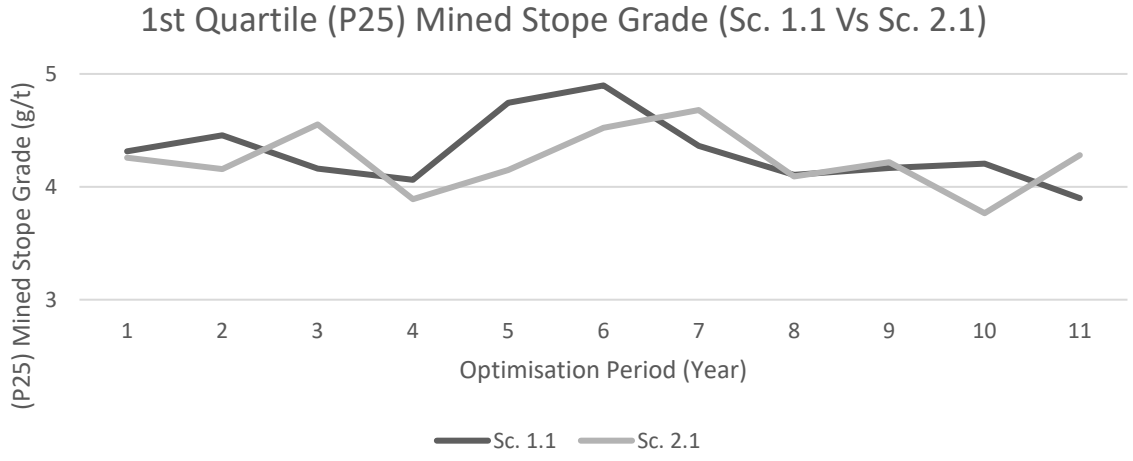


FIG 4 – Scenario comparison of mined stope grades over time.

In the early years of the static gold price scenario (Sc. 1.1), the optimised inventory exhibited a tighter and generally higher-grade distribution, reflecting the optimiser's preference for allocating constrained mining and processing capacity toward higher-margin material.

Under the elevated early metal price conditions (Sc. 1.2), the optimiser favoured the inclusion of additional lower grade economic material, improving utilisation of available processing capacity and increasing early cash flows. However, the operation's ability to respond to higher prices was ultimately constrained by the physical characteristics of the resource and the availability of development access. As a result, the benefit of elevated pricing was realised primarily through improved timing of production rather than a fundamentally different long-term inventory strategy.

While the magnitude and timing of these effects varied across optimisation scenarios, the example illustrates an important characteristic of the vCOG methodology: inventory selection can respond dynamically to changing pricing environments rather than being governed by a single cut-off grade applied uniformly across the project timeline.

The analysis supports the broader conclusion that the economically preferred cut-off grade is not a fixed value, but a planning parameter that varies according to the operating conditions and strategic objectives of the mining system.

LIMITATIONS AND CONSIDERATIONS

Input quality and deterministic optimisation

The optimisation is only as good as the inputs provided. It cannot compensate for bias or errors in the block model, stope designs, recovery assumptions or cost estimates. For this reason, optimisation results should always be considered alongside reconciliation performance, geological understanding and operational experience.

The optimisation was developed using a Mixed Integer Linear Programming (MILP) framework. While well suited to solving large-scale mine planning problems, MILP requires complex mining systems to be represented using simplified linear relationships. In reality, mining operations contain numerous non-linear behaviours and step changes associated with development access, infrastructure utilisation, equipment requirements and operational efficiency.

These behaviours can be represented in increasing levels of detail; however, doing so significantly increases the size and solve time of the optimisation problem. In practice, a balance must be struck between accurately representing the mining system, maintaining reasonable solve times, and evaluating a sufficient number of strategic scenarios. The framework adopted in this study focused on the major value drivers and constraint interactions, recognising that additional complexity would likely result in significantly longer runtimes for relatively small improvements in the optimisation outcome.

While conditional simulation results of key geological domains were reviewed as part of the broader planning process and provided confidence in the direction of the optimisation outcomes, the results presented in this paper represent a deterministic solution to an inherently uncertain problem.

Computational considerations

High-resolution optimisation runs for RLO typically required several days of cloud computation time and converged to an optimality gap of approximately 15 per cent, highlighting the practical trade-off between scenario breadth, solve time and solution quality.

This was managed through a staged workflow, with initial screening runs used to identify high-value scenarios before undertaking detailed optimisation. Improvements in computing capability and optimisation workflows are expected to reduce this constraint over time.

Development representation

The current implementation assumes a common development design across all cut-off grade alternatives for a given stope location. As a result, the optimiser does not fully capture potential

changes in development requirements associated with different stope geometries and cut-off grade selections. This is managed with post optimisation designs refinements by the engineer.

Future workflows may incorporate cut-off-specific development designs or development scaling factors to better represent the interaction between inventory selection and development requirements.

Translation to operational planning

The APEX output is a strategic schedule rather than a detailed resource-levelled operational plan. Additional planning is required to account for equipment availability, workforce constraints, and operational practicalities.

Accordingly, the cut-off grade outputs are best viewed as economically informed planning boundaries rather than prescriptive operating instructions.

CONCLUSION

The application of APEX and its vCOG capability at RLO demonstrated a practical approach for integrating cut-off grade optimisation into Life-of-Mine planning.

The Strategic Optimisation process established a preferred operating direction through the evaluation of alternative ore sources, capital investments, infrastructure configurations, and sequencing strategies. The subsequent LOM optimisation provided an independent mathematical validation of that direction within a detailed scheduling environment while generating the practical planning parameters required for implementation.

A key feature of the approach was that cut-off grade decisions were made within the optimisation itself rather than being prescribed as fixed planning assumptions. By presenting multiple inventory alternatives for each stope location, the optimiser was able to select the inventory that maximised value within the context of the broader mining system. This allowed cut-off grade selection to respond dynamically to schedule timing, operating constraints, capital decisions, and competing opportunities across the operation.

The optimisation confirmed that cut-off grade is not a fixed economic threshold, but a response to the prevailing constraints and opportunities within the mining system. While the optimisation operated at the stope level, subsequent analysis of the results identified distinct cut-off grade behaviours across different orebodies and mining domains, allowing practical planning ranges to be developed for operational use.

Beyond cut-off grade selection, the optimisation provided valuable insight into the relative importance of key operational and capital constraints. Paste backfill capacity, development sequencing, ore source strategy, and capital timing were all shown to have a material influence on project value, reinforcing the importance of evaluating these decisions within an integrated optimisation framework.

The orebody-level planning parameters derived from the optimisation provide a practical bridge between strategic planning and operational execution. Rather than relying on a single uniform cut-off grade across the operation, planners were provided with economically informed decision boundaries that better reflect the characteristics of individual ore systems and their role within the broader production schedule.

The work also highlights the importance of input quality in optimisation-based planning. As with any optimisation methodology, the quality of the output is fundamentally dependent on the quality of the underlying geological, mining, processing and economic assumptions. Continued improvements in geological confidence, reconciliation practices, stope design methodology, and cost estimation are expected to further enhance the quality of future optimisation outcomes.

More broadly, this work demonstrates that the practical barriers to implementing dynamic cut-off grade optimisation at operational scale have largely been overcome. Modern optimisation tools now allow inventory selection and cut-off grade determination to be evaluated simultaneously within the scheduling process, creating a stronger link between strategic intent, schedule optimisation, and Life-of-Mine planning outcomes.

For RLO, the vCOG methodology provided a practical mechanism for validating strategic planning decisions and translating them into an executable LOM plan. For the broader mining industry, it demonstrates a repeatable framework for treating cut-off grade as an output of the planning process rather than a fixed input, enabling more informed inventory decisions and better alignment between strategy and execution.

ACKNOWLEDGEMENTS

The authors wish to thank Evolution Mining for their support of the RLO strategic LOM optimisation programme and for approval to publish the findings of this work. Particular thanks are extended to the Group Technical and RLO site teams for their engagement with the planning process and their openness to challenging the *status quo* on mine planning methodology. The authors also acknowledge the contributions of the broader Deswik® team in advancing the vCOG feature within the APEX platform, and the AusIMM Strategic Mine Planning Conference organising committee for the opportunity to present this work.

REFERENCES

Evolution Mining Limited, 2025. Mineral Resource and Ore Reserve Statement Evolution Mining: Sydney.

Hall, B E, 2014. *Cut-off Grades and Optimising the Strategic Mine Plan*, Spectrum 20 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Lane, K F, 1988. *The Economic Definition of Ore: Cut-off Grades in Theory and Practice* (Mining Journal Books: London).

Mortimer, J J, 1950. Cut-off grade for different conditions of mine and mill operation, *Transactions of the Institution of Mining and Metallurgy*, 60:84–90.

Fleet, equipment and blasting in mine planning

Diesel hazards to zero-emission safety – techno-economic feasibility of mining fleet electrification at AngloGold Ashanti's Iduapriem gold mine, Ghana

A Acheampong[1,2], E O Ansah[3,4], L K Dadzie[5] and C Abuntori[6]

1. MPhil Fellow, University of Mines and Technology, Tarkwa, Ghana. Email: aacheampong2017@gmail.com
2. Superintendent Short- and Medium-Term Planning, AngloGold Ashanti Iduapriem Mine, Tarkwa, Ghana. Email: aacheampong3@aga.gold
3. Research Fellow, WH Bryan Mining Geology Research Centre, Sustainable Minerals Institute, The University of Queensland, Indooroopilly Qld 4068, Australia.
4. Research Fellow, School of Civil Engineering, Faculty of Engineering, Architecture and Information Technology, The University of Queensland, St Lucia Qld 4072, Australia.
5. PhD Fellow, Centre for Justice, Queensland University of Technology, Brisbane Qld 4000, Australia.
6. Senior Lecturer, University of Mines and Technology, Tarkwa, Ghana.

ABSTRACT

Diesel-powered haul trucks remain the dominant source of occupational health risks in gold mining, contributing high levels of diesel particulate matter (DPM), NOx, noise (up to 95 dB), heat, and vibration. These hazards drive lost-time injury rates, increased ventilation costs, and significant Scope 1 emissions. This study presents a techno-economic analysis of fleet electrification at AngloGold Ashanti's Iduapriem mine, evaluating the health, safety, operational, and financial implications of transitioning from a fully diesel fleet to electrified alternatives.

Using the mine's ten-year plan (2025–2035), three scenarios were compared: Base Case (current diesel fleet), Case 1 (diesel-electric trucks with trolley assist), and Case 2 (full battery-electric trucks). Cases 1 and 2 were assessed at 50 per cent and 100 per cent adoption levels. Key parameters included CAPEX, OPEX (capital and operational expenditure), fuel/electricity consumption and site-specific production data.

The results demonstrate substantial safety and health benefits. Case 2 (full electric) eliminates tailpipe emissions, reducing DPM and NOx exposure to near zero. Noise levels drop by 40–50 dB and heat generation by two-thirds, projected to reduce lost-time injuries by 40–50 per cent compared to the Base Case. Operationally, both Cases 1 and 2 improved cycle times by 10–30 per cent through higher torque and regenerative braking, increasing gold output by 5–15 per cent (75 000–225 000 additional ounces) and extending mine life by 2–5 years.

Economically, Case 2 yields the highest net present value (NPV) (US$650 million, +44 per cent versus Base), with payback in four years and cumulative OPEX savings of US$90 million. CO_2 emissions is predicted to fall by 70–90 per cent (200 000–720 000 t avoided). Sensitivity analysis confirms robustness across ±20 per cent variations in fuel prices and CAPEX.

Overall, battery-electric fleets are not only technically feasible but deliver superior health and safety outcomes while remaining economically attractive at Iduapriem. This study's lessons are transferable to other Ghanaian and African gold mines, supporting zero-emission mining operations.

INTRODUCTION

The global mining industry is currently navigating a critical transition from traditional diesel-powered haulage to electrified fleet systems. While diesel trucks offer operational reliability, their use results in significant environmental, health, and economic burdens. At the AngloGold Ashanti (AAIL) Iduapriem Mine in Ghana, reliance on conventional diesel machinery exposes workers to chronic health hazards, including high noise levels (95 dB), mechanical vibrations, and diesel particulate matter (DPM) exposure, while simultaneously driving high operational costs through volatile fuel prices.

This paper evaluates the techno-economic feasibility of electrifying the haulage fleet at the Iduapriem Mine. The research is driven by the urgent need to align mining operations with corporate net-zero commitments (30 per cent reduction in Scope 1 and 2 emissions by 2030) and national frameworks, such as Ghana's 2023 Green Minerals Policy (Ministry of Lands and Natural Resources, 2023). Using the mine's ten-year plan (2025–2035) (AngloGold Ashanti, 2024) and discounted cash flow modelling, three scenarios were compared: Base Case (current diesel fleet), Case 1 (diesel-electric trucks with trolley assist), and Case 2 (full battery-electric trucks). By modelling the transition from diesel to trolley-assisted (Case 1) and battery-electric (Case 2) systems, this study provides a decision-support framework for sustainable open pit mining in West Africa.

METHODOLOGY

The methodology employs a deterministic techno-economic model based on the Iduapriem ten-year mine plan. The model integrates production data, including total volume mined, stripping ratios (up to 45), and gold output, with financial parameters such as fuel/electricity pricing and maintenance costs. Three distinct scenarios were modelled:

1. **Base Case:** Continued use of current diesel-powered haulage (mechanical and diesel-electric trucks).
2. **Case 1:** Retrofitting diesel-electric trucks with trolley-assist systems for steep uphill segments.
3. **Case 2:** A full transition to battery-electric vehicles (BEVs) featuring fast-charging/swapping capabilities.

Safety impacts were assessed by correlating the reduction of DPM, noise, and vibration with Lost-Time Injury (LTI) data from the Iduapriem site and regional industry benchmarks.

RESULTS AND DISCUSSION

Figure 1 shows the ten-year plan data used for modelling the three scenarios: Base Case (current diesel fleet), Case 1 (diesel-electric trucks with trolley assist), and Case 2 (full battery-electric trucks).

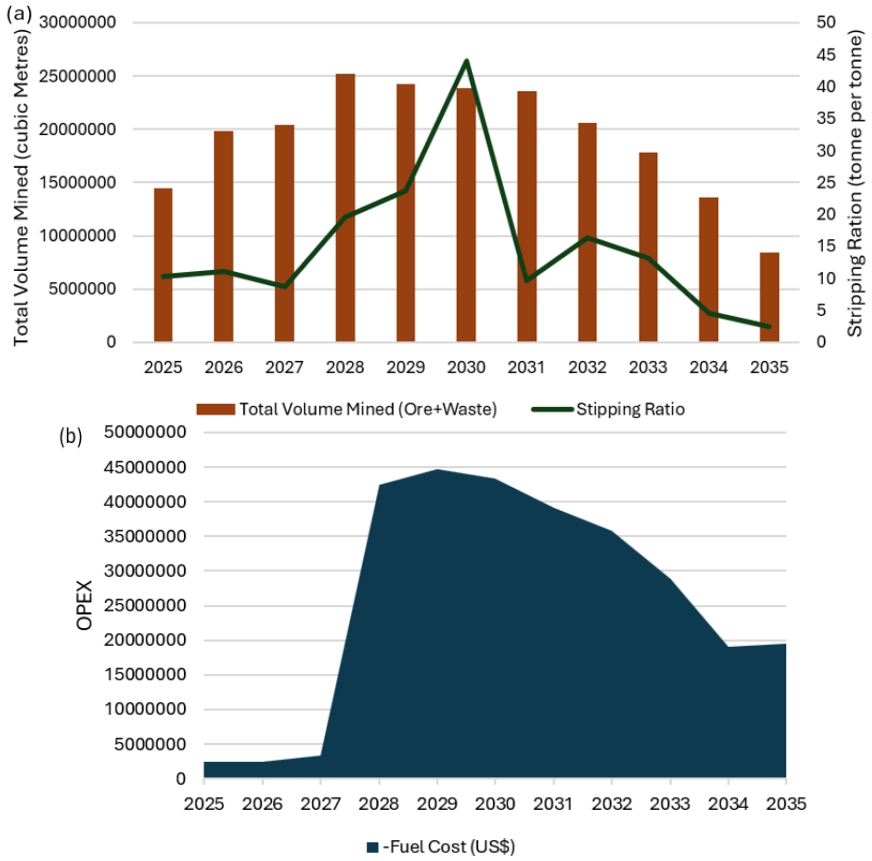


FIG 1 – (a) Total ore mined against stripping ratio; and (b) fuel cost over ten years.

Techno-economic implications

The financial analysis reveals that electrification is highly viable (Table 1). The Base Case, while requiring no immediate capital investment, remains vulnerable to fuel price volatility. In contrast, Case 2 (100 per cent battery-electric) requires a higher initial capital expenditure (CAPEX) of about US$100 million but yields the highest long-term net present value (NPV) of US$650 million.

TABLE 1

Summary of techno-economic performance (2025–2035).

Metric	Base case	Case 1 (100%)	Case 2 (100%)
Total CAPEX (US$M)	0	25	100
Cumulative OPEX savings	0	240	480
NPV (US$M)	450	600	650
Payback (years)	N/A	3–4	4

Safety and health impacts

The transition to Case 2 offers transformative health benefits. The elimination of diesel tailpipe emissions directly addresses respiratory risks for site personnel. Furthermore, the significant reduction in heat generation. Specifically, electric drivetrains producing only one-third of the heat of diesel counterparts, which mitigates heat stress in the tropical Ghanaian climate. Also, there is predicted 40–50 dB reduction in operational noise and the decrease in vibration levels are expected to contribute to a 40–50 per cent reduction in LTIs.

Operational efficiency and sustainability

Electrification also unlocks operational efficiencies. The higher torque of electric vehicles (Case 2) results in 20–30 per cent faster acceleration, which, combined with regenerative braking, improves cycle times by 10–30 per cent. This efficiency allows for a throughput increase, potentially adding 150 000–225 000 ounces of gold over the life-of-mine. Environmentally, Case 2 achieves a predictable 70–90 per cent reduction in CO_2 emissions compared to the diesel baseline.

CONCLUSIONS

This study demonstrates that the electrification of haulage fleets at the Iduapriem Mine is technically and economically superior to the *status quo*. While Case 1 (trolley-assist) offers a lower CAPEX entry point for partial decarbonisation, Case 2 (100 per cent battery-electric) provides the most robust path toward net-zero targets and long-term economic stability. For Ghanaian mining operations, this transition is supported by favourable policy frameworks such as the Green Minerals Policy and duty waivers on electric vehicle components. Future research should focus on the integration of autonomy and the long-term monitoring of battery performance under Ghana's tropical conditions to refine these findings.

ACKNOWLEDGEMENTS

The author wishes to acknowledge AngloGold Ashanti Iduapriem Limited for providing access to operational data and the Department of Mining Engineering at the University of Mines and Technology, Tarkwa, for their support.

REFERENCES

AngloGold Ashanti, 2024. Business plan 2025, Iduapriem Mine, internal company document.

Ministry of Lands and Natural Resources (MLNR), 2023. Green Minerals Policy, Government of Ghana.

Incorporating maintenance-driven equipment costs into mine planning

D Kaykov[1] and M Kumral[2]

1. PhD Candidate, Department of Mining and Materials Engineering, McGill University, Montreal QC H3A 0E8, Canada. Email: dimitar.kaykov@mail.mcgill.ca
2. Professor, Department of Mining and Materials Engineering, McGill University, Montreal QC H3A 0E8, Canada. Email: mustafa.kumral@mcgill.ca

ABSTRACT

The purpose of this paper is to show how different levels of sustaining capital costs, as well as different maintenance costs due to alternative maintenance scenarios can impact the results of strategic mine planning. With regard to maintenance costs, two alternative strategies for performing maintenance procedures were assumed, involving a different ratio between corrective and preventive maintenance operations. Sustaining capital costs and maintenance costs per annum are based on the piecewise decaying availability of the equipment according to the specified Generalised Renewal Process (GRP) model. Such formulation allows to simulate equipment failures and the associated maintenance procedures on an extended period of several years of life-of-mine operation, rather than assuming constant and time invariant equipment-related costs.

As a result, the selection of maintenance strategy should allow not only technical factors to be considered, but also economic, financial and mine-planning implications of maintenance activities. For example, rebuilding or overhaul of the equipment will lead to a sharp improvement in its reliability, while also temporarily reducing maintenance costs for sequential periods by restoring part of the equipment condition. But with ageing equipment and accumulated degradation, probability of failure increases and further maintenance activities are required. Thus, this model makes it possible to take into account the interrelations between the current equipment condition, its maintenance strategy, cost accumulation and production scheduling.

Finally, to show the potential implications of the suggested approach for strategic mine planning, a case study analysis is conducted on the benchmark Marvin deposit. Specifically, this research explores how various maintenance-related costs, involving both preventive and corrective activities, will influence strategically important mine planning results such as pit limit determination, mine scheduling, cash flow formation and the ultimate value of the project. Consequently, this paper aims to provide a practical tool to mine planners by presenting a model for estimating equipment-related costs in more realistic terms.

INTRODUCTION

Mobile equipment is a critical asset for mining operations. Meeting production targets specified in mine plans largely depends on the availability of individual pieces of equipment and the whole system. Due to ageing, wear, tear, corrosion and harsh operating environments maintaining above a constant level of equipment availability generally requires increasing maintenance, rebuilds and sustaining capital costs over time. In current mine planning, availability is usually assumed constant. This treatment could be problematic because operations can employ a wide range of equipment fleet configurations and maintenance strategies to achieve this target but realising different equipment reliability during the life-of-mine. Hence, this non-unique input can result in significantly different annual costs due to varying shares of equipment capital, sustaining and maintenance costs, all of which affect the project's net present value (NPV).

In addition to that, traditional strategic mine planning models use constant unit costs for their depiction of mining costs (in USD/t). This is a suitable representation since the costs in question are directly related to the volume of materials handled. However, this approach may not be ideal for costs that are independent of or only weakly related to the amount of mined ore mass. Costs related to the choice of maintenance strategy, which are often neglected, are explicitly considered in the proposed model. These costs are generated by equipment degradation, stochastic failure behaviour, planned and unplanned maintenance interventions and the timing of rebuilding and replacement, rather than by tonnes mined alone.

The recent literature in open pit mine planning has continued to formulate pit optimisation and mine scheduling as multi-period optimisation problems, where the extraction sequence is determined to maximise project value subject to constraints such as destination policies and operational requirements (Newman *et al,* 2010; Khan, 2024; Jélvez *et al,* 2023). Contemporary literature has expanded the scope of the mine scheduling problem to include waste dumping and in-pit waste or tailings disposal (Rimélé, Gamache and Doyon-Poulin, 2018; Maremi and Ben-Awuah, 2016), stockpiling (Levinson and Dimitrakopoulos, 2020; Montiel and Dimitrakopoulos, 2013), processing streams and products (Goodfellow and Dimitrakopoulos, 2017; Dimitrakopoulos and Lamghari, 2022) and sustainability-oriented aspects such as greenhouse-gas-aware scheduling and environmental costs integration (Rahnema *et al,* 2024; Mirzehi and Moradi Afrapoli, 2024; Lotsu, Bimpong and Boakye, 2026). At the same time, strategic mine planning research has increasingly adopted data-driven, artificial-intelligence and learning-based approaches to improve optimisation, uncertainty handling and decision support (Noriega and Pourrahimian, 2022; Quelopana and Navarra, 2021; Yaakoubi and Dimitrakopoulos, 2022).

On the other hand, mining-equipment performance studies show a direct relationship between reliability, availability, maintainability and maintenance planning (Barabady and Kumar, 2008; Moniri-Morad, Pourgol-Mohammad and Sattarvand, 2014). Recent maintenance literature has also emphasised on preventive, reliability-based and condition-based maintenance strategies for mining equipment and other industrial assets (Robatto Simard, Gamache and Doyon-Poulin, 2023; Sarıgül, 2022; Javadnejad *et al,* 2022; Sharma, Mittal and Soni, 2022). However, despite the clear link between strategic mine planning and equipment maintenance management through equipment availability, production capacity and maintenance costs, limited attention has been given to explicitly incorporate equipment deterioration, maintenance strategy selection and time-dependent maintenance costs into pit limit definition and long-term mine scheduling models.

Therefore, this paper aims to address the existing methodological gap by further extending the links between the concept of reliability-driven maintenance-cost modelling to that of strategic mine planning. In conventional mine planning, costs are usually represented in terms of fixed unit costs (USD/t). However, in maintenance and reliability problems, equipment performance is usually considered independent of its production rate. This logic can be extended to pit optimisation and mine scheduling. Hence, the current study attempts to link these two areas by obtaining time-dependent maintenance and sustaining capital costs based on the performance of a typically used haul truck, extrapolates their reliability on a fleet level and assess the effect of the selection of a maintenance strategy for pit limits and mine scheduling.

PROPOSED RELIABILITY-DRIVEN MAINTENANCE COST MODEL

It is well known that two machines with the same chronological age may have different likelihoods of failure when their operating and maintenance histories are distinct. In this approach, the dynamics of the equipment state are represented by a Generalised Renewal Process (GRP) based on virtual age, where virtual age acts as the state variable describing the accumulated degradation condition of the equipment after each failure or maintenance event (Kijima and Sumita, 1986; Kijima, 1989). The use of GRP and virtual-age-based models is well established in reliability literature, particularly for repairable systems and imperfect maintenance modelling. These models have been applied to represent failure recurrence, distinguish between minimal, perfect and imperfect repair effects, assess preventive-maintenance policies and estimate maintenance-related life cycle costs (Doyen and Gaudoin, 2004; Wu, Niknam and Kobza, 2015).

Weibull distribution as choice of underlying failure time distribution

The Weibull distribution, as commonly used one in reliability engineering, was selected because it can flexibly represent increasing failure rates associated with wear-out behaviour in ageing mining equipment. For example, a new machine begins with a virtual age of $V(0) = 0$. As the machine operates, its virtual age increases. This implies that the machinery becomes older, but may be less or more vulnerable to breakdowns despite its chronological age. The implication of this is that reliability can be modelled as a function of virtual age and can be described using the two-parameter Weibull distribution as follows:

$$R(V(t)) = exp\left[-\left(\frac{V(t)}{\eta}\right)^{\beta}\right] \quad (1)$$

Where:

V is the virtual age of the equipment at time t, engine hours

β shape parameter of the baseline Weibull failure-time distribution of the GRP

η scale parameter of the baseline Weibull failure-time distribution of the GRP, engine hours

In this approach, the degradation of the system is characterised using virtual age instead of time. The virtual age can be seen as a state variable that accumulates during use but is partially depleted upon performing any kind of maintenance. In this regard, reliability is determined using a Weibull distribution with two parameters. When the Weibull shape parameter β is greater than 1, the failure rate increases as virtual age increases. The wear out phenomenon is explained in the model, which allows to mimic how aged or deteriorated units become increasingly prone to failure from the influence of unspecified factors such as fatigue, abrasion, contamination, thermal effects, parts wear and others.

Virtual age decomposition

An important feature of the model is the decomposition of virtual age into two categories, one being permanent and the other being repairable. This is important because it allows for the inclusion of the fact that maintenance may remove some of the ageing but not all of it. The part that cannot be removed by maintenance might include issues such as fatigue, frame wear, thermal effects, among others, but can be mitigated to some extent by rebuilds/overhauls.

$$V(t) = V_{perm}(t) + V_{rep}(t) \quad (2)$$

Where:

$V_{perm}(t)$ irreversible virtual age due to degradation in time t

$V_{rep}(t)$ repairable virtual age component

Realistically, $V_{rep}(t)$ can be reduced through preventive maintenance (PM), corrective maintenance (CM) or rebuild activities (R). During equipment operation, virtual age accumulates according to:

$$dV = (1 + \gamma)dt \quad (3)$$

The increment is divided between permanent and repairable ageing as follows:

$$dV_{perm} = \alpha\, dV \quad (4)$$

$$dV_{rep} = (1 - \alpha)dV \quad (5)$$

Where:

γ ageing acceleration factor, $\gamma \in R$ (assumed that $\gamma > 0$)

α irreversible ageing factor, $\alpha \in [0, 1]$

The parameter γ is defined to represent increased equipment wear under more severe operating conditions compared to a reference one, while α controls the proportion of ageing that becomes irreversible and damage that remains accumulated until a rebuild or overhaul is performed. This structure is important because it prevents the model from assuming that maintenance can make old equipment fully equivalent to new equipment, ie even after strong maintenance actions or rebuilds, accumulated permanent age remains. As the parameter α serves modelling purposes and does not have an explicit physical meaning, its value should be based on the best consideration and expertise of the maintenance planner.

Imperfect maintenance

The model also implements an imperfect-maintenance logic based on the Kijima virtual-age framework. As discussed, maintenance does not necessarily restore the equipment to an as-good-

as-new condition. Instead, it partially reduces the repairable component of virtual age. For the assumed categories of maintenance activities, the model applies a Kijima Type I rule to the repairable component of virtual age (Kijima, 1989):

$$V_{rep}{}^{+}(t) = V_{rep}{}^{-}(t) + \rho_{repair} \times W \quad (6)$$

Where:

$V_{rep}{}^{+}(t)$	repairable virtual age in time t after repair, engine hours
$V_{rep}{}^{-}(t)$	repairable virtual age in time t before repair, engine hours
W	working time interval between two sequential repairs, engine hours
ρ_{repair}	Kijima Type I repair effectiveness, $\rho_{repair} \in [0, 1]$

A smaller value closer to 0 for ρ_{repair} indicates stronger rejuvenation, whereas a value close to 1 indicates weak maintenance effectiveness.

For rebuilds, the model applies a stronger rejuvenation rule to the repairable component using the Kijima Type II rule (Kijima, 1989). Additionally, two separate parameters are introduced for repairable and permanent accumulated virtual age:

$$V^{+}(t) = \rho_{perm|Rebuild} \times V_{perm}{}^{-}(t) + \rho_{rep|Rebuild} \times V_{rep}{}^{-}(t) \quad (7)$$

Where:

$\rho_{perm\|Rebuild}$	Kijima Type II repair effectiveness of rebuilds/overhauls related to the permanent virtual age component, $\rho_{perm\|Rebuild} \in [0, 1]$
$\rho_{rep\|Rebuild}$	Kijima Type II repair effectiveness of rebuilds/overhauls related to the repairable virtual age component, $\rho_{rep\|Rebuild} \in [0, 1]$
$V^{+}(t)$	virtual age in time t after rebuild, engine hours
$V_{perm}{}^{-}(t)$	permanent virtual age in time t before rebuild/overhaul, engine hours
$V_{rep}{}^{-}(t)$	repairable virtual age in time t before rebuild/overhaul, engine hours

Therefore, rebuilds reduce repairable degradation substantially, but realistically, they might not always significantly decrease the irreversible damage of the asset. This provides the possibility for the model to reproduce a more flexible configuration of diminishing restorative impact from rebuilds as equipment ages.

Generation of stochastic failure events on engine-hour timeline

Corrective maintenance events arise endogenously from the reliability process. At the beginning of each operating period or after a maintenance event, the model evaluates the current value of total virtual age using a time increment (dt). During operation, virtual age increases and the reliability function decreases, reflecting the progressive degradation of the equipment. At each increment, a random uniformly distributed random variable (u) on the interval [0, 1] is drawn. This aims to mimic the mechanism when during operation, virtual age increases and the reliability decreases. A corrective-maintenance event is triggered when $R\big(V(t)\big) \leq u$. This procedure is equivalent to sampling the residual time to failure conditional on the current degradation state. The rationale behind this implementation is that it models the real-life behaviour of an aged system, namely that failure events can become more recurrent, thereby increasing corrective maintenance costs and reducing the system's remaining useful life compared with a new system.

Simulation of stochastic maintenance events on calendar time

The simulation module of the model combines both a deterministic (planned maintenance activities) and stochastic component (corrective maintenance activities). Planned maintenance is performed at fixed engine-hour intervals, while corrective maintenance is triggered by virtual-age-dependent failure occurrences. In both cases, the duration of the maintenance time is simulated using a Weibull

distribution for the maintainability of the equipment based on the category type of the maintenance job. For this case study, five categories of maintenance activities were assumed and two rebuild/ overhaul activities, as shown in Table 1. Additionally, the Weibull distribution parameters were derived from the assumed mean and coefficient of variation (COV) for the maintenance duration.

TABLE 1

Maintenance activity categories, repair effectiveness assumptions and downtime duration distributions used in the equipment degradation model.

Maintenance category	Role in the model	Assumed mean duration (h)	Assumed COV
Service	Provides condition monitoring, lubrication, adjustment and minor servicing, with negligible or very short downtime	2	0.30
Minor preventive maintenance	Provides limited rejuvenation and short downtime	6	0.35
Moderate preventive maintenance	Provides grater rejuvenation than minor maintenance	12	0.40
Major preventive maintenance	Provides substantial rejuvenation and restores a larger portion of the equipment condition, but requires longer downtime	36	0.45
Corrective maintenance	Triggered by stochastic failure through the GRP	24	0.80
Minor rebuild/ overhaul	Provides minor rejuvenation of repairable and partially permanent degradation	120	0.50
Major rebuild/ overhaul	Provides major rejuvenation of repairable and partially permanent degradation	240	0.40

After maintenance events are generated in the engine-hour domain using the GRP, they are mapped into calendar time by assigning a stochastic duration to each maintenance activity. The occurrence time of each subsequent event is then adjusted to account for equipment downtime. If the duration of a maintenance activity overlaps with a later event, the later event is shifted to the earliest possible time after the preceding maintenance activity has been completed.

As a result, the simulated calendar timeline alternates between two general equipment states: available and unavailable. This makes it possible to estimate the availability of a single equipment unit over time. Hence, in the scope of the applied model, availability comes as a result of a combination between the age dependent behaviour of the reliability function, maintenance frequency, the effectiveness of repairs and the stochastic duration of maintenance.

Maintenance strategy annual costs estimation

The cost model integrates maintenance costs (related to the corrective and planned preventive maintenance activities) and sustaining capital costs (related to the rebuild/overhaul activities). The total annual cost is the sum of all relevant cost categories. The model is able to take into account not only the cost of maintenance but also the broader economic implications of equipment reliability. As both failure occurrences and maintenance durations are stochastic, the costs are evaluated using Monte Carlo simulation. Each simulation run uses two independent random seeds: 1) for the generation of events using GRP; and 2) for maintenance-duration generation. The model then produces annual series for maintenance costs and capital sustaining costs in calendar time domain, which can be used as inputs for mine scheduling. The base formula used for estimating the annual maintenance costs for each maintenance activity category is as follows:

$$C_{m_i} = N_{m_i} . c_m \quad (8)$$

Where:

m	activity from the assumed set of maintenance and rebuild/overhaul activities
c_m	mean costs of maintenance and rebuild/overhaul activities, USD
N_{m_i}	simulated number of maintenance and rebuild/overhaul activities of type m for year i
C_{m_i}	annual costs for activity of type m for year i, USD

MAINTENANCE STRATEGY COSTS IMPACT ON PIT OPTIMISATION AND MINE SCHEDULING

Incorporation of annual maintenance costs and capital sustaining costs into the mine planning optimisation framework

Conventionally, mining costs expressed in USD/t assume that each additional tonne mined generates the same incremental cost and that this cost can be assigned consistently across all blocks and scheduling periods. Maintenance-strategy costs do not necessarily satisfy this assumption. The same quantity of material may be mined under different maintenance strategies. However, each strategy can result in different annual maintenance expenditures, downtime losses, rebuild timing and, in some cases, the need for equipment replacement. When the above expenses are aggregated into a uniform block economic value, the model will lack the ability to distinguish between the effect of selected maintenance strategy from the timing of the costs related to them. As a result, time-dependent costs are incorrectly treated as a constant block-level parameter.

This distinction does not imply that all equipment-related costs should be removed from unit mining costs. Variable costs related to moving materials and the use of machinery, like fuel, tire usage, oil and operating expenses based on use, can still be expressed in USD/t. However, costs associated with sustaining capital and replacement decisions are better represented as annual or period-specific costs. This is due to the extent to which such an increase will happen will depend not only on how much ore has been mined, but also on how long the equipment has been used, its long-term accumulated degradation and condition.

Alternatively, embedding maintenance-strategy costs in a USD/t mining cost would require annual or life-of-mine maintenance expenditures to be divided by the tonnes mined. However, when assessing the influence of maintenance strategy on strategic pit planning, this approach can introduce an averaging bias. In traditional pit optimisation procedures, where the maintenance cost is stated in terms of USD/t, this cost element becomes part of the production cost and economic evaluation of blocks, meaning that the value of time-based maintenance costs is averaged out against the tonnage across all periods. This can mask the influence of expenses, which are due to equipment ageing and more frequent failures. Consequently, annual maintenance costs were distributed over the planned tonnes to be mined in the corresponding year to derive period-specific maintenance cost-per-tonne values. The average value across all periods was used for pit optimisation, while the actual annual cost-per-tonne values were applied during scheduling. Sustaining capital costs were also introduced at the scheduling stage as period-specific costs. This ensures that the pit optimisation stage remains consistent with the traditional cost per tonne paradigm, whereas the scheduling stage will have a more direct emphasis on time-based maintenance cost implications.

It is necessary to mention, though, that for later stages of development and particularly at the end of life of a certain operation, such calculations can be less reliable because they depend on planned movement of material, condition of the equipment and maintenance activities. However, the impact of these inaccuracies on project NPV is to some extend mitigated by discounting, since costs occurring in the final years have a lower present value. As a result, potential errors in late-period maintenance-cost estimates may have a limited effect on the overall economic comparison, although they may still be important from an operational reliability and scheduling perspective.

Application in mine planning software

The distinction between direct costs assigned to mined blocks and costs arising due to time-dependent performance of the equipment can be incorporated within the context of GEOVIA Whittle's model. The software version used for this case study was GEOVIA Whittle 2022 (Dassault Systèmes, 2022). Time-varying costs related to equipment, such as dynamic costs for maintenance are included in the calculations by summing them with the assumed constant mining costs and added to the unit costs for mining. Double-counting was avoided by differentiating between time-varying unit costs and absolute-value equipment capital costs. In this way, the model can reflect changes in the economic burden of maintenance over time. Under this framework, GEOVIA Whittle can evaluate not only the quantity and value of material mined, but also the timing impact of equipment-related cash flows. A schedule may be penalised if material extraction is delayed into periods of higher equipment degradation and greater maintenance burden or periods of higher costs related to major fleet overhauls. Conversely, the mining schedule, cash-flow timing, project NPV may be affected by the chosen maintenance strategy. Additionally, a maintenance strategy which is generally more expensive may also influence the selection of the ultimate pit shell. This ensures the proposed framework to preserve the causal structure of the equipment cost model while reflecting the economic consequences of different maintenance policies in strategic mine planning outcomes. At the same time, the approach proves to be practical and compatible with the existing industry-adopted methods of mine planning, making it usable for optimisation and scheduling purposes.

CASE STUDY

Inputs and assumptions

The proposed framework is evaluated through a case study in which two alternative maintenance strategies are incorporated as two separate scenarios into the strategic mine planning model. The aim of the study is to examine how maintenance-related costs affect decisions related to strategic mine planning based on the selection of pit limits, mine scheduling outcomes and project value. For this purpose, the open-source Marvin deposit block model was used, which is frequently assumed as a benchmark for strategic mine planning experiments and proof-of-concept model development (MiningMath, 2026). Maintenance activities that are incorporated into the study and their frequency of repetition according to the different maintenance policies are shown in Table 2.

TABLE 2

Maintenance activity periodicities and unit cost assumptions for the two maintenance strategies.

Maintenance category	Maintenance Strategy 1 periodicity (engine hours)	Maintenance Strategy 2 periodicity (engine hours)	Mean activity cost (USD)
Service	200	500	773.41
Minor preventive maintenance	500	1000	3374.89
Moderate preventive maintenance	2500	3000	14 062.04
Major preventive maintenance	5000	7500	50 623.33
Corrective maintenance	GRP	GRP	42 186.11
Minor rebuild/overhaul	10 000	15 000	210 930.55
Major rebuild/overhaul	30 000	35 000	421 861.10

Table 3 shows the assumed repair effectiveness values for the elements of the considered set of maintenance and rebuild/overhaul activities. The selected repair-effectiveness ranges reflect the assumed increasing restorative capacity of maintenance interventions, from negligible rejuvenation during service and inspections to stronger virtual-age reduction during major preventive

maintenance and rebuilds, while still preventing the model from unrealistically restoring aged equipment to an as-good-as-new condition.

TABLE 3

Maintenance activities repair effectiveness assumptions for the two maintenance strategies.

Maintenance category	Maintenance Strategy 1 repair effectiveness	Maintenance Strategy 2 repair effectiveness
Service	$\rho_{repair\|PM_{Service}} = 1.00$	$\rho_{repair\|PM_{Service}} = 1.00$
Minor preventive maintenance	$\rho_{repair\|PM_{Minor}} = 0.75$	$\rho_{repair\|PM_{Minor}} = 0.70$
Moderate preventive maintenance	$\rho_{repair\|PM_{Moderate}} = 0.60$	$\rho_{repair\|PM_{Moderate}} = 0.55$
Major preventive maintenance	$\rho_{repair\|PM_{Major}} = 0.45$	$\rho_{repair\|PM_{Major}} = 0.40$
Corrective maintenance	$\rho_{repair\|CM} = 0.65$	$\rho_{repair\|CM} = 0.55$
Minor rebuild/overhaul	$\rho_{rep\|R_{Minor}} = 0.45$ $\rho_{perm\|R_{Minor}} = 0.95$	$\rho_{rep\|R_{Minor}} = 0.35$ $\rho_{perm\|R_{Minor}} = 0.93$
Major rebuild/overhaul	$\rho_{rep\|R_{Major}} = 0.25$ $\rho_{perm\|R_{Major}} = 0.87$	$\rho_{rep\|R_{Major}} = 0.20$ $\rho_{perm\|R_{Major}} = 0.82$

The GRP model parameters were selected to represent the degradation behaviour of a 150 t mining haul truck. The assumed parameter values used in the simulations are presented in Table 4. The ageing acceleration factor (γ=0.30) represents a moderate increase in ageing with accumulated operating hours (for every hour of operating time, the equipment ages 30 per cent faster), while the irreversible ageing factor (α=0.15) accounts for the portion of equipment ageing that remains after maintenance (15 per cent) and cannot be recovered through repair actions.

TABLE 4

Assumed GRP parameter values for an average 150 t mining haul truck.

Generalised renewal process (GRP) parameter	Value
GRP Weibull shape (β)	4.25
GRP Weibull scale (η)	15 000
Ageing acceleration factor (γ)	0.30
Irreversible ageing factor (α)	0.15

A 24-haul-truck fleet was assumed to be required for the deposit. During the first seven years of the life-of-mine, the operation uses 12 haul trucks to cover an annual material production rate of 25 Mt. From Year 8 onward, an additional 12 haul trucks are introduced, increasing the total fleet size to 24 haul trucks for the remaining mine life and the material throughput to 40Mt. Under this assumption, each individual truck will be considered separately with the use of the GRP model, where similar reliability parameters, repair effectiveness parameters and maintenance cost will apply to all trucks. Both maintenance policies were considered appropriate because the average truck availability was always higher than 0.85 for any mining period. Under the assumption that haul trucks are identical and statistically independent, the expected fleet-level annual maintenance cost is additive and the expected annual cost for one representative truck was multiplied by the number of active trucks in each mining period to obtain the total annual fleet maintenance cost. The total number of Monte Carlo simulation runs was 100 for developing the virtual age and reliability profiles, as well as the

costs per annum dynamics using the suggested maintenance strategy schedules. The use of such a method allows depicting the annual cost curve as an expected one and not just a single scenario.

Results

The first step in the evaluation of the simulated maintenance strategies was the consideration of their influence on the condition and reliability of the equipment. As illustrated by Figure 1, virtual age and reliability change differently depending on the assumed values of the maintenance intervals in each of the two strategies.

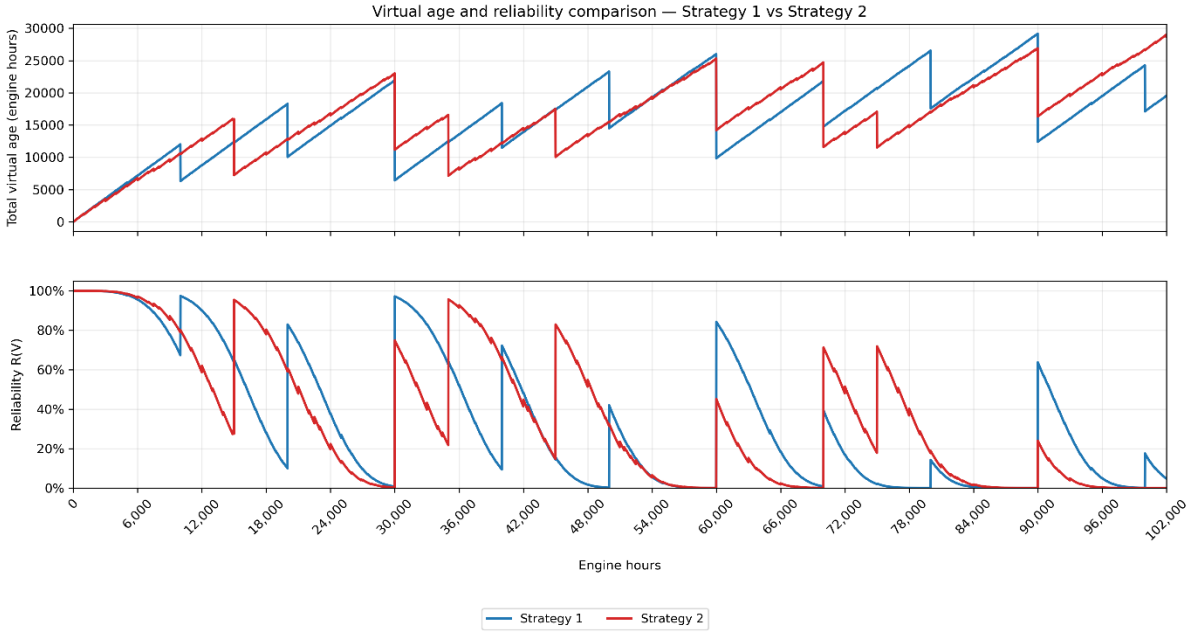


FIG 1 – Comparison of virtual age and reliability evolution for a single haul truck under the two maintenance strategies.

Under Strategy 1, a slower decreasing reliability is observed along with increased stability of the reliability level. In turn, Strategy 2 leads to periods of higher equipment degradation and a more pronounced decline in reliability. The effect that these differences have in terms of economics is illustrated by Figure 2 showing maintenance and sustaining capital costs for the fleet of mining trucks.

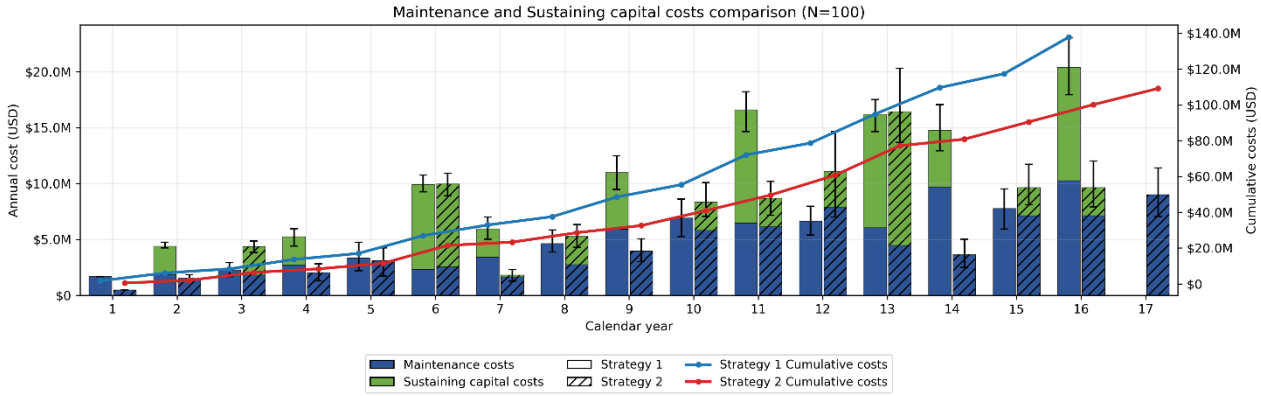


FIG 2 – Comparison of annual and accumulated maintenance and capital sustaining costs related to haul truck fleet under the two maintenance strategies.

Figure 2 shows that maintenance and sustaining capital costs do not follow a uniform annual pattern. Instead, both strategies exhibit an increasing cost structure with irregular peaks across the life-of-mine. The trends shown by the stacked cost bars show that the cost of maintenance increases over

time, where the discrepancies between the two strategies become even more distinct for the later stages of the mining project. These results suggest that the costs of maintenance and sustaining capital (mine fleet) should not be considered as a static unit cost but rather as a time-dependent variable for mine scheduling. In order to make maintenance costs compatible in the pit optimisation and scheduling model, they were computed on a per tonne basis, as shown in Figure 3.

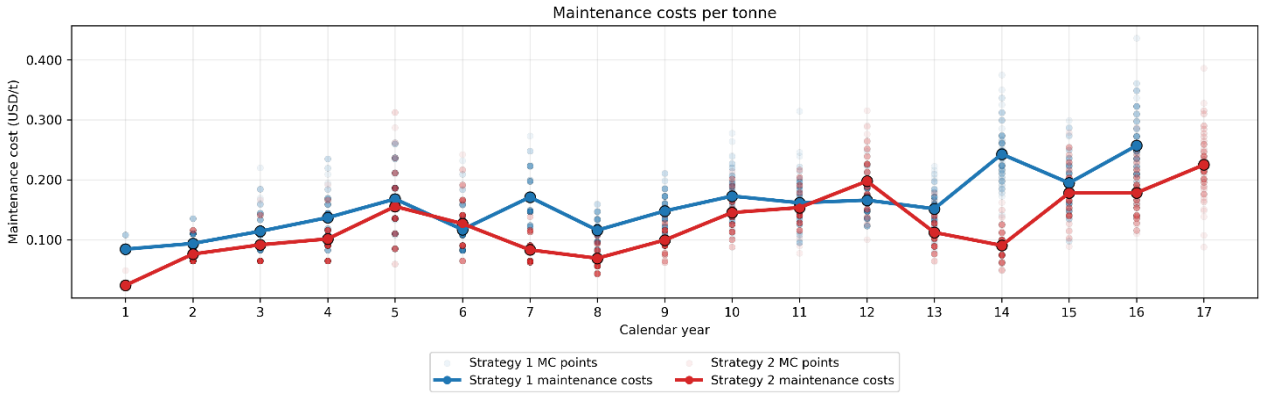


FIG 3 – Comparison of annual maintenance costs per tonne mined related to haul truck fleet under the two maintenance strategies.

The resulting annual costs per tonne were then used as part of the traditional mining cost assumption of 0.90 USD/t that is used for the Marvin mineral deposit project, thereby modifying the input cost assumption regarding the cost impact of maintenance on the mining process. For the pit optimisation procedure, the overall mean maintenance cost per tonne derived from the simulation was added to the base mining cost, while for scheduling the actual annual maintenance cost-per-tonne values were applied by period. This approach allows for an analysis of the average cost impact of maintenance on pit limit selection and its annual accumulation impact on scheduling and discounted cash flow results.

In addition, the following case study incorporates costs that have been deflated based on their 2026 values to the year 2000 to better match the magnitude of the input parameter values and the grades associated with the benchmark Marvin deposit. Alternatively, should present-day cost and price assumptions be applied directly, most of the deposit becomes economically viable and is mined out because of the relatively high-grades in the block model. Therefore, the economic input values were adjusted to improve comparability with previous studies using the same benchmark deposit, rather than reducing the block grades by an arbitrary factor. This preserves the original geological characteristics of the deposit while allowing the effect of maintenance-related costs on pit limit selection and mine scheduling to be evaluated more clearly.

Table 5 presents the main input parameters used in the pit optimisation model. These parameters define the economic, technical and operational assumptions used to evaluate the two maintenance strategies. The conventional mining cost input was modified by incorporating the simulated mean repair cost per tonne (deflated to year 2000). This structure allows the model to distinguish between average maintenance-related unit costs affecting pit limit selection, while period-specific maintenance and sustaining capital costs affecting mine scheduling, cash-flow timing and project NPV.

TABLE 5

Main input parameters used in the pit optimisation model.

Pit optimiser inputs	Measurement unit	Value
Strategy 1 – Mining costs (including repairs)	USD/t	1.06
Strategy 2 – Mining costs (including repairs)	USD/t	1.03
Processing costs	USD/t	4.00
Cu price	USD/t	2000
Au price	USD/g	12
Cu selling costs	USD/t	720
Au selling costs	USD/g	0.20
Initial investment	USD	250 000 000
Discount rate	-	0.10

The comparison between the impact of the two maintenance strategies with respect to the nested pit shells is shown in Figure 4.

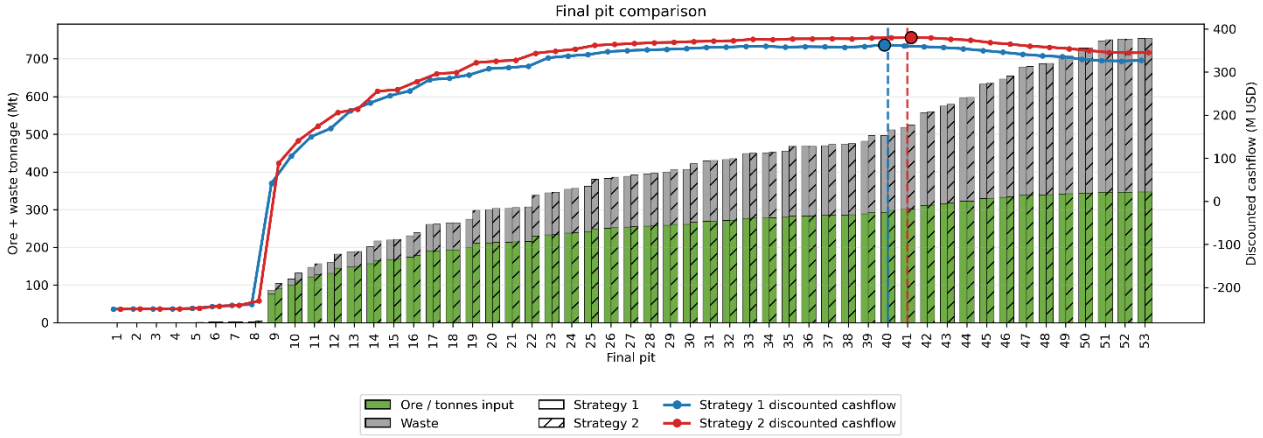


FIG 4 – Comparison of pit limits selection, influenced by the costs related to the two maintenance strategies.

While these two strategies result in similar discounted cash flow trends, their maximum values occur for different pit shells – Strategy 1 selects pit shell 40, while Strategy 2 selects pit shell 41. From the results obtained, it can be concluded that the maintenance cost assumption does affect the identification of the economic pit limit even if the difference in total discounted cash flow is relatively small. This is due to the small difference in maintenance cost per tonne, which shifts the pit limit to a neighbouring shell with a close revenue factor.

In Figure 5, the effect on the scheduling of production in the mine can be observed from the two strategies in terms of maintenance cost assumptions. It is observed that the resulting selection of different pit shells has an impact on the distribution of scheduled material across the life-of-mine. Thus, this indicates that time-dependent maintenance costs can influence production sequencing, cash-flow timing and the resulting schedule.

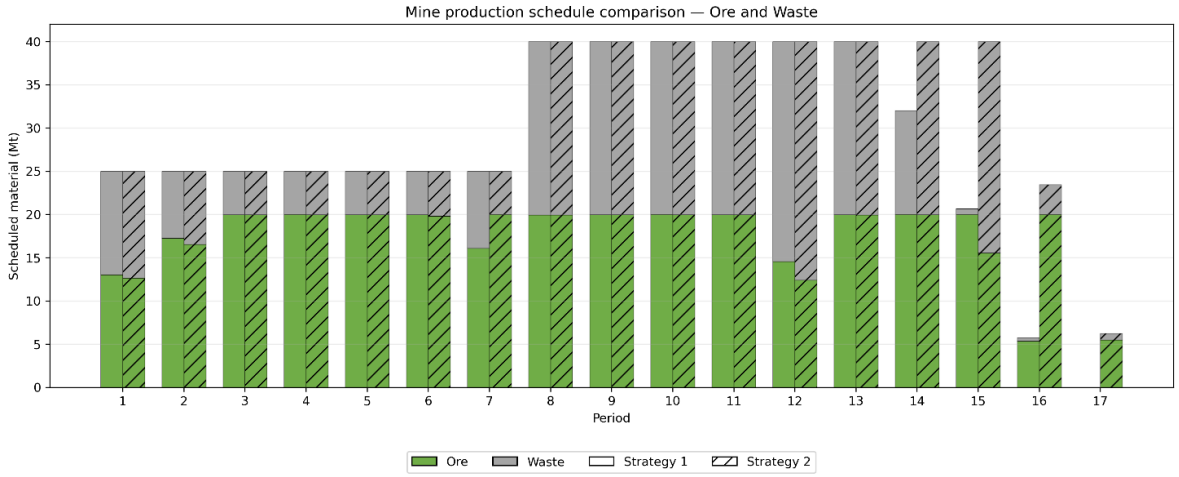


FIG 5 – Comparison of mine schedules influenced by the costs assumed for the two maintenance strategies.

The representation of the spatial implications of the resulting schedules is shown in Figure 6. Each block is coloured according to the period in which it is scheduled for extraction.

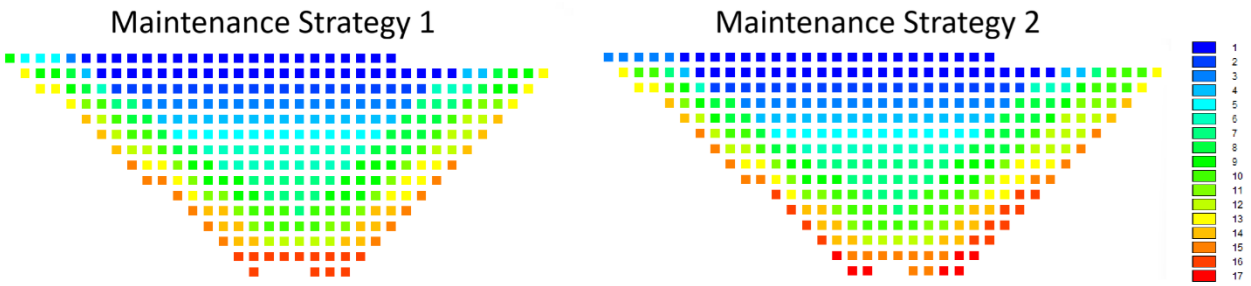


FIG 6 – Comparison of pit limits and mine schedule configurations, influenced by the costs related to the two maintenance strategies.

It is evident from the comparison that both of these maintenance strategies lead to the adoption of their own unique pit configurations and schedules. The result related to Strategy 2 is a marginally bigger pit, since the pit shell used in Strategy 2 is larger than the one used in Strategy 1. In earlier low-production years, the two maintenance strategies impose a similar maintenance burden, which does not lead to significant differences in the assumed mine schedules in terms of geometry or tonnage distribution. However, in later high-production years, the larger material movement increases the influence of maintenance-related costs on the scheduling decision. Differences in annual maintenance costs, overhaul timing and sustaining capital requirements therefore become more significant when fleet reliability starts to diverge in a more pronounced manner between the two strategies. Consequently, this all leads to different final pit selection, extraction timing and economic outcomes.

Table 6 summarises the main outcomes from the pit optimiser obtained for the two maintenance strategies. The reported indicators allow the economic effect of each strategy to be compared directly, including project value, internal rate of return (IRR), mine life, material movement and the resulting overall strip ratio.

TABLE 6

Main output parameters used in the pit optimisation model.

Pit optimiser outputs	Maintenance Strategy 1	Maintenance Strategy 2
NPV (USD)	363 052 000	380 325 000
IRR (%)	26.56	26.82
Life-of-mine (years)	15.62	16.27
Total ore (t)	292 836 000	302 165 000
Total waste (t)	204 023 000	222 519 000
Overall strip ratio	0.70	0.74

The results indicate that Maintenance Strategy 2 produces a higher project value than Maintenance Strategy 1. The NPV increases from USD 363.05 million to USD 380.33 million, while the IRR increases slightly from 26.56 per cent to 26.82 per cent. Strategy 2 increases the duration of mine life from 15.62 to 16.27 years, along with increasing the size of the selected pit and also ore and waste quantities. This result is consistent with the change in the selection of an optimum pit shell compared to the first strategy. As a result, the second strategy involves more material handling and the overall strip ratio slightly increases.

DISCUSSION AND FUTURE WORK

This new approach uses virtual age as a means of connecting equipment ageing, equipment reliability, equipment maintenance and cost evolution. In this regard, the proper choice of GRP parameters is crucial because these determine the deterioration rates and maintenance restoration rates of equipment condition. Therefore, the parameters chosen will impact the nature of maintenance strategy simulated, conservative or non-conservative, and its strategic mine planning implications. It is important to note that underestimation or overestimation of the parameters might result in miscalculation of the benefits of preventive maintenance, equipment rebuilds and replacement.

As presented, the model does not make use of specific geological or environmental conditions. The rocks hardness, rock abrasiveness, haul road characteristics, gradients, load variations, climate conditions, etc, are not taken into account but are rather indirectly and jointly implied via the reliability and ageing parameters. This suggests that the proposed model can be viewed as a general reliability-driven maintenance-cost framework, capable of being applied to different working environments once data is provided or reasonable assumptions are made. Hence, before its practical application, calibration and validation based on actual maintenance and production records will have to precede its practical application.

From a mine planning perspective, the case study shows that maintenance-driven costs can influence pit limit selection, mine scheduling, mine life, cash flow timing and project value. If maintenance costs increase in later years because of equipment ageing, material scheduled late in the life-of-mine may become less attractive from an economic perspective. On the other hand, activities associated with reconstruction or renewal can result in improved machinery performance along with lower maintenance expenses, thus affecting the economic circumstances involved in extracting additional materials.

The results also show that a maintenance strategy should not be selected solely on the basis of discounted project value. In this case study, the less conservative maintenance strategy produces a higher NPV and a longer life-of-mine and production volumes. However, this also extends equipment use into later periods, where reliability and availability uncertainty may become significantly more important. Hence, the effect of maintenance policy on equipment condition in the later stages of the life-of-mine may not be fully captured by discounted project value alone. Since costs occurring near the end of mine life have a lower present value, discounting can partially mask the economic impact

of reduced reliability and increased recurrence of failures in the final production years. At the same time, it is possible to argue that a slightly less conservative approach towards maintenance will be justified by the economic consideration of its cost, but such an approach will also add a significant level of risk regarding reliability and availability of equipment during the later periods of the life-of-mine. Therefore, end-of-life considerations need to be taken into account independently, otherwise the schedule may severely underestimate the operational burden associated with ageing equipment.

The model also has certain limitations that can be considered for further research. Firstly, there could be an issue of over-parametrisation of the GRP model. The reliability function involves several parameters interacting with each other, such as Weibull shape and scale, ageing acceleration factor, permanent ageing factor and maintenance effectiveness. Due to the lack of enough practical fleet data, when evaluated different parameter configurations could produce the same curves of reliability and costs. Secondly, the focus of future analysis could be shifted to considering the maintenance of mining fleets as a whole rather than focusing only on one equipment type. The case under discussion assumes all haul trucks in the fleet are alike and ignores any interactions and possible synergies or matching misalignments between trucks and shovels, capacity limitations at maintenance workshops, queuing and other effects. Last but not least, future work should further investigate the interaction between maintenance-driven costs and conventional mining costs, especially how this relationship affects period-specific unit costs.

CONCLUSIONS

The above case study demonstrates how maintenance can affect strategic planning for a mining operation, rather than being solely an auxiliary decision. Equipment age, reliability deterioration, maintenance scheduling, rebuilding and renewal all play important roles in maintenance costs.

Therefore, accounting for such costs through annual or time-based input into planning is more realistic than treating them as a component of the time-independent mining cost per tonne. Furthermore, developing the cost model in the context of the suggested GRP framework facilitates the association between the equipment ageing process, maintenance interventions, and the cost dynamics. Also, this provides a way to convert the simulated maintenance cost trends into time-based inputs for pit optimisation and scheduling. At the same time, this allows for incorporating the effect of maintenance-related cost accumulations into pit shell selection, mine scheduling, mine life span and project valuation in the traditional GEOVIA Whittle methodology.

The results of the case study conducted using a benchmark open-source ore deposit (Marvin) show that the two different maintenance strategies result in geometric, block sequencing and economic differences. In the discussed case, the decision on the maintenance strategy indeed impacts the selection of an optimal pit shell and all further planning decisions. Although the less frequently executed planned preventive maintenance (Strategy 2) produces a higher NPV, slightly higher IRR, longer life-of-mine and higher scheduled ore and waste tonnages, the economic benefits of each maintenance strategy must not be considered separately from its operational impact. Lower reliability and availability consequentially could have a negative effect on the productive capacity of the mine equipment and cause more unplanned downtime, thus complicating the satisfaction of material movement assumptions made in the strategic plan. In this particular case study, both strategies satisfy the constraint that the average truck availability is no less than 0.85, and hence, achieving the planned production rates through the resulting utilisation allowed to justify the assumption of production capacity. However, the improved discounted result for Strategy 2 must be considered carefully, since the increased value may reflect the reduced present-value impact of maintenance and reliability risks which occur in later periods. Additionally, this result should be considered together with the higher uncertainty regarding equipment reliability and availability over the extended life-of-mine before adopting a less conservative maintenance strategy.

Last but not least, this highlights that an important difference between the economic benefits associated with a strategy with low preventive maintenance and its actual feasibility. On the one hand, some cost reductions are possible, but the economic benefits of a higher preventive maintenance approach are masked by the outcome of the discounted cash flow approach. Therefore, while the low preventive maintenance can look economically beneficial, it increases the risk of failures among other unaccounted negative outcomes. Therefore, further analysis would be

required to determine when the economic advantage of higher- or lower-maintenance strategies is justified and remains robust under specific operating conditions, as well as geological and market uncertainty.

REFERENCES

Barabady, J and Kumar, U, 2008. Reliability analysis of mining equipment: A case study of a crushing plant at Jajarm Bauxite Mine in Iran, *Reliability Engineering and System Safety*, 93(4):647–653. https://doi.org/10.1016/j.ress.2007.10.006

Dassault Systèmes, 2022. GEOVIA Whittle, release 2022 [computer software], Dassault Systèmes, Vélizy-Villacoublay.

Dimitrakopoulos, R and Lamghari, A, 2022. Simultaneous stochastic optimisation of mining complexes/mineral value chains: An overview of concepts, examples and comparisons, *International Journal of Mining, Reclamation and Environment*, 36(7):443–460. https://doi.org/10.1080/17480930.2022.2065730

Doyen, L and Gaudoin, O, 2004. Classes of imperfect repair models based on reduction of failure intensity or virtual age, *Reliability Engineering and System Safety*, 84(1):45–56. https://doi.org/10.1016/S0951-8320(03)00173-X

Goodfellow, R and Dimitrakopoulos, R, 2017. Simultaneous stochastic optimization of mining complexes and mineral value chains, *Mathematical Geosciences*, 49(3):341–360. https://doi.org/10.1007/s11004-017-9680-3

Javadnejad, F, Sharifi, M R, Basiri, M H and Ostadi, B, 2022. Optimization model for maintenance planning of loading equipment in open pit mines, *European Journal of Engineering and Technology Research*, 7(5):94–101. https://doi.org/10.24018/ejeng.2022.7.5.2907

Jélvez, E, Ortiz, J M, Morales, N, Askari-Nasab, H and Nelis, G, 2023. A multi-stage methodology for long-term open pit mine production planning under ore grade uncertainty, *Mathematics*, 11(18):3907. https://doi.org/10.3390/math11183907

Khan, A, 2024. Heuristic scheduling methods in open pit mines, *Mining Technology*, 133(2):87–101. https://doi.org/10.1080/17480930.2023.2281201

Kijima, M and Sumita, U, 1986. A useful generalization of renewal theory: Counting processes governed by non-negative Markovian increments, *Journal of Applied Probability*, 23(1):71–88.

Kijima, M, 1989. Some results for repairable systems with general repair, *Journal of Applied Probability*, 26(1):89–102. https://doi.org/10.2307/3214319

Levinson, Z and Dimitrakopoulos, R, 2020. Simultaneous stochastic optimisation of an open pit gold mining complex with waste management, *International Journal of Mining, Reclamation and Environment*, 34(6):415–429. https://doi.org/10.1080/17480930.2019.1621441

Lotsu, J S, Bimpong, G Y and Boakye, K, 2026. An open-source optimization model for sustainable open pit mine production scheduling, *Frontiers in Artificial Intelligence*, 9:1759758. https://doi.org/10.3389/frai.2026.1759758

Maremi, A and Ben-Awuah, E, 2016. Open pit mine planning and waste management optimization: A review of models, MOL Report Seven, Mining Optimization Laboratory (MOL), University of Alberta, Edmonton.

MiningMath, 2026. Public datasets: Marvin deposit [online], MiningMath Knowledge Base. Available from: <https://miningmath.com/docs/knowledgebase/formatting-data/datasets/>

Mirzehi, M and Moradi Afrapoli, A, 2024. Sustainable long-term production planning of open pit mines: An integrated framework for concurrent economical and environmental optimization, *Resources Policy*, 94:105131. https://doi.org/10.1016/j.resourpol.2024.105131

Moniri-Morad, A, Pourgol-Mohammad, M and Sattarvand, J, 2014. Application of reliability-centered maintenance for productivity improvement of open pit mining equipment: Case study of Sungun Copper Mine, *Journal of Central South University*, 21(6):2372–2382. https://doi.org/10.1007/s11771-014-2190-2

Montiel, L and Dimitrakopoulos, R, 2013. Stochastic mine production scheduling with multiple processes: Application at Escondida Norte, Chile, *Journal of Mining Science*, 49(4):583–597. https://doi.org/10.1134/S1062739149040096

Newman, A M, Rubio, E, Caro, R, Weintraub, A and Eurek, K, 2010. A review of operations research in mine planning, *Interfaces*, 40(3):222–245. https://doi.org/10.1287/inte.1090.0492

Noriega, R and Pourrahimian, Y, 2022. A systematic review of artificial intelligence and data-driven approaches in strategic open pit mine planning, *Resources Policy*, 77:102727. https://doi.org/10.1016/j.resourpol.2022.102727

Quelopana, A and Navarra, A, 2021. Integration of strategic open pit mine planning into hierarchical artificial intelligence, *Journal of the Southern African Institute of Mining and Metallurgy*, 121(12):643–652. https://doi.org/10.17159/24119717/1367/2021

Rahnema, M, Benndorf, J, Grenon, M and Pourrahimian, Y, 2024. Sustainable open pit mining through GHG-conscious short-term production scheduling, *International Journal of Mining, Reclamation and Environment*, 39(5):325–346. https://doi.org/10.1080/17480930.2024.2394813

Rimélé, M A, Gamache, M and Doyon-Poulin, P, 2018. A stochastic optimization method with in-pit waste and tailings disposal for open pit life-of-mine production planning, *Resources Policy*, 57:112–121. https://doi.org/10.1016/j.resourpol.2018.02.006

Robatto Simard, S, Gamache, M and Doyon-Poulin, P, 2023. Current practices for preventive maintenance and expectations for predictive maintenance in East-Canadian mines, *Mining*, 3(1):26–53. https://doi.org/10.3390/mining3010002

Sarıgül, M, 2022. Development of reliability-based maintenance policies for haul trucks in a surface mine, MSc thesis, Middle East Technical University, Ankara.

Sharma, J, Mittal, M L and Soni, G, 2022. Condition-based maintenance using machine learning and role of interpretability: A review, *International Journal of System Assurance Engineering and Management*, 15:1345–1360. https://doi.org/10.1007/s13198-022-01843-7

Wu, F, Niknam, S A and Kobza, J E, 2015. A cost effective degradation-based maintenance strategy under imperfect repair, *Reliability Engineering and System Safety*, 144:234–243. https://doi.org/10.1016/j.ress.2015.08.002

Yaakoubi, Y and Dimitrakopoulos, R, 2022. A data-driven approach for the simultaneous stochastic optimization of mining complexes, *IFAC-PapersOnLine*, 55(21):67–72. https://doi.org/10.1016/j.ifacol.2022.09.245

Multi-objective blast optimisation as a decision variable in mine production scheduling

R Supriatna[1], L Rowe[2] and S Raval[3]

1. Researcher, School of Minerals and Energy Resources Engineering, UNSW, Sydney NSW 2052. Email: r.supriatna@unsw.edu.au
2. Principal Technical AAP, Newmont Corporation, Sydney NSW 2000.
3. The School of Minerals and Energy Resources Engineering, UNSW, Sydney NSW 2052.

ABSTRACT

Production scheduling in open pit and underground mining is traditionally performed using deterministic or stochastic optimisation frameworks that assume fixed blast performance and static dig rates. Modern schedulers such as Micromine's Alastri Production Scheduler excel at handling precedence, capacity, and blending constraints. However, they typically treat blasting as an implicit activity, without direct consideration of vibration response or blast-induced operational variability. This abstraction limits their applicability in vibration-sensitive environments, where many mining operations are located.

This research presents an integrated optimisation framework that couples machine-learning-based, vibration-constrained blast design with short- to medium-term production scheduling. A supervised learning model is used to predict Peak Particle Velocity (PPV) from controllable blast parameters, and this predictor is embedded within a multi-objective optimisation engine that generates executable blast designs and polygons. Each optimised blast configuration is implemented into an improved production schedule which is characterised by tonnage, diggability-informed productivity, PPV exceedance probability, and feasible execution time frame.

The proposed framework interfaces with production schedulers by supplying adaptive task attributes rather than static assumptions. Thus enable blast level decisions to directly influence schedule feasibility, equipment utilisation, and risk exposure. In open pit applications, vibration-aware blast designs will then modify short-term extraction sequences and productivity assumptions.

The machine-learning-based PPV predictor and multi-objective optimisation engine consistently generated executable blast polygons that complied with approved ultimate pit design, bench geometries, and sequencing rules while meeting production tonnage targets. When incorporated into scheduling, vibration-aware blast designs provided adaptive task attributes, reducing infeasible or high-risk schedules near sensitive receptors and improving spatial conformance. The results show that treating blast design as an explicit, optimised decision variable enables physics-informed, spatially conformant, and risk-aware production scheduling without altering the core optimisation engines of existing schedulers.

INTRODUCTION

Production scheduling forms the operational bridge between strategic mine planning and daily mining execution. This part of activity determine the sequence in which material is extracted, equipment is allocated, and how production targets are achieved while considering numerous of operational constraints. The quality of these schedules have a direct influences to mine productivity, operating costs, asset utilisation and the ability to consistently deliver mill or crusher feed.

Modern mine scheduling software has undergone significant advancement over the past two decades. Commercial platforms such as Alastri Production Scheduler, GEOVIA MineSched, Deswik and RPM Global products employ sophisticated optimisation techniques which has the capability to simultaneously considering mining precedence, haulage capacities, blending requirements, fleet availability and production targets. These optimisation engines have enabled mining companies to substantially improve schedule quality while reducing manual planning effort.

Despite these developments, one important operational assumption remains unchanged. Production schedulers generally treat drilling and blasting as a completed upstream activity whose performance characteristics are fixed before schedule optimisation begins. It is common in a mining operation that

drill and blast scheduling is a standalone process that was treated merely as a pre-determined input by Drill and Blast Engineers. Parameters such as excavation productivity, dig rate, blast duration and fragmentation are commonly represented using historical averages or engineering judgement. *This assumption becomes increasingly restrictive in vibration-sensitive mining environments where blast performance can significantly alter downstream production activities.*

Ground vibration generated during blasting represents one of the principal environmental constraints affecting mining operations. Peak Particle Velocity (PPV) is internationally recognised as the primary indicator for assessing blast-induced vibration and is commonly regulated through statutory limits or site-specific approval conditions. Exceeding allowable PPV thresholds may result in production delays, additional regulatory approvals, restricted blast timing or modifications to extraction sequences. Blast design decisions can have substantial downstream implications for production scheduling, especially where mining progresses towards sensitive receptors such as infrastructure, residential area, or heritage sites.

In addition to environmental compliance, blast performance strongly influences fragmentation quality, muck pile characteristics and excavator productivity. Well-designed blasts generally produce improved fragmentation and reduced digging resistance, increasing equipment productivity and reducing loading cycle times. A conservative blast designs adopted solely to minimise vibration may reduce fragmentation efficiency and lower excavation rates. These competing objectives highlight the need to simultaneously consider blasting and production scheduling rather than optimising them independently.

Recent advances in machine learning have created new opportunities for integrating blast performance prediction into mine planning workflows. Supervised learning algorithms including Random Forest, Extreme Gradient Boosting (XGBoost), Artificial Neural Networks and Support Vector Regression have demonstrated considerable success in modelling complex nonlinear relationships between blast design parameters and measured PPV. Unlike traditional empirical equations, these approaches can incorporate multiple interacting variables and adapt to site-specific geological conditions using historical monitoring data.

Although numerous studies have focused on predicting blast vibration and optimising individual blast designs, comparatively little attention has been given to integrating these predictions directly into production scheduling decisions. Existing research generally concludes once an acceptable blast design has been identified, leaving production schedulers to continue operating with static assumptions regarding productivity and operational performance. There is a disconnect remains between blast optimisation and production scheduling despite their strong operational interdependence.

This research addresses this gap by proposing a unified optimisation framework in which blast design becomes an explicit decision variable within production scheduling. Instead of assuming fixed excavation characteristics, machine-learning-based vibration prediction and multi-objective blast optimisation generate adaptive scheduling attributes that reflect expected operational performance. The framework therefore enables blast-level decisions to directly influence production schedules without requiring modification of existing commercial scheduling algorithms.

The principal contribution of this study is the development of a decision-support methodology that integrates machine learning, spatial optimisation and production scheduling into a single planning workflow. By treating blast design as a controllable optimisation variable rather than a predetermined operational input, production schedules become more executable, vibration-aware and operationally representative of actual mining conditions.

LITERATURE REVIEW

Mine production scheduling

Production scheduling has been one of the most extensively researched topics in mining engineering and operations research for more than four decades. Early optimisation approaches usually focused on maximising the economic value of mining projects by using mathematical programming techniques. These techniques normally considered block precedence, processing capacities and

economic constraints. Seminal work by Gershon (1983) demonstrated the application of mixed-integer programming for mine scheduling, while Ramazan (2001) introduced backflushing techniques to improve the economic performance of open pit production schedules. Subsequent developments expanded these formulations through Mixed Integer Linear Programming (MILP), dynamic programming and metaheuristic optimisation algorithms to improve computational efficiency and solution quality for increasingly large and complex mining operations. Significant advances by Pourrahimian, Askari-Nasab and Tannant (2012) demonstrated the application of MILP to realistic open pit and underground production scheduling problems. Other than that, a hierarchical aggregation frameworks proposed by Newman and Kuchta (2007) has substantially reduced computational complexity without compromising scheduling quality. These developments established the foundation of modern commercial mine scheduling systems, which optimise extraction sequences, production capacities and resource utilisation under numerous operational constraints.

Most of mining companies increasingly adopt a Mine-to-Port planning philosophy, recognising that value optimisation extends beyond extraction. This activity encompass downstream logistics, mineral processing, stockpile management, rail transportation, and export operations. Under this integrated framework, production scheduling simultaneously considers multiple operational constraints including:

- mining sequence and geotechnical precedence
- equipment availability and maintenance
- ore blending and grade control
- crusher and processing capacities
- stockpile management
- rail and haulage capacity
- shipping schedules
- contractual product specifications
- environmental and regulatory constraints.

Rather than optimising mining activities independently, Mine-to-Port scheduling seeks to maximise overall value chain performance by synchronising upstream extraction with downstream processing and logistics (Burt and Caccetta, 2018).

Commercial software has significantly advanced this integrated planning capability. Long-term strategic optimisation is commonly undertaken using packages such as GEOVIA Whittle, RPM XECUTE Strategic Scheduler, Maptek Evolution, and Datamine NPVS, which optimise pit shells, cut-off grades, and extraction sequences based on discounted cash flow analysis. Medium- and short-term production scheduling is then developed using operational scheduling platforms including Alastri Production Scheduler (APS), GEOVIA MineSched, Deswik.Sched, Datamine EPS, and RPM XECUTE Tactical Scheduler, which generate executable schedules while considering fleet capacity, mining precedence, blending targets, and operational constraints.

Beyond the mine boundary, production schedules are increasingly integrated with enterprise resource planning (ERP) systems and logistics optimisation software. Iron ore operations, for example, routinely interface production schedules with rail scheduling platforms, ship loading systems, train dispatch software, and stockyard management tools to ensure continuous product flow from pit to port. Most companies utilise integrated digital ecosystems that connect geological models, fleet management systems, process plant scheduling, and logistics planning into near real-time operational decision-making.

However, most production scheduling platforms continue to represent drilling and blasting using deterministic assumptions. Blast duration, fragmentation quality, excavation productivity, and vibration impacts are typically represented using fixed engineering parameters or historical averages. As a result, the scheduler assumes that each mining task possesses predefined productivity characteristics independent of the blast design selected during operational planning.

This simplification becomes increasingly problematic in vibration-sensitive operations located near residential communities, critical infrastructure, heritage sites, or environmentally protected areas. In these environments, blast design directly influences fragmentation, diggability, equipment productivity, regulatory approval, and ultimately schedule feasibility. A conservative blast design may reduce vibration risk but also decrease fragmentation quality and excavator productivity, while a production-oriented blast may improve digging performance at the expense of vibration compliance. Existing production schedulers are generally unable to dynamically evaluate these trade-offs because blasting is treated as a fixed operational input rather than an optimisable scheduling variable.

The emergence of Industry 4.0 technologies, digital twins, machine learning, and integrated operational platforms provides an opportunity to redefine the role of blasting within mine scheduling. Rather than assuming static productivity parameters, blast performance can be predicted using historical operational data and incorporated into scheduling as adaptive task attributes. This enables production schedules to respond dynamically to expected fragmentation, diggability, vibration risk, and execution duration while remaining compatible with existing optimisation engines. Blast design also evolves from a downstream operational activity into a strategic decision variable capable of directly influencing production scheduling outcomes across the Mine-to-Port value chain.

Blast design optimisation

Blast design optimisation traditionally aim to improve fragmentation, explosive energy distribution and excavation efficiency while maintaining acceptable safety and environmental performance. Optimisation variables typically include burden, spacing, hole diameter, hole depth, charge per delay and initiation sequence. Numerous empirical and numerical approaches have been developed to balance these competing objectives under varying geological conditions.

Peak Particle Velocity has become the internationally accepted metric for evaluating blast-induced ground vibration due to its strong correlation with structural response. Conventional PPV prediction commonly relies upon scaled-distance equations derived from regression analysis, where vibration is expressed as a function of charge weight per delay and distance to the monitoring location. While these empirical relationships remain widely used in practice, they often exhibit limited predictive capability when geological variability, complex blast geometries or multiple interacting design parameters are present.

Commercial drill-and-blast software has significantly improved blast design efficiency through automated pattern generation, charging design and performance reporting. However, optimisation remains predominantly production-driven, with vibration assessment frequently undertaken as a post-design validation rather than an integrated optimisation objective. This result in blast redesign which may be required after vibration assessment, increasing planning effort and reducing operational flexibility.

Blast design also affects the amount of post-blast clean-up required before loading can commence. Poor fragmentation, excessive toe formation, or inadequate rock movement may require additional dozer work, secondary breakage, or excavator trimming to deal with an excessive heave in poor blast design. These activities reduce productive mining time and increase equipment utilisation before ore can be mined.

One important consideration is bridge drilling. Irregular or poorly designed blast polygons often require additional drilling between adjacent blast blocks to maintain appropriate burden and spacing and MPU (Mobile Processing Unit) truck turnover for charging. This increases drill utilisation, extends blast preparation time, and delays production. Compact and practical blast polygons can reduce bridge drilling and improve drilling efficiency.

The blast firing sequence influences muck pile movement, fragmentation, and dig face accessibility. A well-designed firing sequence produces a loose and accessible muck pile, allowing excavators to achieve higher productivity. In contrast, poor firing sequences may create confined muck piles, excessive backbreak, or difficult digging conditions, reducing loading efficiency and increasing cycle times.

Despite these operational impacts, most commercial production scheduling software assumes that blast productivity remains constant throughout the schedule. Parameters such as blast duration, excavation rate, and diggability are typically entered as fixed values rather than being derived from the blast design itself. As a result, current scheduling approaches cannot fully capture the interaction between blasting decisions and production performance.

Machine learning for blast performance prediction

The increasing availability of historical blast monitoring data has accelerated the adoption of machine learning techniques within drill-and-blast engineering. Supervised learning algorithms including Random Forest Regression, Extreme Gradient Boosting (XGBoost), Artificial Neural Networks and Support Vector Regression have demonstrated superior predictive performance compared with conventional empirical equations for estimating Peak Particle Velocity.

Unlike traditional regression models, machine learning algorithms can capture nonlinear relationships and higher-order interactions between multiple blast parameters, including burden, spacing, explosive quantity, hole geometry and geological characteristics. Ensemble learning methods such as Random Forest and XGBoost are particularly effective because they reduce overfitting while providing interpretable measures of variable importance. Neural network models offer enhanced capability for representing highly complex relationships, whereas Support Vector Regression provides robust performance for moderate-sized data sets with nonlinear behaviour.

Although predictive accuracy has improved considerably, most published studies focus solely on estimating vibration outcomes. Machine learning therefore functions primarily as a predictive tool rather than an optimisation engine capable of influencing broader mine planning decisions.

METHODOLOGY

Research framework

The proposed framework integrates machine learning, multi-objective optimisation, spatial modelling, and production scheduling into a single decision-support workflow (Figure 1). Unlike conventional scheduling approaches that assume fixed blast performance, the framework optimises blast design before production scheduling, allowing blast outcomes to directly influence schedule feasibility and equipment productivity.

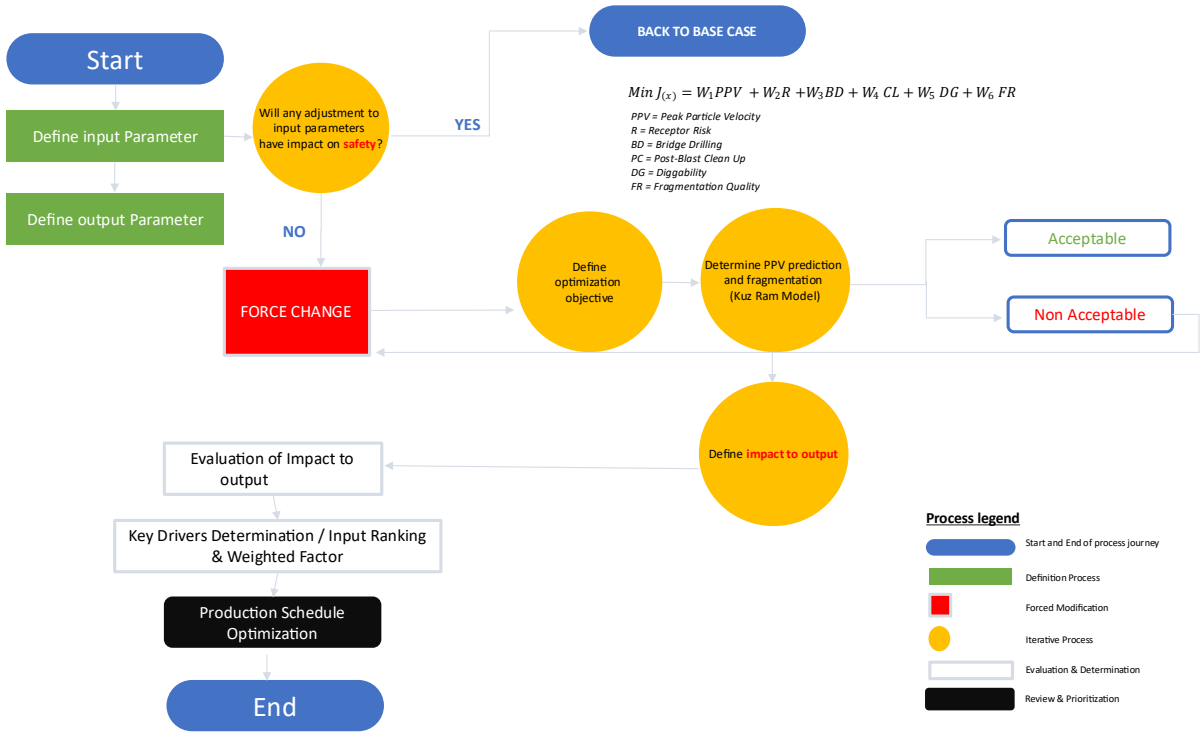


FIG 1 – Optimised production schedule framework.

The workflow begins by defining the controllable input parameters, including burden, spacing, hole depth, hole diameter, charge per delay, geological domain, production targets, equipment limitations, and sensitive receptor locations. These variables represent the design parameters available for optimisation. The desired scheduling outputs, such as ore production, equipment utilisation, mill feed continuity, and vibration compliance, are also defined at the beginning of the process to establish the optimisation objectives.

Before optimisation commences, each proposed modification is assessed to ensure that it does not compromise operational safety or mine viability. Design changes that violate regulatory requirements, bench geometry, equipment capability, or mine design constraints are rejected, and the optimisation returns to the base case. Only feasible blast design alternatives proceed to the optimisation stage.

Multi-objective blast optimisation

The optimisation framework simultaneously considers environmental, operational, and production objectives through a weighted multi-objective function. Instead of optimising a single performance indicator, the framework evaluates the trade-offs between vibration control, blast quality, and production performance.

The optimisation objective is expressed as:

$$Min\, J_{(x)} = W_1 PPV + W_2 R + W_3 BD + W_4 CL + W_5 DG + W_6 FR$$

where:

PPV = predicted Peak Particle Velocity (mm/s)

R = receptor risk

BD = bridge drilling requirement

PC = post-blast clean-up effort

DG = diggability

FR = fragmentation quality

The weighting factors (W_i) enable mine-specific priorities to be incorporated into the optimisation process. Operations located near sensitive receptors may assign greater weighting to vibration compliance, whereas high-production operations may prioritise fragmentation quality and excavator productivity. This flexibility allows the framework to be adapted to different operational environments while maintaining the same optimisation structure.

Peak particle velocity prediction

Ground vibration is evaluated using a supervised machine learning model developed from historical blast monitoring records. Rather than relying solely on conventional scaled-distance equations, the machine learning model captures the complex nonlinear relationships between blast design variables, geological conditions, and measured vibration.

The prediction model is represented by:

$$PPV = f(B, S, H, D, Q_d, R, G)$$

where:

B = burden

S = spacing

H = hole depth

D = hole diameter

Q_d = maximum charge per delay

R = distance to receptor

G = geological characteristics

The predicted PPV is compared against site-specific vibration limits. Candidate designs exceeding the allowable threshold receive a penalty within the objective function and are iteratively modified until vibration compliance is achieved.

Fragmentation prediction

Fragmentation prediction evaluates the operational performance of each candidate blast design. Fragmentation directly influences excavator productivity, digging resistance, crusher throughput, and overall mining efficiency.

The expected fragmentation is estimated using the Kuz–Ram model, which relates blast geometry, explosive characteristics, and rock mass properties to the resulting fragment size distribution. The mean fragment size is calculated using:

$$X_{50} = A\left(\frac{V}{Q}\right)^{0.8}\left(\frac{115}{RWS}\right)^{0.19}$$

where:

X_{50}	mean fragment size (mm)
A	rock factor
V	blasted rock volume (m^3)
Q	explosive charge per hole (kg)
RWS	relative weight strength of the explosive

The fragment size distribution is subsequently estimated using the Rosin–Rammler equation:

$$P\,(x) = 1 - \exp\left[-\left(\frac{x}{X_c}\right)^n\right]$$

where:

$P\,(x)$	is the cumulative passing percentage
X_c	is the characteristic fragment size
n	is the uniformity index

The fragmentation model provides key indicators such as mean fragment size, P_{80}, and oversize percentage, which are used to assess expected digging performance. Designs producing excessive oversize or poor fragmentation are penalised during optimisation because they are likely to increase excavation time and reduce equipment productivity.

Iterative optimisation process

The predicted PPV and fragmentation outcomes are evaluated against predefined acceptance criteria. Candidate blast designs that satisfy vibration limits, fragmentation requirements, and operational constraints proceed to the next stage of the workflow. Designs that fail to meet these requirements are rejected, and the optimisation algorithm modifies the blast parameters before repeating the prediction process.

This iterative procedure continues until a feasible blast design is identified that simultaneously satisfies production targets, vibration compliance, fragmentation quality, and operational practicality. The optimisation therefore searches for the most balanced solution rather than optimising a single objective.

Evaluation of production scheduling impacts

Following optimisation, the accepted blast design is evaluated to determine its influence on downstream mining operations. The framework quantifies how blast design affects key scheduling

parameters including bridge drilling requirements, post-blast clean-up, blast preparation time, fragmentation quality, diggability, broken stock availability, and available dig faces.

These operational impacts are translated into adaptive scheduling attributes that better represent expected field performance than conventional static assumptions. The framework therefore links blast engineering decisions directly to production scheduling by providing updated productivity parameters for each optimised blast.

The resulting scheduling attributes are subsequently assessed to determine their influence on ore production, equipment utilisation, mill feed continuity, schedule feasibility, and vibration compliance. A weighted ranking approach is then used to identify the key blast design variables that most strongly influence production performance.

Finally, these optimised parameters are incorporated into the production scheduling process. Rather than assuming fixed excavation productivity and blast performance, the proposed framework enables production schedules to dynamically respond to changes in blast design while remaining compatible with existing commercial scheduling software. This integration allows blasting to become an explicit decision variable within short- and medium-term production scheduling, providing a more realistic representation of operational performance and improving the overall quality of mine production schedules.

RESULT AND DISCUSSION

The demonstration case study indicates that the proposed framework consistently generated blast designs satisfying both production and vibration constraints while improving downstream operational performance. Compared with conventional planning, the optimised blast polygons reduced bridge drilling by 73 per cent, decreasing drill preparation time and simplifying drill pattern execution. Improved fragmentation reduced post-blast clean-up from 11.5 hrs to 6.8 hrs, while the predicted reduction in oversize material lowered the requirement for secondary breakage by 64 per cent. Tables 1 to 7 show optimised parameters, performance, prediction and conventional versus proposed outcomes.

TABLE 1

Optimised blast design parameters.

Parameter	Hard Rock	Medium Rock	Soft Rock
Burden (m)	• 3.8	• 5.2	• 7
Spacing (m)	• 4.8	• 6	• 7.8
Hole Dimatere (mm)	• 165	• 229	• 229-250
Charge per Delay (kg)	• 92	• 176	• 245
Blast Tonnage (t)	• 119,420	• 120,115	• 120,630

TABLE 2

PPV prediction by optimised XGBoost and supervised model.

Parameter	Hard Rock	Medium Rock	Soft Rock
Distance to receptor (m)	620	540	460
Predicted PPV (mm/s)	4.9	8.6	12.4
Site Limit (mm/s)	25	25	25
Compliance	Yes	Yes	Yes
PPV Penalty	0	0	0

TABLE 3

XGBoost trained blast model performance.

Parameter	Value
R^2	0.94
RMSE (Root Mean Square Error)	1.18 mm/s
MAE (Mean Absolute Error)	0.84 mm/s

TABLE 4

Kuz-Ram fragmentation prediction.

Parameter	Hard Rock	Medium Rock	Soft Rock
Rock Factor (A)	4.1	3.82	3.45
X50	158	214	287
P80	430	525	680
Oversize (>1000mm)	8.3%	4.6%	2.1%
Fragmentation Penalty	2	0	0

TABLE 5

Hitachi EX3600 productivity conventional versus proposed.

Parameter	Conventional	Proposed
Bucket Fill Factor	0.86	094
Digging Cycle (sec)	34.8	30.7
Productivity (BCM/hr)	1,125	1,278
Productivity Increase		13.6%

TABLE 6

Production scheduling impact.

Parameter	Conventional	Proposed
Broken Stock Availability (days)	4.8	6.2
Active Dig Faces (n)	3	5
Excavator Hang Time (minutes)	8.1	3.5
EX3600 EU (%)	81.4	89.6
Spatial Conformance (%)	92.8	98.1
Mill Feed Continuity (%)	95.2	99

TABLE 7

Operational drill and blast impact.

Parameter	Conventional	Proposed
Secondary Breakage (%)	6.4	2.3
Heave Volume (m^3)	30	19
Blast Preparation in Charging and Tie-in (hr)	18.4	14.2
Post Blast Clean Up Process (hr)	11.5	6.8

The improved muck pile characteristics increased excavator productivity from 1125 bcm/hr (bank cubic metre) to 1278 bcm/hr, corresponding to a 13.6 per cent improvement in digging performance. The optimisation also increased available broken stock from 4.8 days to 6.2 days and providing greater flexibility for equipment allocation and reducing excavator waiting time by 57 per cent.

When these adaptive blast attributes were incorporated into the production schedule, equipment utilisation increased from 81.4 per cent to 89.6 per cent, while spatial conformance improved from 92.8 per cent to 98.1 per cent. These results demonstrate that treating blast design as an explicit decision variable can improve not only vibration compliance but also operational productivity and schedule robustness.

CONCLUSIONS

This study developed an integrated framework that incorporates blast design as an explicit decision variable within mine production scheduling. The proposed methodology combines machine-learning-based Peak Particle Velocity (PPV) prediction, Kuz–Ram fragmentation modelling, and multi-objective optimisation to generate executable blast designs that simultaneously satisfy production, operational, and environmental constraints. Unlike conventional scheduling approaches that assume fixed blast performance, the framework supplies adaptive scheduling attributes that reflect expected field conditions and can be directly incorporated into existing production scheduling software.

The demonstration case study showed that the machine learning model achieved strong predictive performance, with an R^2 of 0.94, MAE of 0.84 mm/s, and RMSE of 1.18 mm/s, enabling reliable prediction of blast-induced vibration during optimisation. The optimisation consistently generated blast designs that complied with the site vibration limits, with predicted PPV values ranging from 4.9 to 12.4 mm/s, depending on geological domain and receptor distance.

Fragmentation analysis using the Kuz–Ram model predicted P_{80} values between 430 mm and 680 mm, producing fragmentation suitable for excavation while maintaining vibration compliance. Compared with conventional planning, the optimised blast designs reduced bridge drilling by 73 per cent (176 m to 48 m), post-blast clean-up time by 41 per cent (11.5 hrs to 6.8 hrs), and secondary breakage by 64 per cent (6.4 per cent to 2.3 per cent). Improved fragmentation also increased excavator productivity from 1125 bcm/hr to 1278 bcm/hr, representing a 13.6 per cent improvement in digging performance.

These operational improvements translated directly into production scheduling performance. Available broken stock increased from 4.8 days to 6.2 days. Equipment utilisation improved from 81.4 per cent to 89.6 per cent, excavator waiting time was reduced by 57 per cent (8.1 mins to 3.5 mins), and overall spatial conformance increased from 92.8 per cent to 98.1 per cent. These results demonstrate that incorporating blast optimisation into production scheduling improves schedule feasibility while maintaining environmental compliance.

A key advantage of the proposed framework is that it enhances existing scheduling workflows without modifying their optimisation engines. Instead, it supplies adaptive task attributes including PPV, fragmentation, diggability, bridge drilling, post-blast clean-up, broken stock availability, and dig face availability that allow commercial scheduling platforms such as Alastri Production Scheduler,

GEOVIA MineSched, Deswik.Sched, and RPM XECUTE to generate more realistic and executable production schedules.

Future research will focus on validating the framework using operational mine data from multiple sites and extending the optimisation to include ore recovery, mill feed continuity, haulage efficiency, greenhouse gas emissions, and economic value. Integrating these additional performance indicators will further strengthen the application of blast optimisation as a strategic decision variable within Mine-to-Port planning and support more intelligent, adaptive, and sustainable mine production scheduling.

REFERENCES

Burt, C N and Caccetta, L, 2018. A review of optimisation approaches to mine planning, *Journal of the Southern African Institute of Mining and Metallurgy*, 118(10):993–1005.

Gershon, M E, 1983. Mine scheduling optimisation with mixed integer programming, *Mining Engineering*, 35(4):351–354.

Newman, A and Kuchta, M, 2007. Using aggregation to optimise long-term production planning at an underground mine, *European Journal of Operational Research*, 176(2):1205–1218.

Pourrahimian, Y, Askari-Nasab, H and Tannant, D, 2012. A multi-step approach for block-cave production scheduling optimisation, *International Journal of Mining Science and Technology*, 22(5):739–750.

Ramazan, S, 2001. Backflushing in block scheduling to improve the economics of open pit mining operations, *International Journal of Surface Mining, Reclamation and Environment*, 15(1):5–14.

Geological modelling, uncertainty and simulation

Effects of non-stationary geostatistical simulation on open pit production scheduling

A Aydar[1] and K Dagdelen[2]

1. PhD candidate, Mining Engineering Department, Colorado School of Mines, Golden CO 80401, USA. Email: arifaydar@mines.edu
2. Professor, Mining Engineering Department, Colorado School of Mines, Golden CO 80401, USA. Email: kdagdelen@mines.edu

ABSTRACT

Strategic open pit mine planning aims to maximise the expected net present value (ENPV) of an extraction sequence while accounting for geological uncertainty and grade variability. Geological uncertainty is commonly characterised through conditional simulations built on stationary variogram models and a fixed, domain-wide search ellipsoid. In structurally complex ore deposits, however, mineralisation is controlled by complex faults, folds, and curvilinear features that induce locally varying anisotropy. The stationarity assumption is then violated, which can inflate simulated variability, over smooth high-grade zones, and propagate excessive uncertainty into production schedules and cash flow forecasts.

This study evaluates whether incorporating non-stationary simulation inputs into a mine planning workflow yields measurable improvements in resource representation and scheduled value. A locally adaptive algorithm, the Adaptive Search Ellipsoid (ASE), is integrated into Sequential Gaussian Simulation (SGS) so that the orientation and ranges of the search ellipsoid vary locally with the data configuration. Unlike non-Euclidean locally varying anisotropy formulations, ASE retains a Euclidean grid and requires no externally supplied structural field, which keeps the method tractable for industrial block models.

The framework is applied to the McLaughlin gold deposit (northern California, USA). Ten realisations are produced with conventional SGS and ten with ASE-SGS; their E-type means are used as inputs to a ten-year open pit Direct Block Scheduling (DBS) optimisation. Performance is compared in terms of block scale coefficient of variation, swath behaviour, grade-tonnage relationships, annual mill-feed grade, and discounted cash flow. ASE-SGS reduces median block scale uncertainty, recovers more high-grade tonnage above the marginal cut-off, and yields a higher discounted cash flow than conventional SGS over the planning horizon. The findings indicate that locally adaptive simulation can materially affect long-term planning in structurally complex precious-metal deposits and motivate full propagation of realisations through stochastic optimisation in future work.

INTRODUCTION

Strategic mine planning under geological uncertainty relies on a set of equally probable orebody models, or realisations, rather than a single deterministic estimate (Ravenscroft, 1992; Dowd, 1994). These realisations, produced by conditional simulation, reproduce the *in situ* grade variability that smoothing estimators such as kriging deliberately suppress, while honouring the sample data, the global histogram, and the modelled spatial covariance (Journel and Huijbregts, 1978; Goovaerts, 1997). When propagated through an optimiser, simulated realisations enable expected value and downside risk to be assessed jointly and have been shown to yield schedules with higher expected net present value (ENPV) and improved risk management relative to deterministic approaches (Dimitrakopoulos, Farrelly and Godoy, 2002; Godoy and Dimitrakopoulos, 2004; Dimitrakopoulos, 2011). Simulated block grades are therefore the primary geological input to the stochastic mine planning process.

Conditional simulation can be carried out by several established methods, including the turning-bands method, sequential indicator simulation, and Sequential Gaussian Simulation (SGS) (Goovaerts, 1997; Deutsch and Journel, 1998; Chilès and Delfiner, 2012). Among these, SGS is one of the most widely used techniques in mineral resource modelling owing to its conceptual simplicity, its faithful reproduction of the modelled covariance, and its compatibility with standard

kriging workflows (Isaaks, 1990; Deutsch and Journel, 1998). In SGS, the sample grades are first transformed to standard normal scores with marginal distribution $N(0,1)$. Grid nodes are then visited along a random path, and at each unsampled location u a value is drawn from a Gaussian conditional distribution whose mean and variance are supplied by ordinary kriging of the surrounding data and previously simulated nodes.

$$Y(u) \sim N(\mu_{OK}(u), \sigma^2_{OK}(u)) \quad (1)$$

Where $\mu_{OK}(u)$ is the ordinary kriging estimate and $\sigma^2_{OK}(u)$ the kriging variance. Each simulated value is added to the conditioning set before the next node is visited, and the completed field is back transformed to grade units. Because the conditional mean and variance derive directly from the kriging system, the quality of an SGS realisation depends strongly on the kriging configuration and the fitted variogram model that defines spatial continuity (Goovaerts, 1997; Rossi and Deutsch, 2014).

Conventional SGS assumes second-order stationarity: a single variogram model and a fixed search ellipsoid are applied across the entire domain. In structurally complex deposits affected by folding, faulting, or curvilinear mineralisation controls, the direction and magnitude of grade continuity vary spatially, and this assumption is violated (Journel and Rossi, 1989; Boisvert, Manchuk and Deutsch, 2009). Locally varying anisotropy (LVA) approaches address this issue by allowing the anisotropy orientation and ranges to change from one location to another. Boisvert, Manchuk and Deutsch (2009) and Boisvert and Deutsch (2011) introduced kriging and SGS algorithms in which continuity is governed by an LVA field through non-Euclidean distances, while Machuca-Mory and Deutsch (2013) proposed a distance-weighted, location-dependent formulation of the statistics used in non-stationary modelling. The performance of ordinary kriging under several local-anisotropy formulations has been examined by Boisvert and Deutsch (2011), Machuca-Mory and Deutsch (2013), and Aydar and Dagdelen (2025), and its extension to SGS has been studied by Aydar (2026).

The uncertainty of a simulated block is most commonly quantified by the coefficient of variation computed across the set of L realisations. For a block at location u with simulated grades $z_l(u)$, $l = 1, \ldots, L$, the E-type mean $m(u)$ and the between-realisation standard deviation $s(u)$ are defined by Equation 2, and the coefficient of variation (CV) by Equation 3.

$$m(u) = (1/L) \sum z_l(u), \; s(u) = \sqrt{[(1/L) \sum (z_l(u) - m(u))^2]} \quad (2)$$

$$CV(u) = s(u) / m(u) \quad (3)$$

Aydar (2026) reported that incorporating locally varying anisotropy into SGS reduces block-scale uncertainty, expressed as a lower median CV, relative to conventional stationary SGS. A reduction in simulated uncertainty at the block scale, however, need not translate proportionally into downstream planning outcomes, because the optimiser aggregates block values over space and time subject to capacity, access, and discounting constraints. The downstream effect of Adaptive Search Ellipsoid (ASE)-aided SGS on stochastic mine planning therefore remains unquantified.

This study addresses that gap. Conditional realisations of the McLaughlin gold deposit are generated using both conventional stationary SGS and ASE-aided SGS, and the E-type mean of each realisation suite is used as input to a ten-year open pit production plan solved with the Direct Block Scheduling (DBS) approach, in which extraction and processing decisions are assigned directly at the block level rather than through predefined pushbacks (Menabde *et al,* 2007). By comparing the two simulation approaches in terms of block-scale uncertainty, grade-tonnage behaviour, annual mill-feed grade, and discounted cash flow, the study quantifies whether the reduction in block-scale uncertainty achieved by ASE-aided SGS propagates into measurable differences in long-term production planning decisions for a structurally complex precious-metal deposit.

SEQUENTIAL GAUSSIAN SIMULATION

Sequential Gaussian Simulation proceeds in three stages: a normal-score transformation, a sequential kriging-based draw along a random visit path, and a back-transformation to data units (Isaaks, 1990; Deutsch and Journel, 1998). Let $z(u)$ denote the regionalised grade variable at

location u. The forward transform $\varphi(\cdot)$, constructed from the empirical cumulative distribution of the samples, maps the data to normal scores $y(u) = \varphi(z(u))$ with marginal distribution $N(0,1)$.

Grid nodes are then visited in a random order. At each unsampled node u_i, the surrounding data $\{y(u_\alpha),\ \alpha = 1, \ldots, n(u_i)\}$, comprising both original samples and previously simulated nodes falling within the search ellipsoid, supply the conditional moments through the ordinary-kriging system given in Equation 4.

$$\mu_{OK}(u_i) = \sum \lambda_\alpha\, y(u_\alpha),\ \sigma^2_{OK}(u_i) = C(0) - \sum \lambda_\alpha\, C(u_\alpha - u_i) - \mu_L \tag{4}$$

where the weights λ_α solve the standard ordinary-kriging linear system, μ_L is the Lagrange multiplier enforcing $\sum \lambda_\alpha = 1$, and $C(h)$ is the covariance corresponding to the fitted variogram model $\gamma(h)$ (Goovaerts, 1997). A single value at u_i is drawn from the conditional Gaussian distribution as given in Equation 5.

$$y_l(u_i) = \mu_{OK}(u_i) + \sigma_{OK}(u_i)\,\varepsilon,\ \varepsilon \sim N(0,1) \tag{5}$$

and added to the conditioning set before the next node is visited. Once all nodes are simulated, the field is back-transformed by $z_l(u) = \varphi^{-1}(y_l(u))$. Repeating the procedure with L independent random paths and seeds produces a set of equiprobable realisations that honour the data, reproduce the modelled covariance, and provide a Monte Carlo characterisation of grade uncertainty. Figure 1 illustrates ten SGS realisations at elevation 1610 ft of the McLaughlin gold deposit.

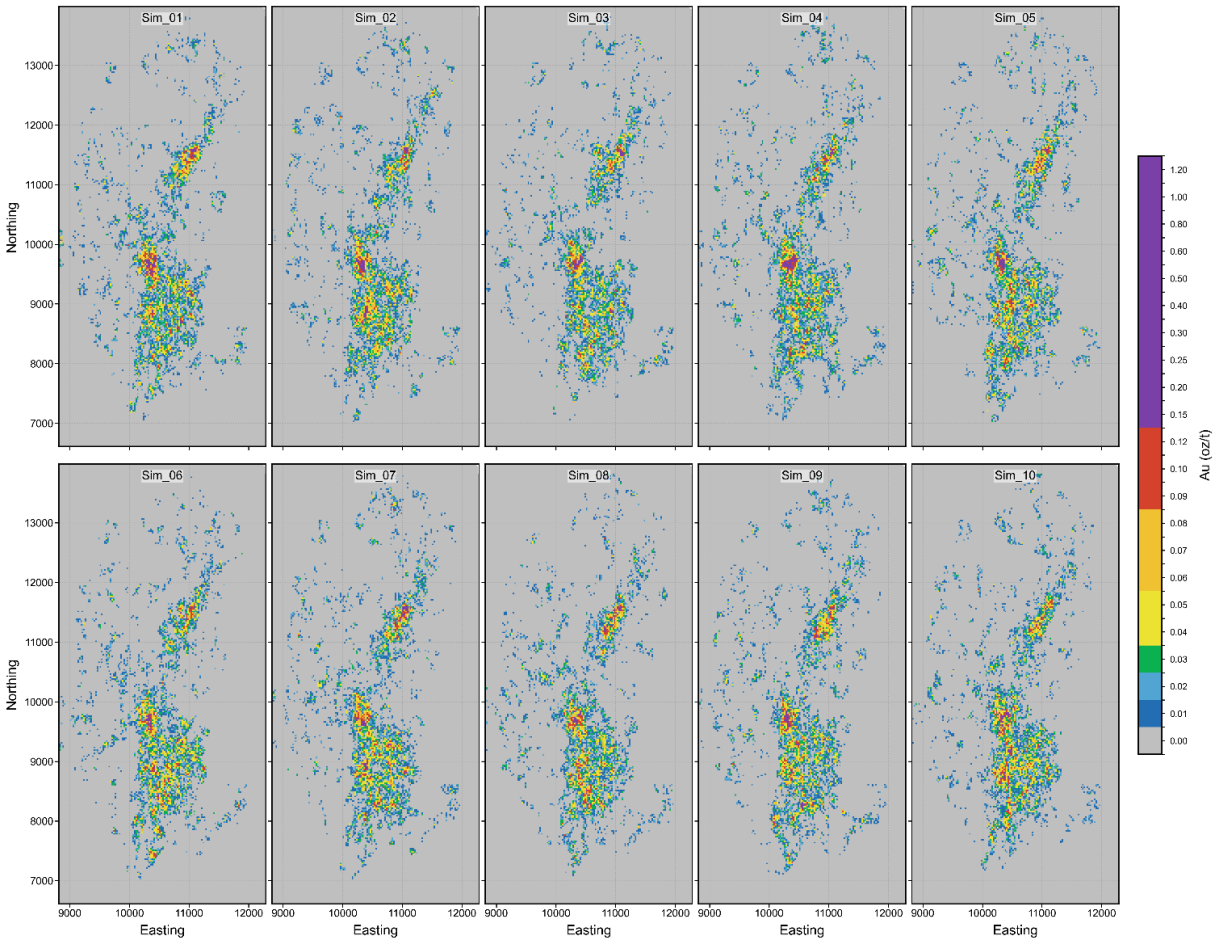


FIG 1 – Ten conventional SGS realisations of Au grade (oz/t) at elevation 1610 ft, McLaughlin gold deposit.

ADAPTIVE SEARCH ELLIPSOID SGS

Conventional SGS evaluates the kriging system at every node using a single, domain-wide variogram model $\gamma(h)$. The Adaptive Search Ellipsoid SGS (ASE-SGS) retains the sequential framework but replaces this stationary model with a location-dependent one (Aydar, 2026). For each block u a local variogram $\gamma_u(h)$ is supplied by a separate, prior fitting step (Aydar and

Dagdelen, 2025), which specifies the local anisotropy axes (azimuth, dip, and plunge; see Figure 2) and three range parameters $a_{maj}(u)$, $a_{min}(u)$, and $a_{vert}(u)$.

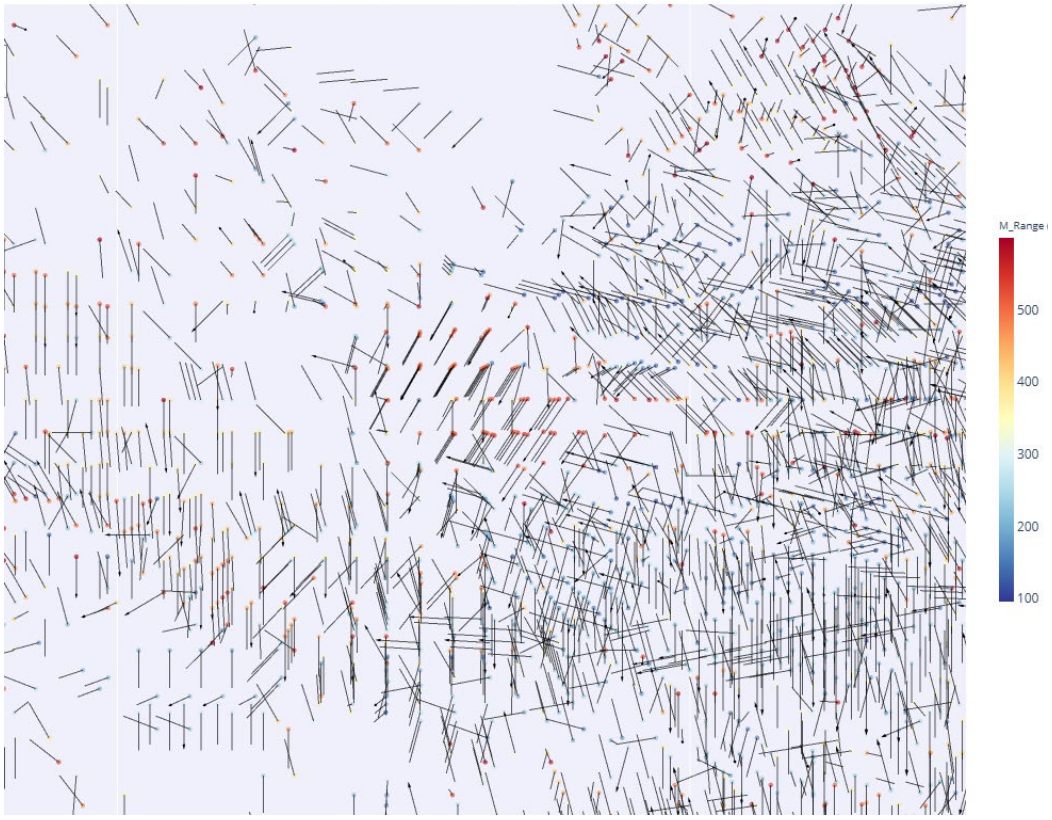


FIG 2 – Three-dimensional representation of locally fitted anisotropy directions and major-axis ranges (colour scale, ft) at the McLaughlin gold deposit.

The local anisotropy is encoded in a rotation matrix $\boldsymbol{R}(u)$ and a range matrix $\boldsymbol{A}_u = \mathrm{diag}(a_{maj}, a_{min}, a_{vert})$, which together define a reduced ellipsoidal distance as given in Equation 6.

$$h_{red}(\boldsymbol{h}; u) = \sqrt{(\boldsymbol{h}^\top \boldsymbol{R}(u)\, \boldsymbol{A}_u^{-2}\, \boldsymbol{R}^\top(u)\, \boldsymbol{h})} \tag{6}$$

With this reduced distance, the local covariance is defined by Equation 7, where $C(\cdot)$ is the standard covariance corresponding to $\gamma(h)$:

$$C_u(\boldsymbol{h}) = C(h_{red}(\boldsymbol{h}; u)) \tag{7}$$

The kriging system at each node is solved using the local covariance in place of the global one; the random visit order is prioritised for the blocks selected for local anisotropy; the normal-score transform, conditioning rule, and back-transform are otherwise unchanged from standard SGS. Figure 2 shows a 3D representation of the deposit in which circles mark blocks with locally fitted anisotropies, arrows indicate the major-axis direction, and the colour scale encodes the major-axis range.

Two design choices keep the algorithm tractable for industrial block models. First, the simulation operates in Euclidean space; only the orientation and scaling of the search ellipsoid vary across the grid, rather than the metric of the space itself. This formulation avoids the shortest-path computations required by the fully non-Euclidean approach of Boisvert, Manchuk and Deutsch (2009), at the cost of an inability to follow strongly curvilinear features that bend across more than a few correlation ranges. Second, blocks for which the local variogram fit is judged unreliable, based on a residual-error threshold applied during the prior fitting step, revert to the global variogram model, which ensures well-behaved kriging weights in data-sparse regions. The resulting realisations honour the histogram and global covariance of conventional SGS, while allowing local continuity to follow the geological grain implied by the data. Figure 3 shows ten ASE-SGS realisations at elevation 1610 ft for direct visual comparison with the conventional realisations of Figure 1.

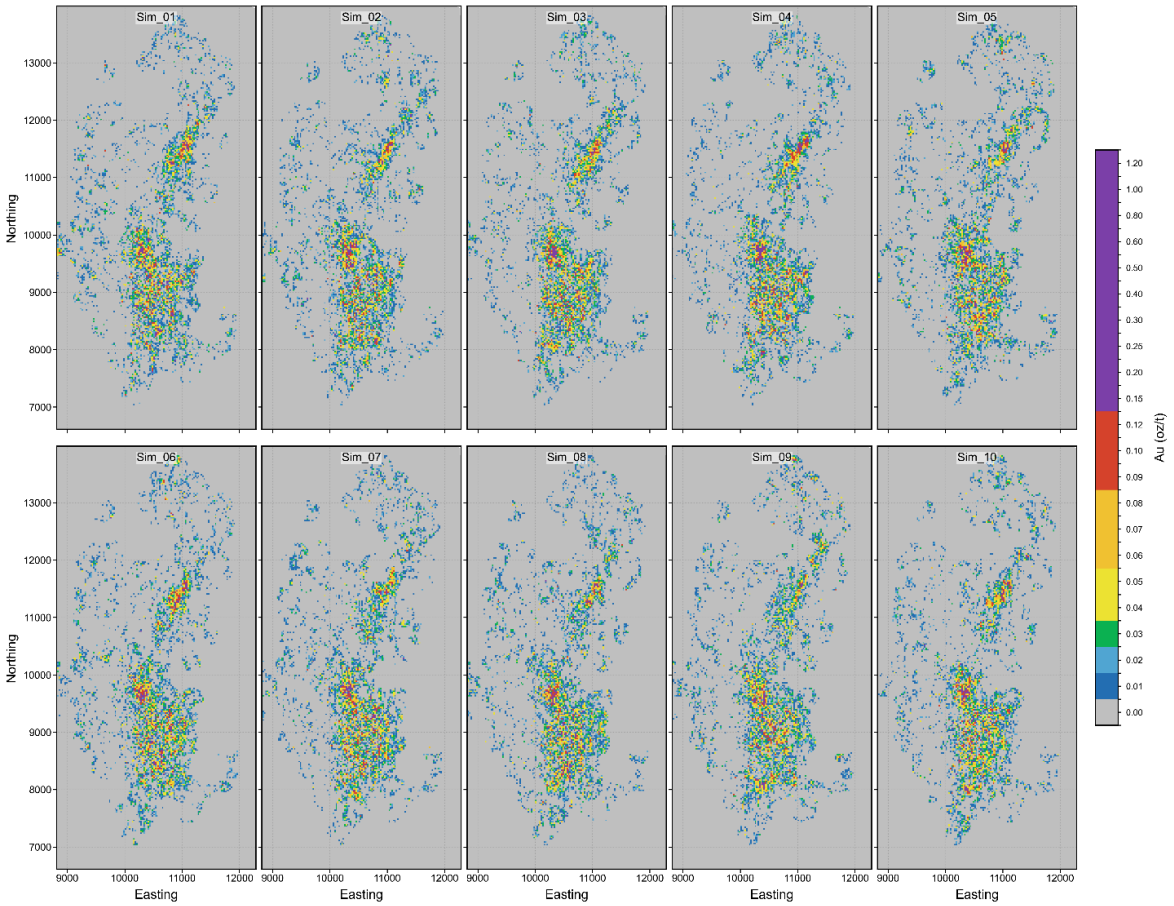


FIG 3 – Ten ASE-SGS realisations of Au grade (oz/t) at elevation 1610 ft, McLaughlin gold deposit.

DIRECT BLOCK SCHEDULING

Conventional open pit production planning is a two-stage exercise: an ultimate pit limit is first identified by the Lerchs–Grossmann graph algorithm or a max-flow equivalent (Lerchs and Grossmann, 1965), and the resulting volume is then subdivided into nested pushbacks that are sequenced into periods using ordering heuristics. Direct Block Scheduling (DBS) bypasses this decomposition and assigns extraction and destination decisions directly at the block level by solving a single mixed-integer programme over the full block model (Johnson, 1968; Dagdelen and Kawahata, 2008; Aras, 2018; Aras, Dagdelen and Johnson, 2019). Because DBS retains the temporal interactions between blocks, most importantly the discounting of future cash flows, it provides a natural framework for measuring how geological inputs propagate through to the schedule (Menabde *et al,* 2007; Dimitrakopoulos and Ramazan, 2008).

Mathematical formulation

Let $\boldsymbol{B}$ denote the set of blocks in the model, $\boldsymbol{T} = \{1, \ldots, T\}$ the set of scheduling periods, and $\boldsymbol{D} = \{m, w\}$ the set of destinations, where m denotes the processing plant and w the waste dump. For each block $b \in \boldsymbol{B}$, period $t \in \boldsymbol{T}$, and destination $d \in \boldsymbol{D}$, a binary decision variable $x_{b,t,d}$ is defined that equals 1 if block b is extracted in period t and sent to destination d, and 0 otherwise.

No cut-off grade is supplied to the optimiser as an input. Instead, the destination of each extracted block is itself a decision variable: the optimiser determines, block by block and period by period, whether the block is routed to the processing plant or to the waste dump, according to which assignment maximises the discounted objective under the operational constraints (Johnson, 1968). The destination-dependent block economic values are given in Equation 8.

$$v_{b,m} = T_b\,(z(b)\,r\,p - c_p),\; v_{b,w} = -T_b\,c_d \tag{8}$$

where the symbols are defined in Table 1. The DBS problem maximises the discounted net cash flow over the scheduling horizon, as given in Equation 9.

TABLE 1

Economic and operational parameters used in the DBS optimisation.

Parameter	Symbol	Value
Gold price	p	US$2000/oz
Processing cost (ore)+mining cost	c_p	US$40/t
Mining cost (waste)	c_d	US$3.00/t
Metallurgical recovery	r	0.9
Annual discount rate	i	12.5%
Mill capacity	Mmill	2.25 *Mt/a
Mining rate	–	Unconstrained
Waste dump capacity	–	Unconstrained
Maximum overall pit slope	–	45°
Scheduling horizon	|T|	10 years
Au grade of block b in the E-type input model	z(b)	-
Block tonnage	T_b	-

$$\max \text{NPV} = \sum_t (1 / (1 + i)^{t-1}) \sum_b \sum_d v_{b,d}\, x_{b,t,d} \tag{9}$$

with i the annual discount rate. The maximisation in Equation 9 is subject to four families of constraints.

Reserve constraint. Each block can be extracted at most once across the horizon (Equation 10):

$$\sum_{t \in T} \sum_{d \in D} x_{b,t,d} \le 1,\ \forall b \in \boldsymbol{B} \tag{10}$$

Slope (precedence) constraint. For every block b and every predecessor block $b' \in P(b)$ implied by the maximum overall pit slope, b' must be removed in or before the same period as b (Equation 11):

$$\sum_{\tau=1}^{t} \sum_d x_{b,\tau,d} \le \sum_{\tau=1}^{t} \sum_d x_{b',\tau,d},\ \forall b \in \boldsymbol{B},\ \forall b' \in P(b),\ \forall t \in \boldsymbol{T} \tag{11}$$

The predecessor set $P(b)$ is generated geometrically from the chosen overall slope angle using a nine-point cone template.

Processing capacity constraint. Total ore tonnage routed to the plant in any period cannot exceed the mill throughput M_{mill} (Equation 12):

$$\sum_{b \in B} T_b\, x_{b,t,m} \le M_{mill},\ \forall t \in \boldsymbol{T} \tag{12}$$

Mining and dump capacities. Both are left unbounded in the present formulation, which corresponds to the assumption that the operation is processing-limited rather than equipment or storage-limited; the slope and processing constraints alone determine the schedule. Figure 4 summarises the resulting two-destination flow.

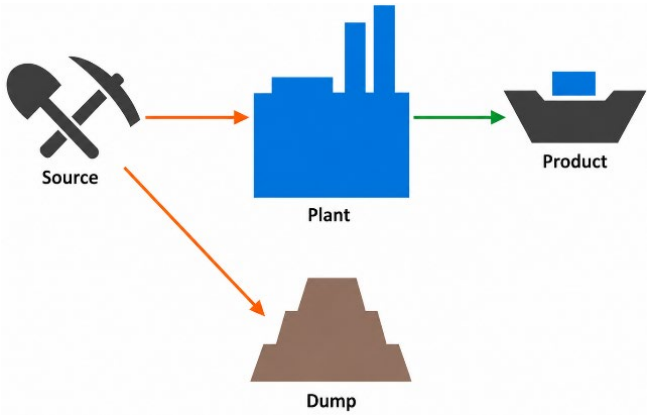


FIG 4 – Flow chart of block destinations adopted in the case study; the routing of each block to mill feed or to the waste dump is assigned by the DBS optimiser rather than by a pre-set cut-off grade.

Because block destinations are decision variables rather than the product of a pre-set cut-off grade, the marginal cut-off grade is an outcome of the optimisation rather than an input to it. A block is routed to the plant only when its plant value exceeds its dump value, that is, when $z(b) > (c_p - c_d)/(p\ r)$; whenever the processing capacity constraint binds, the opportunity cost of mill capacity raises this implied cut-off further, period by period. Cut-off determination in the DBS formulation is therefore made endogenously, block by block, under the reserve, slope, and capacity constraints.

In a fully stochastic implementation, the mixed-integer programme defined by Equations 8 to 12 would be solved independently on each conditional realisation, producing a distribution of schedules and NPV outcomes (Dimitrakopoulos and Ramazan, 2008; Ramazan and Dimitrakopoulos, 2013). In the present study, in order to isolate the systematic effect of the simulation method on planned value rather than the realisation-to-realisation variability, the DBS problem is solved on the E-type mean of each simulation suite. This deterministic comparison provides a direct, like-for-like assessment of conventional SGS against ASE-SGS as inputs to long-term scheduling. Full propagation of all realisations through the optimiser is identified as a key next step and is discussed under Limitations.

Economic and operational parameters

Table 1 lists the economic and operational parameters adopted for the comparative study. The values represent a generic open pit gold operation rather than the historical economics of McLaughlin, and the same set is applied uniformly to both the conventional-SGS and ASE-SGS E-type means, so that any difference in schedule outcomes is attributable to the simulation method rather than to differences in input economics. Because the DBS formulation assigns block destinations endogenously, no cut-off grade is supplied to the optimisation as an input: the routing of each block to the plant or to the waste dump is decided by the optimiser, and the marginal cut-off grade is an outcome of the schedule rather than a parameter of it. Under the parameters of Table 1 and the destination values of Equation 8, the implied plant-versus-dump breakeven grade is $(c_p)/(p^*r) \approx 0.02$ oz/t Au, and the discounted cash flows apply the metallurgical recovery of 0.9 to the scheduled grades.

CASE STUDY – THE MCLAUGHLIN GOLD DEPOSIT

Geological setting

The McLaughlin deposit is a Plio-Pleistocene hot-spring type gold deposit located in the Knoxville mining district of the northern Coast Ranges of California, USA, at the contact between the Middle Jurassic Coast Range ophiolite and the overlying Great Valley sequence (Sherlock *et al,* 1995). Mineralisation is centred on a sheeted vein complex developed along the moderately north-east-dipping Stoney Creek fault, with gold occurring almost entirely as submicron particles in quartz veins and the surrounding silica–adularia–sericite alteration envelope (Tosdal *et al,* 1993; Sherlock *et al,* 1995). Operated by Homestake Mining Company between 1985 and 2002, the mine

recovered approximately 3.4 million ounces (Moz) of gold (Sherlock *et al,* 1995). The combination of strong vertical continuity along the vein complex, a steeply concentrated high-grade core, and a long, low-grade halo makes McLaughlin a useful test case for assessing the effect of non-stationary modelling assumptions on stochastic mine planning.

Drill hole data and block model

The data set used in this study consists of 10 197 composite assay records from the Homestake exploration drilling campaigns. Each composite carries an east, north, elevation position (X, Y, Z) in feet, downhole interval bounds, and an Au grade z in troy ounces per short ton. Of these composites, 7445 (73 per cent) record non-zero gold. The marginal distribution is strongly positively skewed: the Q1 (first quartile) is 0 oz/t, median Au grade is 0.005 oz/t, Q3 is 0.03 oz/t, the 99.5th percentile is 0.549 oz/t, the 99.9th percentile is 1.232 oz/t, and the maximum is 4.37 oz/t. This long-tailed distribution motivated the use of a normal-score transformation, with an upper-tail hyperbolic extrapolation during back-transformation to prevent the generation of unrealistically extreme draws.

The estimation grid is a regular block model of 140 × 296 × 68 = 2 817 920 blocks, with block dimensions $\Delta x = \Delta y = 25$ ft and $\Delta z = 20$ ft. Block centroids span $X \in [8812.5, 12\,287.5]$ ft east, $Y \in [6612.5, 13\,987.5]$ ft north, and $Z \in [910, 2250]$ ft elevation.

Variogram modelling

A global experimental variogram was computed on the normal-score values and modelled as a single exponential structure with unit sill, a 0.15 nugget, and effective ranges of a_{maj} = 1103 ft, a_{min} = 646 ft and a_{vert} = 356 ft. The total fit residual, expressed as the sum of the per-direction RMSE, is 0.25. The major, semi-major, and minor continuity axes have azimuths and dips of 090°/+75°, 000°/000°, and 090°/−15°, respectively. This stationary model is the variogram used in the conventional SGS run.

For the ASE-SGS run, local directional variograms were fitted on a per-block basis over the same normal-score data using the moving-disc approach of Aydar and Dagdelen (2025), with the search radius set to half the global major-axis range. Each fit returns a local rotation matrix R(u), a local range matrix A(u), and a residual error metric expressed as the per-direction RMSE. Blocks whose total fit RMSE exceeds the tolerance threshold are flagged as unreliable and are assigned the global variogram, so the ASE field is locally adaptive where the data support it and stationary elsewhere. The same 10 197 composites condition both simulations, and both use the same random visit path within each realisation seed; the only operational difference between the two runs is the variogram supplied to the kriging system at each block.

Of the 2 146 196 blocks in the model, 110 665 (5.16 per cent) satisfied the tolerance threshold and were assigned a local variogram. Each local fit was selected as the lowest-RMSE candidate among the three commonly used parametric models, spherical, exponential and Gaussian. The exponential model was preferred in 83 989 blocks (75.89 per cent), the spherical in 23 408 (21.15 per cent), and the Gaussian in 3268 (2.95 per cent). Figure 5 shows the distributions of the fitted ranges, RMSE, and azimuths for the major, semi-major, and minor axes of the locally adaptive blocks. Summary statistics are reported in Table 2.

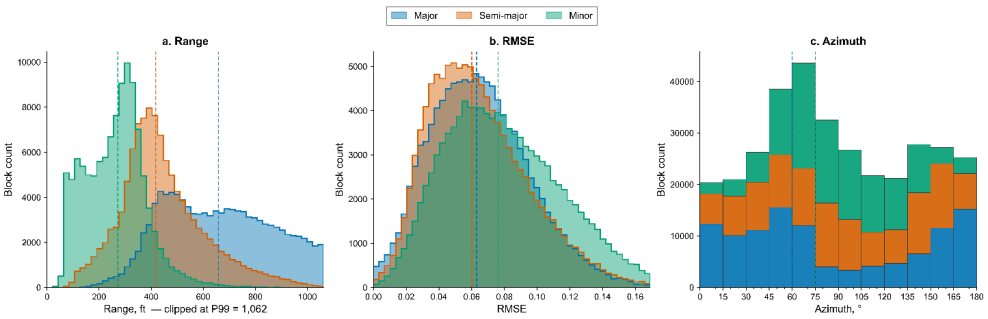


FIG 5 – Histograms of local variogram fits defining the ASE parameters: (a) Range, (b) RMSE of local fits, (c) Azimuth of local variogram fits.

TABLE 2

Summary statistics of local variogram fit parameters.

Axis	Variable	Unit	min	max	mean	median	std	P25	P75
Major	Range	ft	105.6	1103.2	669.6	657.6	214.9	489.7	838.2
Major	RMSE	-	~0	0.21	0.065	0.063	0.030	0.043	0.085
Major	Azimuth	°	0	165	78.1	60	57.8	30	135
Semi-major	Range	ft	63.8	1095.9	445	417.6	163.7	340.6	529.3
Semi-major	RMSE	-	~0	0.21	0.064	0.060	0.031	0.041	0.083
Semi-major	Azimuth	°	0	165	85	75	49.4	45	135
Minor	Range	ft	30.5	993.5	264.1	271.7	115.1	175	334.4
Minor	RMSE	-	~0	0.23	0.080	0.076	0.036	0.053	0.104
Minor	Azimuth	°	0	165	81.8	75	37.6	60	105

Block-scale comparison of SGS and ASE-SGS

For each simulation method, ten realisations were generated using the same composite data and the same random visit path. The ten realisations of each method were then arithmetically averaged to obtain an E-type block grade. The resulting E-type grade maps at elevation 1610 ft are shown in Figure 6 for SGS and ASE-SGS, together with a map identifying the blocks at which the ASE algorithm was activated in the local kriging system.

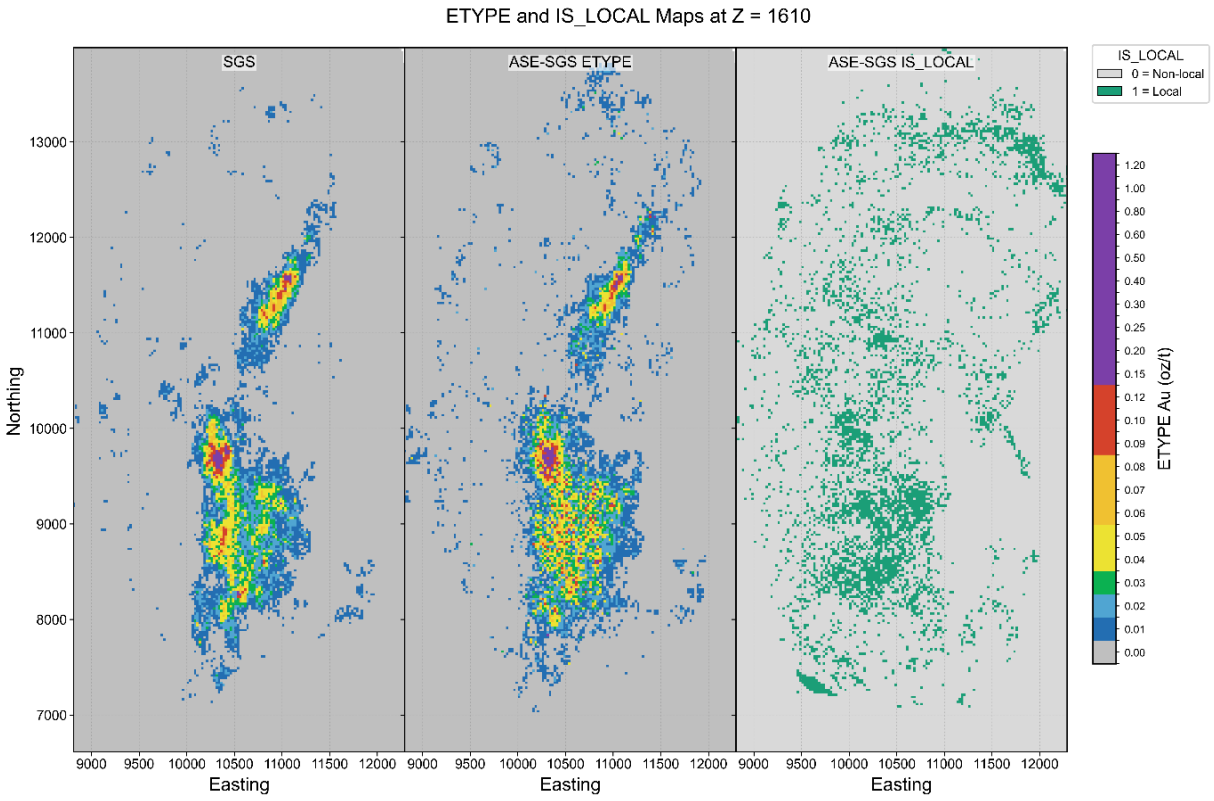


FIG 6 – E-type mean Au grade (oz/t) at elevation 1610 ft for conventional SGS (left) and ASE-SGS (centre); the right-hand map identifies the blocks at which the ASE algorithm was used in the local kriging system.

Swath plots of the two E-type means against composite data at five elevations are shown in Figure 7 for east–west and north–south transects. Where the local composite data indicate elevated grades, conventional SGS systematically underestimates and over smooths the swath

profile. ASE-SGS recovers a closer match to the composite trend in these zones, while tracking the conventional-SGS profile in the more uniformly mineralised intervals. This behaviour reflects the design of ASE-SGS, which uses the locally fitted variogram only where the local fit is judged sufficiently reliable and otherwise reverts to the global model.

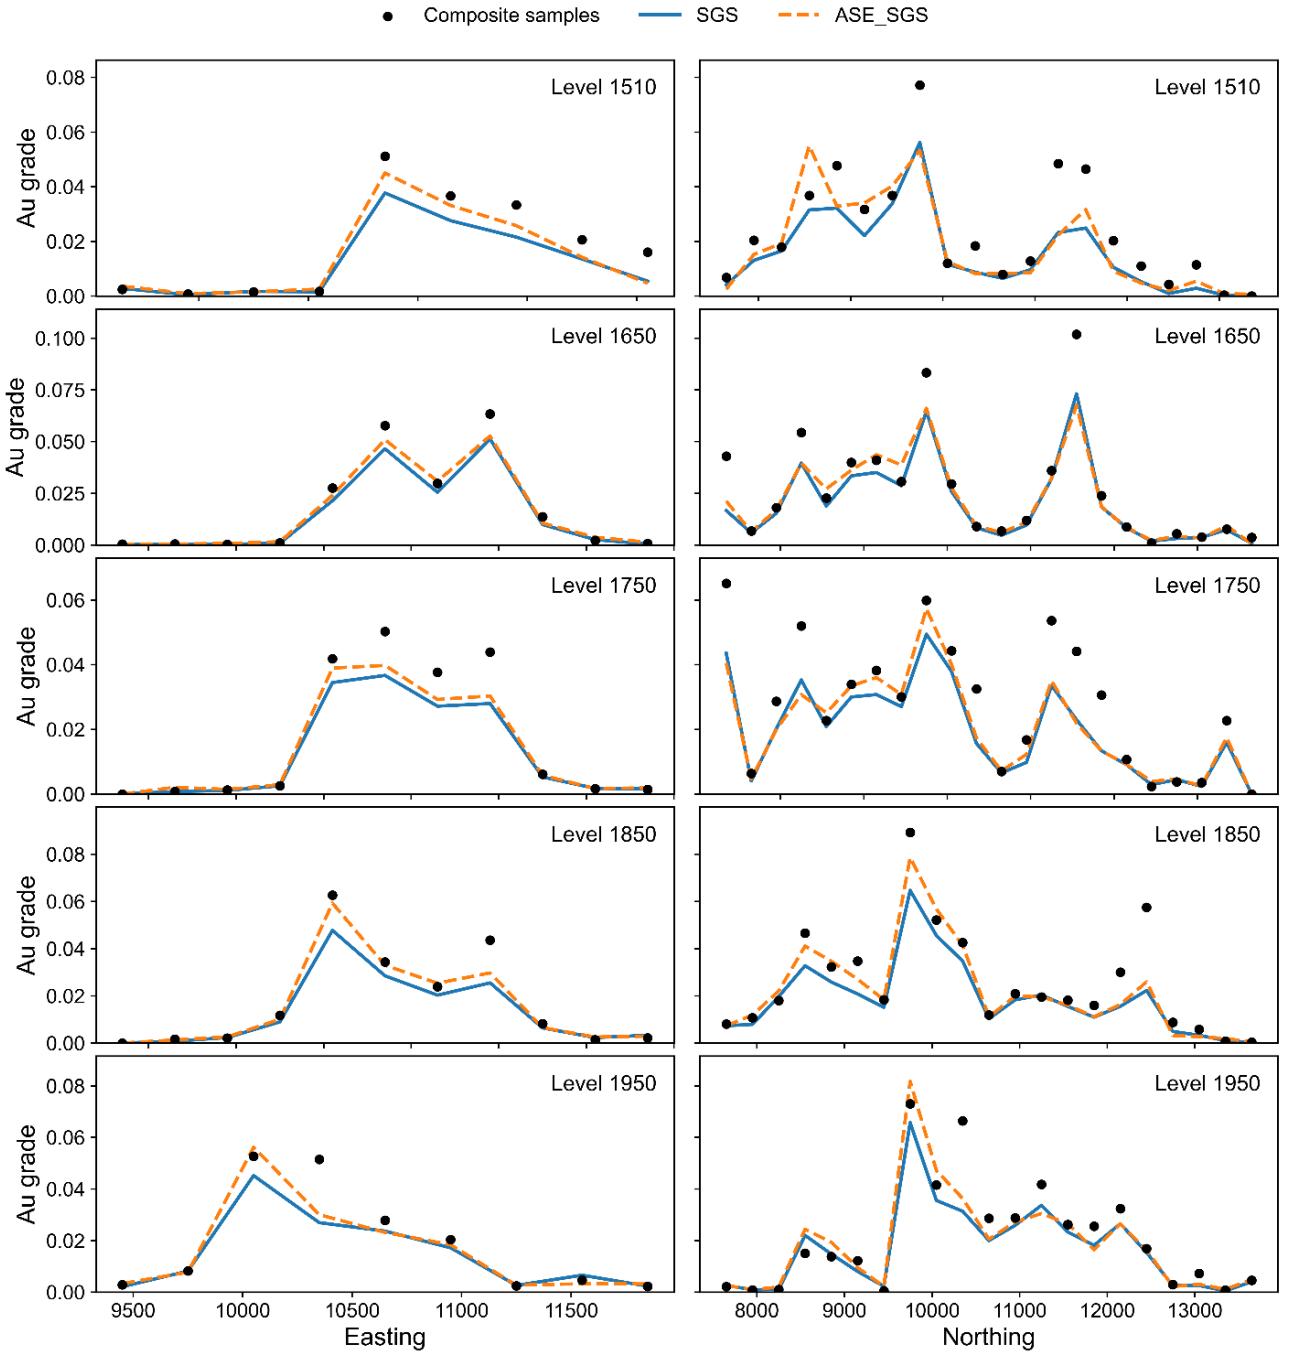
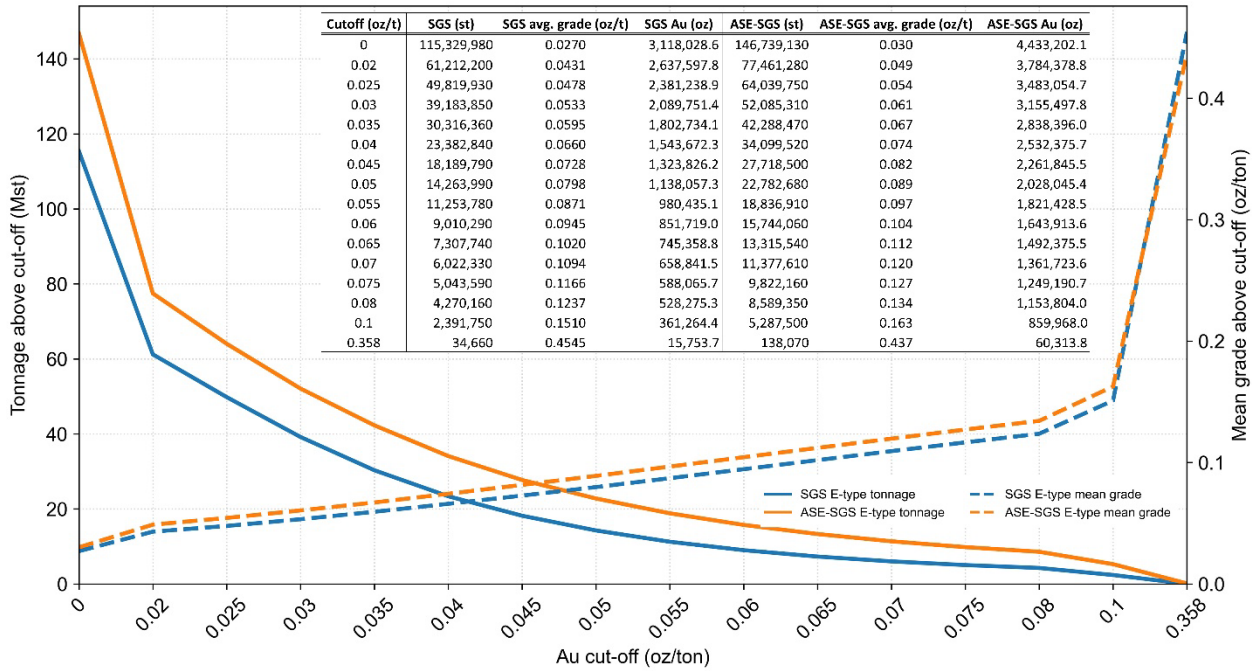


FIG 7 – Swath plots of the E-type mean of conventional SGS and ASE-SGS against composite Au grade in the east–west (left) and north–south (right) directions at five reference elevations.

Grade-tonnage relationships within the ultimate pit limits derived from the two E-types means, using the parameters of Table 1, are compared in Figure 8. For every cut-off grade above the marginal threshold, ASE-SGS recovers a higher tonnage of ore at a higher average grade than conventional SGS. The difference grows with increasing cut-off, consistent with the visual evidence that ASE-SGS better preserves the high-grade tails of the deposit that conventional SGS tends to smear into adjacent low-grade blocks.



Cutoff (oz/t)	SGS (st)	SGS avg. grade (oz/t)	SGS Au (oz)	ASE-SGS (st)	ASE-SGS avg. grade (oz/t)	ASE-SGS Au (oz)
0	115,329,980	0.0270	3,118,028.6	146,739,130	0.030	4,433,202.1
0.02	61,212,200	0.0431	2,637,597.8	77,461,280	0.049	3,784,378.8
0.025	49,819,930	0.0478	2,381,238.9	64,039,750	0.054	3,483,054.7
0.03	39,183,850	0.0533	2,089,751.4	52,085,310	0.061	3,155,497.8
0.035	30,316,360	0.0595	1,802,734.1	42,288,470	0.067	2,838,396.0
0.04	23,382,840	0.0660	1,543,672.3	34,099,520	0.074	2,532,375.7
0.045	18,189,790	0.0728	1,323,826.2	27,718,500	0.082	2,261,845.5
0.05	14,263,990	0.0798	1,138,057.3	22,782,680	0.089	2,028,045.4
0.055	11,253,780	0.0871	980,435.1	18,836,910	0.097	1,821,428.5
0.06	9,010,290	0.0945	851,719.0	15,744,060	0.104	1,643,913.6
0.065	7,307,740	0.1020	745,358.8	13,315,540	0.112	1,492,375.5
0.07	6,022,330	0.1094	658,841.5	11,377,610	0.120	1,361,723.6
0.075	5,043,590	0.1166	588,065.7	9,822,160	0.127	1,249,190.7
0.08	4,270,160	0.1237	528,275.3	8,589,350	0.134	1,153,804.0
0.1	2,391,750	0.1510	361,264.4	5,287,500	0.163	859,968.0
0.358	34,660	0.4545	15,753.7	138,070	0.437	60,313.8

FIG 8 – Grade-tonnage curve and reserve of the ultimate pit at breakeven cut-off grade using the E-type means of conventional SGS and ASE-SGS.

Block-scale uncertainty is summarised in Figures 9 and 10, which show the histogram of the coefficient of variation computed across the ten realisations of each method using Equations 2 and 3. The ASE-SGS distribution shifts toward lower values, and its median CV is lower than that of conventional SGS. The reduction in CV is consistent with the interpretation that allowing local continuity to follow the data, rather than imposing a single domain-wide variogram, decreases the spread of plausible block grades while preserving the global histogram.

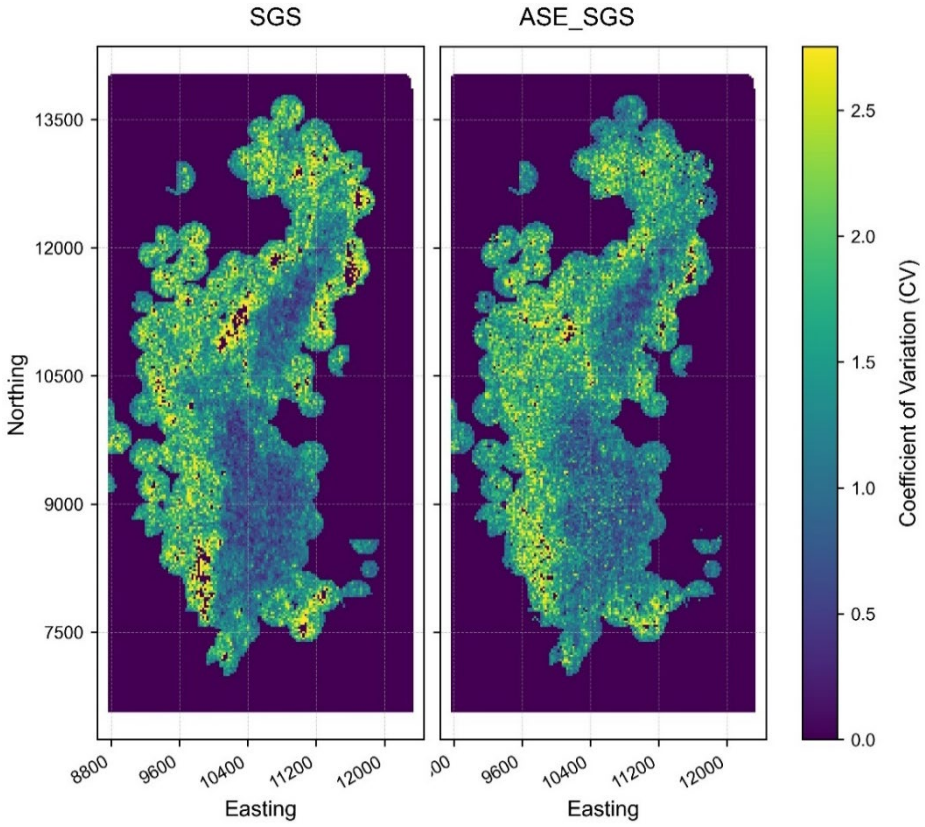


FIG 9 – Coefficient of variation maps across ten realisations for conventional SGS and ASE-SGS at 1610 level.

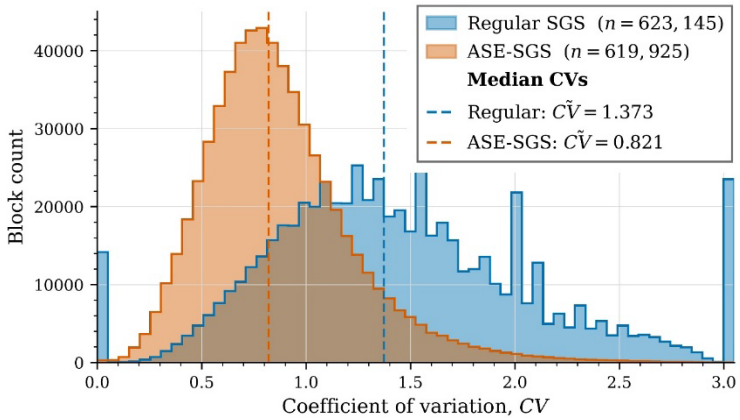


FIG 10 – Coefficient of variation histograms across ten realisations for conventional SGS and ASE-SGS.

Production scheduling outcomes

Production scheduling was performed using the Direct Block Scheduling tool of Deswik.GO with the parameters listed in Table 1 and a two-destination scenario (plant or waste dump), with each block's destination assigned by the optimiser rather than by a pre-set cut-off grade. The E-type means of the SGS and ASE-SGS realisation suites were each used as the deterministic grade input. A ten-year extraction horizon was adopted with a 12.5 per cent annual discount rate, and the optimiser maximised the discounted cash flow subject to the slope and processing constraints in Equations 10 to 12. Plan-view and cross-section visualisations of the two ten-year schedules are shown in Figure 11.

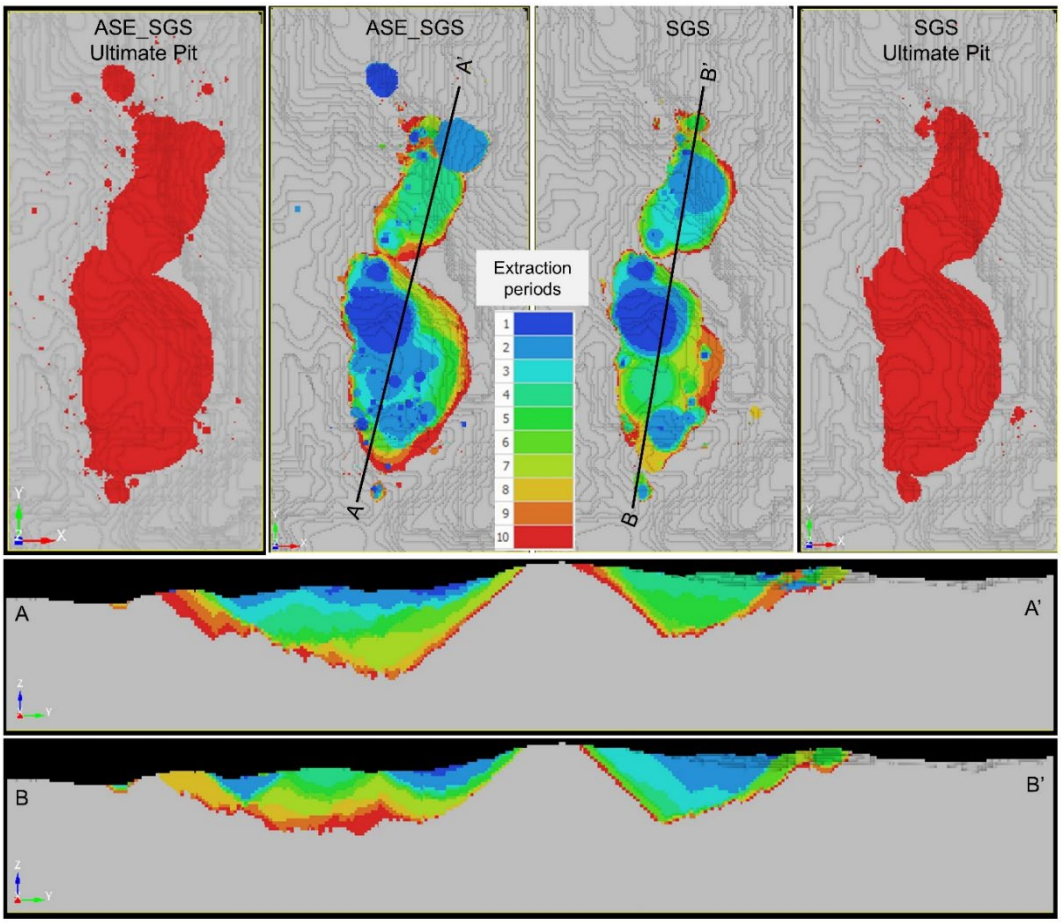


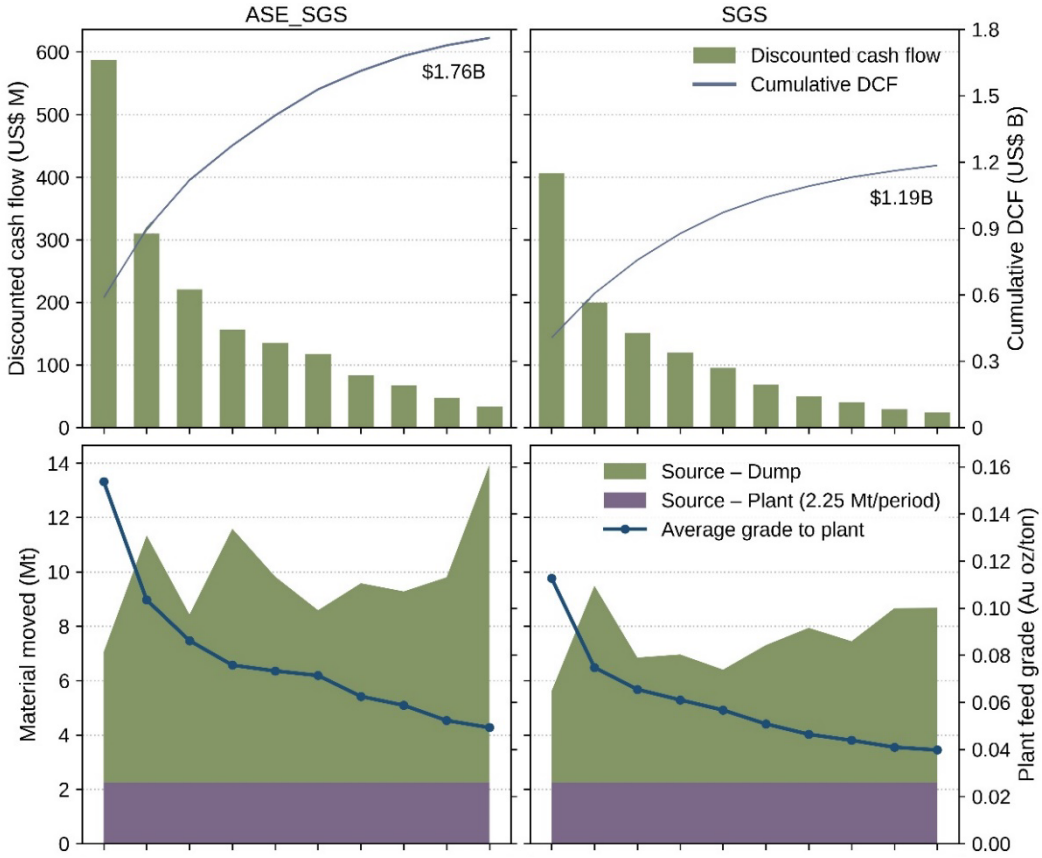
FIG 11 – Ultimate Pits and production period assignments (plan views, top; cross-sections, bottom) for the ten-year extraction periods optimised on the E-type means of conventional SGS (right) and ASE-SGS (left).

Year-by-year tonnages, plant-feed grades, stripping ratios and discounted cash flows for the two schedules are summarised in Table 3 and Figure 12; all tonnages are reported in short tons (st), consistent with the composite database. Because the 2.25 Mst/a plant capacity is binding in every period of both runs, the ore tonnage routed to the plant is identical at 2.25 Mst/a, totalling 22.50 Mst over the ten-year horizon (Figure 12, middle row). The two schedules diverge instead on plant-feed grade, on stripping ratio and on the timing of cash flow generation.

TABLE 3

Year-by-year schedule outcomes for the ten-year extraction plans optimised on the E-type means of conventional SGS and ASE-SGS. Ore tonnage is the material routed to the plant; waste tonnage is the material routed to the waste dump; the stripping ratio (SR) is waste/ore. Discounted cash flow (DCF) is computed at a 12.5 per cent annual discount rate with year 1 undiscounted. Cash flows apply the metallurgical recovery of 0.9 to contained gold.

Period	Ore (Mst)	Avg grade SGS (oz/t)	Avg grade ASE-SGS (oz/t)	Waste SGS (Mst)	Waste ASE-SGS (Mst)	SR SGS	SR ASE-SGS	DCF SGS (USDM)	DCF ASE-SGS (USDM)
1	2.25	0.113	0.154	3.37	4.82	1.50	2.14	406	587
2	2.25	0.075	0.104	7.24	9.09	3.22	4.04	200	310
3	2.25	0.065	0.086	4.60	6.18	2.04	2.75	151	221
4	2.25	0.061	0.076	4.71	9.34	2.09	4.15	120	157
5	2.25	0.057	0.073	4.15	7.57	1.85	3.37	95	135
6	2.25	0.051	0.071	5.06	6.34	2.25	2.82	68	118
7	2.25	0.046	0.063	5.69	7.33	2.53	3.26	50	83
8	2.25	0.044	0.059	5.20	7.03	2.31	3.13	40	67
9	2.25	0.041	0.052	6.41	7.55	2.85	3.36	29	48
10	2.25	0.040	0.049	6.42	11.70	2.85	5.20	24	33
LOM	22.50	0.059	0.079	52.85	76.95	2.35	3.42	1190	1760



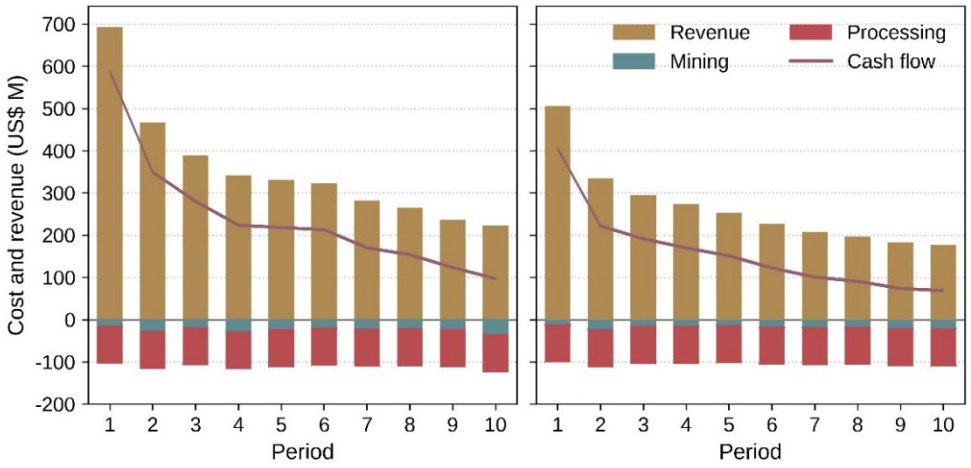


FIG 12 – Discounted cash flow by period (top), material movement to the plant and waste dump with plant-feed grade by period (middle), and revenue and costs by period (bottom) for the ten-year schedules optimised on the E-type means of ASE-SGS (left) and conventional SGS (right).

The ASE-SGS schedule delivers a higher plant-feed grade in every period of the horizon. In year-1 the ASE-SGS feed grade is 0.154 oz/t against 0.113 oz/t for SGS, a relative increase of 36 per cent; the gap persists throughout the schedule, with year-10 feed grades of 0.049 oz/t (ASE-SGS) against 0.040 oz/t (SGS). Both schedules show the declining-grade profile expected under NPV maximisation, in which high-grade material is sequenced early, but the ASE-SGS curve lies above the SGS curve in every period, by between 23 per cent (year 10) and 39 per cent (year 6). Aggregated over the ten-year horizon, the ASE-SGS schedule delivers 1.77 Moz of gold to the plant against 1.33 Moz for SGS, an uplift of 0.44 Moz (+33 per cent) at a life-of-mine average mill-feed grade of 0.079 oz/t versus 0.059 oz/t.

The stripping ratio differs both in magnitude and in temporal sequencing. Over the life-of-mine the conventional schedule strips 52.85 Mst of waste for 22.50 Mst of ore (LOM SR = 2.35), whereas the ASE-SGS schedule strips 76.95 Mst for the same ore tonnage (LOM SR = 3.42). The ASE-SGS schedule therefore moves an additional 24.1 Mst (+32 per cent) of total material to access the higher-grade core that the locally adaptive variogram has resolved into a more contiguous geometry. The two schedules also distribute their stripping differently over time. The conventional schedule incurs its peak stripping in year 2 (7.24 Mst, SR 3.22) and thereafter mines between 4.2 and 6.4 Mst of waste per annum, whereas the ASE-SGS schedule carries a persistently higher stripping load, with pulses in years 2 and 4 (9.09 and 9.34 Mst, SR 4.04 and 4.15) and a final push in year 10 (11.70 Mst, SR 5.20) to reach the high-grade remnants at depth that the optimiser has deferred to the end of the discounted-cash-flow-maximising sequence.

Annual cash flows follow directly from the block values of Equation 8 and the economic parameters of Table 1: a gold price of USD2000 per ounce, a metallurgical recovery of 0.9, a combined mining-and-processing cost of USD40 per ton of ore (c_p) and a mining cost of USD3 per ton of waste (c_d). For each scheduling period t the net cash flow is given by Equation 13 and its discounted value by Equation 14:

$$CF_t = p\, Au_t - c_p\, O_t - c_d\, W_t \tag{13}$$

$$DCF_t = CF_t / (1 + i)^{t-1},\ NPV = \sum_t DCF_t \tag{14}$$

where Au_t denotes the recoverable ounces of gold delivered to the plant in period t, that is, the ounces contained in the mill feed multiplied by the metallurgical recovery r = 0.9, O_t and W_t are the ore and waste tonnages of Table 3, and i = 12.5 per cent is the annual discount rate, with year 1 undiscounted. Because the plant is full in every period, the processing charge is constant at USD90 million per annum in both schedules and the mining charge tracks the waste tonnage; the difference in annual cash flow between the two runs is therefore driven almost entirely by revenue, which in turn follows the plant-feed grade (Figure 12, bottom row).

Applied to the schedules of Table 3, this yields cumulative discounted cash flows of USD1.19 billion for the conventional-SGS schedule and USD1.76 billion for the ASE-SGS schedule, an uplift

of approximately USD0.57 billion, or 48 per cent, in favour of the locally adaptive run (Figure 12, top row). The ASE-SGS discounted cash flow exceeds that of SGS in every period, from USD587 million against USD406 million in year 1 to USD33 million against USD24 million in year 10, and in both schedules the first three years contribute close to two-thirds of the cumulative value (64 per cent in each case), reflecting the combined effect of front-loaded grade and discounting. Because both schedules use identical mining, processing and economic parameters, the entire difference is attributable to the spatial reorganisation of grades introduced by the locally adaptive variogram: the ASE-SGS E-type mean concentrates the contained metal into a more contiguous high-grade core that the optimiser can sequence earlier in the horizon, and the additional 0.44 Moz delivered to the plant more than compensates for the additional 24.1 Mst of waste mined in service of that sequence.

DISCUSSION

Advantages of ASE-SGS for stochastic mine planning

Three features of ASE-SGS account for its improved performance as an input to long-term production scheduling at McLaughlin. First, the use of locally fitted variograms recovers grade continuity that is sub-resolution to a single domain-wide model. In a deposit such as McLaughlin, where mineralisation follows complex and irregular geometries along a moderately dipping fault and develops curvilinear high-grade shoots within an otherwise low-grade halo, a fixed search ellipsoid blends across genuine continuity boundaries. The locally adapted per-block ellipsoids preserve these boundaries, which explains both the improved swath behaviour against composite data and the increased recovered tonnage above the cut-off grade.

Second, the reduction in block-scale CV achieved by ASE-SGS is consistent with the geological interpretation that imposing a single, domain-wide stationary model in a non-stationary deposit inflates simulated variability. When the data themselves contain information about the local direction of continuity, capturing that information at the simulation stage reduces the spread of equiprobable block grades around their conditional mean without changing the global distribution. For downstream planning, lower block-scale CV is desirable when the planning workflow uses an E-type input, because the deterministic schedule is then anchored to a more representative central tendency.

Third, ASE-SGS retains the standard SGS infrastructure: the same normal-score transform, the same conditioning rule, and the same back-transform. The only operational change is the substitution of a local variogram for the global variogram at each kriging system. The method does not require an externally supplied LVA field of the kind used by Boisvert, Manchuk and Deutsch (2009) or Boisvert and Deutsch (2011), and it does not require the shortest-path computations associated with non-Euclidean distance simulation. This makes ASE-SGS comparatively straightforward to integrate into the existing resource estimation workflows used by industry.

Limitations and drawbacks

Several limitations of the present study should be acknowledged when interpreting the reported improvement in scheduled value. The most important is that production scheduling was performed on the E-type mean of each simulation suite, not on the individual realisations. A genuine stochastic mine planning analysis would solve the DBS problem on each realisation and derive a distribution of schedules and NPV outcomes, from which downside risk and chance of meeting production targets could be quantified (Dimitrakopoulos and Ramazan, 2008; Ramazan and Dimitrakopoulos, 2013). The deterministic comparison presented here isolates the systematic effect of the simulation method on planned value but cannot, by construction, quantify the associated risk reduction. Propagating the full set of realisations through the optimiser is the natural next step and is the focus of ongoing work.

Second, ASE-SGS operates in Euclidean space and varies only the orientation and scale of the search ellipsoid, not the underlying metric. The method therefore cannot follow strongly curvilinear features that bend across more than a few correlation ranges. In deposits whose mineralisation describes tightly folded or strongly bending structures, a fully non-Euclidean LVA formulation may

be required to capture continuity accurately. The choice between ASE-SGS and a non-Euclidean alternative is therefore a deposit-specific judgement that depends on the geological style of mineralisation and on the availability of structural information.

Third, the quality of the ASE-SGS output depends on the quality of the prior local variogram fitting step. Blocks for which the local fit is unreliable revert to the global variogram, which prevents pathological kriging behaviour but means that ASE-SGS reduces to conventional SGS in data-sparse regions. The setting of the residual-error threshold used to identify unreliable fits is itself an analyst decision; sensitivity of the downstream schedule to this threshold has not been quantified in the present study and warrants systematic investigation.

Fourth, ASE-SGS is more computationally expensive than conventional SGS, because the local variogram fitting step must be performed once per block prior to simulation and because the kriging weights are recomputed against a location-dependent covariance. For very large block models, this overhead can be substantial and may motivate parallel implementations or the use of approximate local fits.

Fifth, the simulations were conditioned on the full composite database available for the deposit. For prospective application to an unmined deposit, only exploration drilling would be available at the time of long-term planning; with sparser conditioning data, the proportion of blocks passing the local variogram fitting threshold (already only 5.16 per cent of the model in the present study) would be expected to decrease further, leaving a larger, unquantified share of the model to default to the global variogram. In addition, withholding later-stage grade-control drilling from the conditioning set would provide an independent data set against which the realisations of both simulation methods could be validated; such an out-of-sample comparison is recommended as future work.

Sixth, the entire deposit was simulated as a single estimation domain with one global normal-score transformation and one domain-wide variogram model. Industry-standard practice would subdivide the deposit into geological domains, for example by lithology, alteration, or structural position, each with its own refined variogram anisotropy and a tighter grade distribution. Domain-by-domain simulation would be expected to narrow the reported differences between conventional SGS and ASE-SGS, because part of the non-stationarity captured by the locally adaptive ellipsoids would instead be absorbed by the domain boundaries, and would likely further reduce the block-scale coefficients of variation of both methods. The contrast reported here should therefore be read in the context of single-domain modelling.

Finally, the case study findings are based on a single deposit. The McLaughlin gold deposit provides a demanding and well-characterised test case, but the magnitude of the reported improvement in discounted cash flow should not be generalised to all precious-metal operations without further validation on deposits of different geological style and economic profile.

CONCLUSIONS

This study evaluated the effect of incorporating a locally adaptive geostatistical simulation method, the Adaptive Search Ellipsoid SGS (ASE-SGS), as an input to long-term open pit production scheduling for the McLaughlin gold deposit. Ten realisations were generated with conventional stationary SGS and ten with ASE-SGS, and the E-type mean of each suite was used as the grade input to a ten-year DBS optimisation.

ASE-SGS reduced block-scale uncertainty, as measured by the median coefficient of variation across realisations; it produced an E-type grade field that tracked composite data more closely along east–west and north–south swath transects; it recovered higher tonnage and average grade above the marginal cut-off in the grade-tonnage curve; and it yielded a discounted cash flow that exceeded the conventional-SGS schedule by approximately USD0.57 billion (48 per cent) over the ten-year horizon. These improvements are attributable to the use of locally fitted anisotropy directions and ranges at the simulation stage, rather than to any change in the scheduling parameters or economic assumptions.

The results support the proposition that, in structurally complex precious-metal deposits, geostatistical simulation inputs built on locally adaptive continuity assumptions can materially affect

long-term production planning outcomes. The most important limitation of the present analysis is that the scheduling step used the E-type mean of each suite rather than the individual realisations, and full propagation of conditional realisations through the stochastic optimisation framework is the focus of subsequent work. Further validation on deposits of different geological style, and quantification of the sensitivity of the schedule to the local variogram fitting threshold, are also identified as priorities for future research. In addition, future work will validate the simulated grade distributions within a fixed reference volume, such as the ten-year pit derived from the ASE-SGS E-type model, by comparing declustered composites, nearest-neighbour estimates, the two E-type models, and the individual realisations using grade–tonnage curves, histograms, and log-probability plots, to confirm that neither simulation method introduces artificial selectivity or high-grade bias relative to the conditioning data.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the Mining Engineering Department of the Colorado School of Mines for academic support throughout this research, and Homestake Mining Company for making the McLaughlin drill hole data set available for academic use. The authors also thank Deswik for providing access to the Deswik.GO academic license. The constructive feedback of colleagues and reviewers is gratefully acknowledged.

REFERENCES

Aras, C, 2018. A new integer solution algorithm to solve open-pit mine production scheduling problems, PhD thesis (unpublished), Colorado School of Mines, Golden, CO.

Aras, C, Dagdelen, K and Johnson, T, 2019. Generating pushbacks using direct block mine production scheduling algorithm, in *Mining Goes Digital* (eds: C Mueller, W Assibey-Bonsu, E Baafi, C Dauber, C Doran, M J Jaszczuk and O Nagovitsyn), pp 426–436 (Taylor and Francis Group: London).

Aydar, A and Dagdelen, K, 2025. Adaptive search ellipsoid: a new practical algorithm for local anisotropies in mineral resource estimation, in *Proceedings of the 42nd International Symposium on Applications of Computers and Operations Research in the Mineral Industry (APCOM 2025)*, pp 437–452 (Australasian Institute of Mining and Metallurgy: Perth).

Aydar, A, 2026. Simulation and stochastic mine planning in non-stationary domains, PhD thesis (unpublished), Colorado School of Mines, Golden.

Boisvert, J B and Deutsch, C V, 2011. Programs for kriging and sequential Gaussian simulation with locally varying anisotropy using non-Euclidean distances, *Computers and Geosciences*, 37(4):495–510.

Boisvert, J B, Manchuk, J G and Deutsch, C V, 2009. Kriging in the presence of locally varying anisotropy using non-Euclidean distances, *Mathematical Geosciences*, 41(5):585–601.

Chilès, J P and Delfiner, P, 2012. *Geostatistics*: *Modeling Spatial Uncertainty*, second edition (John Wiley and Sons: Hoboken).

Dagdelen, K and Kawahata, K, 2008. Value creation through strategic mine planning and cutoff-grade optimization, *Mining Engineering*, 60(1):39–45.

Deutsch, C V and Journel, A G, 1998. *GSLIB*: *Geostatistical Software Library and User's Guide*, second edition (Oxford University Press: New York).

Dimitrakopoulos, R and Ramazan, S, 2008. Stochastic integer programming for optimising long-term production schedules of open-pit mines: methods, application and value of stochastic solutions, *Mining Technology (Transactions of the Institutions of Mining and Metallurgy, Section A)*, 117(4):155–167.

Dimitrakopoulos, R, 2011. Stochastic optimization for strategic mine planning: a decade of developments, *Journal of Mining Science*, 47(2):138–150.

Dimitrakopoulos, R, Farrelly, C T and Godoy, M, 2002. Moving forward from traditional optimization: grade uncertainty and risk effects in open pit design, *Mining Technology (Transactions of the Institutions of Mining and Metallurgy, Section A)*, 111(1):82–88.

Dowd, P A, 1994. Risk assessment in reserve estimation and open-pit planning, *Transactions of the Institution of Mining and Metallurgy, Section A*: *Mining Industry*, 103:A148–A154.

Godoy, M and Dimitrakopoulos, R, 2004. Managing risk and waste mining in long-term production scheduling of open-pit mines, *SME Transactions*, 316:43–50.

Goovaerts, P, 1997. *Geostatistics for Natural Resources Evaluation* (Oxford University Press: New York).

Isaaks, E H, 1990. The application of Monte Carlo methods to the analysis of spatially correlated data, PhD thesis (unpublished), Stanford University, Stanford.

Johnson, T B, 1968. Optimum open pit mine production scheduling, PhD thesis (unpublished), Operations Research Department, University of California, Berkeley.

Journel, A G and Huijbregts, C J, 1978. *Mining Geostatistics* (Academic Press: London).

Journel, A G and Rossi, M E, 1989. When do we need a trend model in kriging?, *Mathematical Geology*, 21(7):715–739.

Lerchs, H and Grossmann, I F, 1965. Optimum design of open-pit mines, *Transactions CIM*, 58:47–54.

Machuca-Mory, D F and Deutsch, C V, 2013. Non-stationary geostatistical modeling based on distance weighted statistics and distributions, *Mathematical Geosciences*, 45(1):31–48.

Menabde, M, Froyland, G, Stone, P and Yeates, G A, 2007. Mining schedule optimisation for conditionally simulated orebodies, in *Orebody Modelling and Strategic Mine Planning* (ed: R Dimitrakopoulos), pp 379–384 (Australasian Institute of Mining and Metallurgy: Melbourne).

Ramazan, S and Dimitrakopoulos, R, 2013. Production scheduling with uncertain supply: a new solution to the open pit mining problem, *Optimization and Engineering*, 14(2):361–380.

Ravenscroft, P J, 1992. Risk analysis for mine scheduling by conditional simulation, *Transactions of the Institution of Mining and Metallurgy, Section A*: *Mining Industry*, 101:A104–A108.

Rossi, M E and Deutsch, C V, 2014. *Mineral Resource Estimation* (Springer: Dordrecht).

Sherlock, R L, Tosdal, R M, Lehrman, N J, Graney, J R, Losh, S, Jowett, E C and Kesler, S E, 1995. Origin of the McLaughlin Mine sheeted vein complex: metal zoning, fluid inclusion and isotopic evidence, *Economic Geology*, 90(8):2156–2181.

Tosdal, R M, Enderlin, D A, Nelson, G C and Lehrman, N J, 1993. Overview of the McLaughlin precious metal deposit, Napa and Yolo counties, northern California, in *Active Geothermal Systems and Gold/Mercury Deposits in the Sonoma/Clear Lake Volcanic Fields, California* (ed: J J Rytuba), Society of Economic Geologists Guidebook Series, 16:312–329.

Localised conditional simulation in mine planning – bridging deterministic and stochastic workflows

T Elkington[1] and R Santos[2]

1. General Manager, Snowden Optiro, Perth WA 6000.
   Email: tarrant.elkington@snowdenoptiro.com
2. Senior Consultant, Snowden Optiro, Perth WA 6000.
   Email: raphael.santos@snowdenoptiro.com

ABSTRACT

Conditional simulation is widely recognised as one of the most rigorous methods for quantifying geological uncertainty, yet its use within routine mine planning remains limited because planning workflows ultimately require a small number of models for optimisation, design and reporting. Localised conditional simulation (LCS) provides a practical bridge between deterministic and stochastic workflows by generating a single planning model that is intended to reproduce recoverable grade distributions while retaining the underlying simulation realisations for uncertainty analysis.

This paper presents a case study demonstrating the application of LCS within conventional open pit and underground mine planning workflows. The study examines three practical questions: how dilution and selectivity should be represented when moving from resource support to mining support; how pit optimisation outcomes change when alternative grade models are evaluated across multiple geological realisations; and how the same concepts can be applied to underground mineable shape optimisation (MSO).

The results show that support selection and dilution treatment are particularly important for recoverable resource models such as LCS. Application of a floating selective mining unit methodology materially altered the grade-tonnage response of the LCS and simulation models while having comparatively little impact on traditional ordinary kriging and E-type models. Pit optimisation completed on 33 alternative grade models demonstrated substantial variability in pit size and value. Cross-evaluation of candidate pit shells against 30 simulation realisations showed that the LCS-derived shell achieved the highest average value across the uncertainty space considered. Underground MSO analysis produced inventories generally consistent with the simulation ensemble while avoiding some of the continuity assumptions inherent in deterministic estimates.

The study demonstrates that LCS can be integrated into conventional mine planning workflows while retaining a direct link to geological uncertainty. The results support the use of LCS as a practical bridge between recoverable resource estimation and uncertainty-aware mine planning.

INTRODUCTION

Mine planning decisions are made under geological uncertainty. While conditional simulation has long been recognised as one of the most effective methods for quantifying this uncertainty, its application within routine mine planning workflows remains limited. Most planning decisions ultimately require a single block model from which engineers can optimise pits, generate mineable shapes, schedule production and report Ore Reserves. As a result, the benefits of simulation are often confined to resource estimation studies or standalone risk assessments, while day-to-day planning continues to rely on deterministic models.

This creates a practical disconnect between resource modelling and mine planning. Traditional kriged models provide a convenient basis for optimisation and design but are subject to smoothing effects that can distort recoverable grade-tonnage relationships and local selectivity. Recoverable resource estimation methods were developed to address these limitations by estimating the grades and tonnes that could realistically be recovered at a selective mining unit (SMU) scale rather than relying solely on smoothed block estimates. Methods such as Uniform Conditioning (UC), Multiple Indicator Kriging (MIK) and Localised Uniform Conditioning (LUC) have become established approaches for modelling recoverable resources at mining support (Abzalov, 2006; Hardtke, Allen and Douglas, 2011).

In parallel, developments in stochastic mine planning demonstrated the economic significance of geological uncertainty. Studies by Dimitrakopoulos and co-workers showed that pit designs generated from a single deterministic estimate can perform very differently when evaluated across alternative geological outcomes, highlighting the limitations of conventional optimisation approaches under uncertainty (Dimitrakopoulos, Farrelly and Godoy, 2002; Dimitrakopoulos, Martinez and Ramazan, 2007; Leite and Dimitrakopoulos, 2007). While these methods provide a rigorous treatment of uncertainty, their implementation can be challenging within routine planning environments. Elkington and Prott (2013) proposed a practical alternative in which candidate pit shells generated from different geological models were cross-evaluated against a suite of simulation realisations, demonstrating that uncertainty could be incorporated into pit selection using conventional optimisation software. This concept of separating optimisation from uncertainty evaluation provides an important link between deterministic planning workflows and simulation-based decision-making.

Localised conditional simulation (LCS) provides a further step in this evolution. By localising the outcomes of multiple conditional simulation realisations into a single planning model, LCS seeks to reproduce recoverable grade distributions at a practical mining support while retaining the underlying simulation ensemble for uncertainty analysis.

The objective of this paper is not to compare estimation methodologies, but rather to examine how an LCS model behaves when used within typical mine planning workflows. Through a case study based on a large gold deposit, the paper explores three practical questions:

1. How should dilution and selectivity be represented when moving from resource support to mining support?
2. How does the choice of grade model influence pit optimisation outcomes and the robustness of selected pit shells under uncertainty?
3. Can the same concepts be extended to underground mine planning through mineable shape optimisation?

In many respects, the workflow extends the ideas of Elkington and Prott (2013) by combining a localised planning model with simulation-based cross-evaluation. The objective is to assess whether a planning model derived from the simulation ensemble can provide robust mine planning outcomes while retaining the simulations for uncertainty assessment.

More specifically, this study aims to demonstrate how localised conditional simulation can be integrated into conventional open pit and underground mine planning workflows, providing a practical planning model that preserves recoverable resource characteristics while maintaining a direct link to geological uncertainty through the underlying simulation ensemble.

The results demonstrate that the benefits of recoverable resource estimation methods can only be fully realised when modifying factors and planning decisions are applied at an appropriate mining support. More broadly, they show how LCS can provide a practical bridge between conventional deterministic workflows and uncertainty-aware mine planning.

WHAT IS LCS?

Localised conditional simulation (LCS) is a recoverable resource estimation approach that converts a suite of conditional simulation realisations into a single planning model while retaining the underlying simulations for uncertainty analysis. The objective is to create a model that reproduces recoverable grade distributions at a selected planning support and can therefore be used directly within conventional mine planning workflows.

Conceptually, LCS belongs to a broader family of localisation approaches. Localisation methods seek to transform a probabilistic description of a deposit into a single model suitable for mine planning while preserving recoverable grade distributions at an appropriate support. Early applications focused on localising Uniform Conditioning results through ranking variables, leading to the development of Localised Uniform Conditioning (LUC) (Abzalov, 2006). More general localisation frameworks were subsequently proposed that could be applied to probabilistic models derived from simulation, MIK or other recoverable resource methods (Boisvert and Deutsch, 2012). LCS extends

these concepts by applying localisation directly to conditional simulation realisations, allowing uncertainty to be retained explicitly throughout subsequent mine planning studies.

LCS occupies a middle ground between deterministic estimation and full simulation-based planning. A traditional kriged or E-type model provides a single estimate that is convenient for planning but is affected by smoothing, which can reduce apparent grade variability and distort recoverable selectivity. Conditional simulations preserve variability and geological uncertainty but require multiple alternative outcomes to be managed throughout the planning process. LCS combines aspects of both approaches by generating a single planning model that reflects the recoverable characteristics of the simulation ensemble.

The localisation process begins with a set of conditional simulation realisations. Following an appropriate change of support, the grade distribution within each panel is derived from the simulations. Values sampled from these distributions are then assigned back to SMU blocks using the E-type model as a ranking variable. The ranking variable controls the spatial placement of values, while the simulation-derived distributions determine the grades assigned. The resulting model is intended to reproduce the recoverable distribution of grades within each panel while providing a spatial arrangement of high- and low-grade material guided by the E-type ranking variable. The process is shown conceptually in Figure 1.

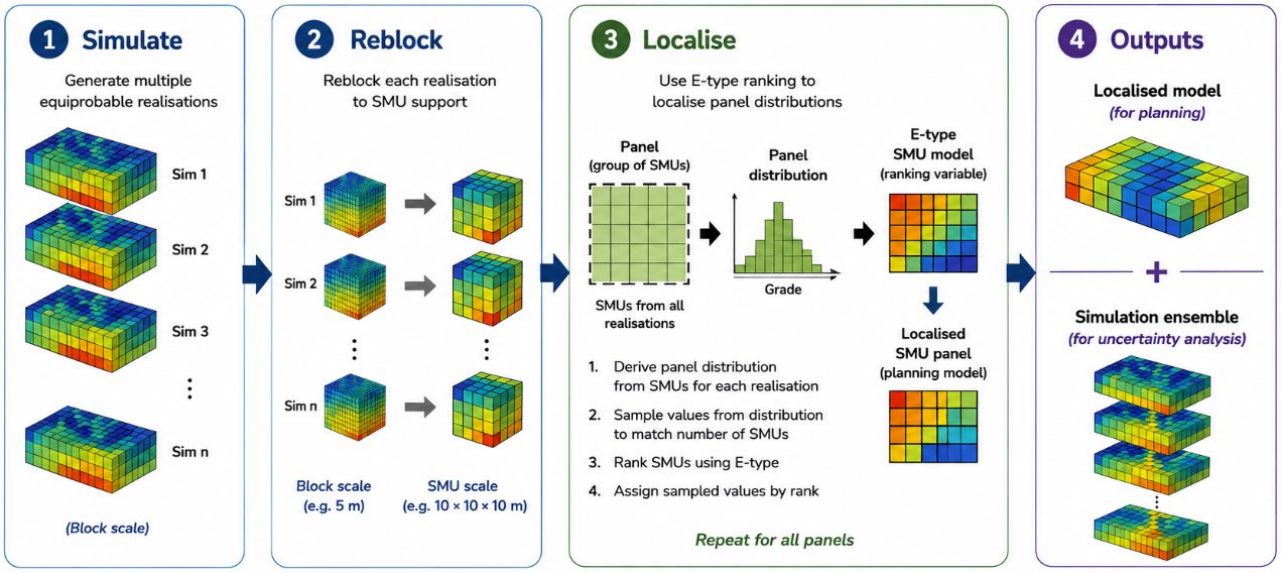


FIG 1 – Localised conditional simulation process.

A key distinction is that LCS is neither an average of the simulation realisations nor a replacement for them. The localised model provides a practical basis for optimisation, design and reporting, while the simulation realisations remain essential for assessing uncertainty and testing the robustness of planning decisions. In this way, LCS should be viewed as a planning model derived from uncertainty rather than a deterministic estimate that removes uncertainty.

This distinction is particularly important in the context of mine planning. Traditional stochastic workflows require planners either to optimise multiple realisations or to evaluate candidate designs across multiple geological outcomes. LCS provides an alternative by creating a single optimisation model that incorporates information from the simulation ensemble, while retaining the realisations for subsequent uncertainty assessment. In this sense, LCS complements approaches such as the cross-evaluation framework described by Elkington and Prott (2013), which demonstrated the value of assessing planning decisions against multiple geological outcomes rather than relying solely on the model from which the design was generated.

As with other recoverable resource approaches, support is critical. The objective is not to reproduce point-scale variability, but rather to reproduce the grade distributions expected at the support corresponding to future mining selectivity. Consequently, modifying factors such as dilution, ore loss and selectivity should be applied at the same practical mining support as the intended planning

decision. Failure to do so can result in planning outcomes that are inconsistent with the recoverable distributions represented by the model.

The value of LCS in mine planning is therefore pragmatic. It enables planners to work with a single model that better represents recoverable selectivity than a conventional deterministic estimate, while retaining a direct connection to the simulation realisations required for uncertainty assessment. The planning model supports optimisation and design, and the simulations provide the uncertainty framework within which alternative decisions can be evaluated.

CASE STUDY

The objective of this case study is to evaluate how a localised conditional simulation (LCS) model behaves when used within conventional mine planning workflows. The study is based on a subset of a large gold deposit for which ordinary kriging (OK), E-type, LCS and 30 conditional simulation realisations were available over a common mineralisation envelope.

The investigation is structured around three practical mine planning questions:

1. How should selectivity and dilution be represented when moving from resource support to mining support?
2. How does the choice of grade model influence pit optimisation outcomes and the robustness of selected pit shells under uncertainty?
3. Can the same concepts be extended to underground mine planning through mineable shape optimisation?

The intent here is not to explain the resource modelling methodologies in detail, but to examine the implications of different styles of model for mine planning workflows.

The same geological data set was used throughout. All models were estimated over the same overall mineralisation domain. The OK implementation used separate grade subdomains within this domain, whereas the LCS and simulation workflows used the combined mineralisation domain. The conditional simulations were generated at 5 × 5 × 2.5 m support and reblocked to 10 × 10 × 5 m SMU blocks, with localisation undertaken within 40 × 40 × 20 m panels. The OK model was estimated using 10 × 10 × 5 m parent blocks and subsequently reblocked to the same dimensions to account for edge subcell effects. Figure 2 presents the resource-scale grade-tonnage relationships for the OK, E-type, LCS and simulation models and illustrates the range of geological uncertainty represented by the simulation ensemble.

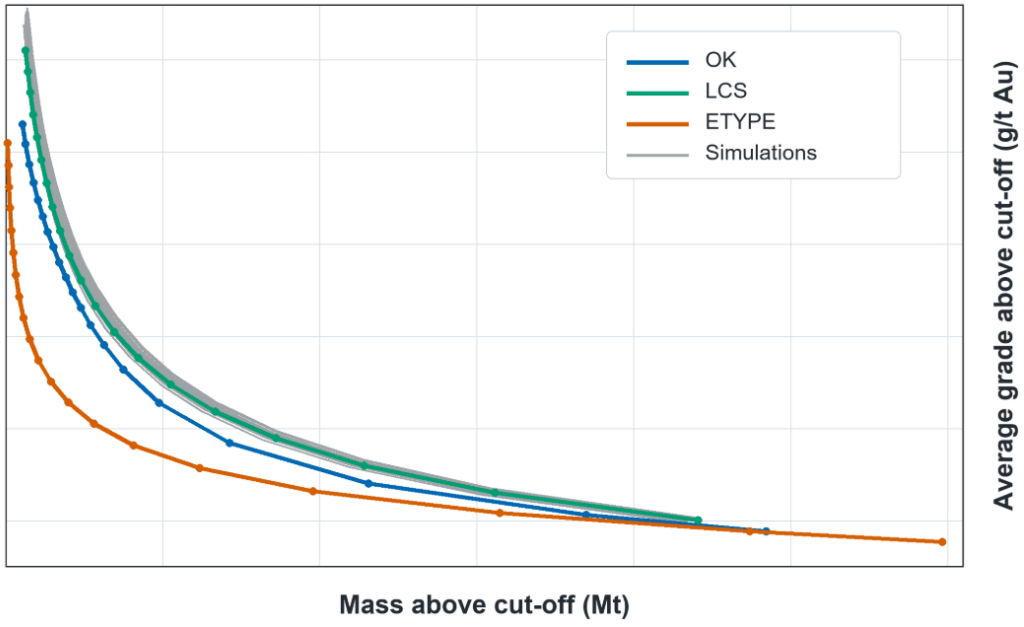


FIG 2 – Resource model grade tonnage curve for OK (blue), LCS (green), E-type (orange) and 30 simulations (grey).

While the resource-scale grade-tonnage responses differ between the modelling approaches, mine planning decisions are not made at resource support. The first question therefore considers how these models behave once translated to practical mining support.

Impact of mining support on recoverable selectivity

Before optimisation, the resource models must be transformed into a representation that reflects practical mining selectivity and dilution. This is particularly important for recoverable resource models such as LCS because they intentionally preserve more local grade variability than traditional kriged estimates.

Dilution is one of the most challenging modifying factors to apply consistently across different estimation methodologies. Common approaches include factoring, skin dilution and reblocking. For bulk mining operations, reblocking is often used as a practical means of approximating mining selectivity. However, reblocking alone does not necessarily produce a realistic representation of the material that would ultimately be mined.

The implications of this issue are illustrated in Figure 3, which compares the bench-scale presentation of the OK and LCS models after reblocking. The smoothing inherent in ordinary kriging produces relatively continuous zones of mineralisation that appear readily mineable. In contrast, the LCS model preserves substantially more local variability, resulting in a fragmented bench presentation that would be difficult to mine directly. This behaviour is not unique to LCS and is inherent to recoverable resource approaches that seek to preserve local variability.

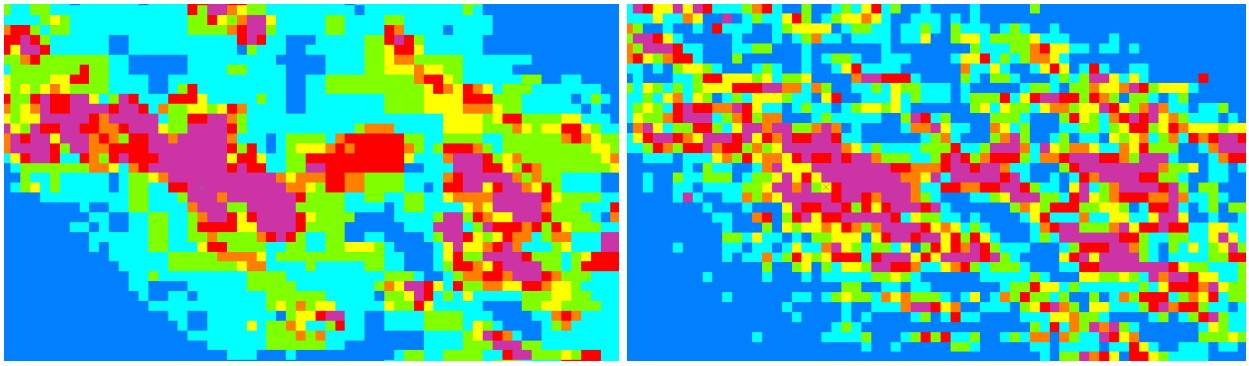


FIG 3 – Bench presentation of OK and LCS models coloured by grade (blue is low, purple is high).

Simply increasing the reblocking support does not adequately address this issue. In practice, mining units are defined through grade control and dig block selection processes that blend material over a larger support than the resource model blocks. To better represent this behaviour, a floating SMU methodology was developed for this study.

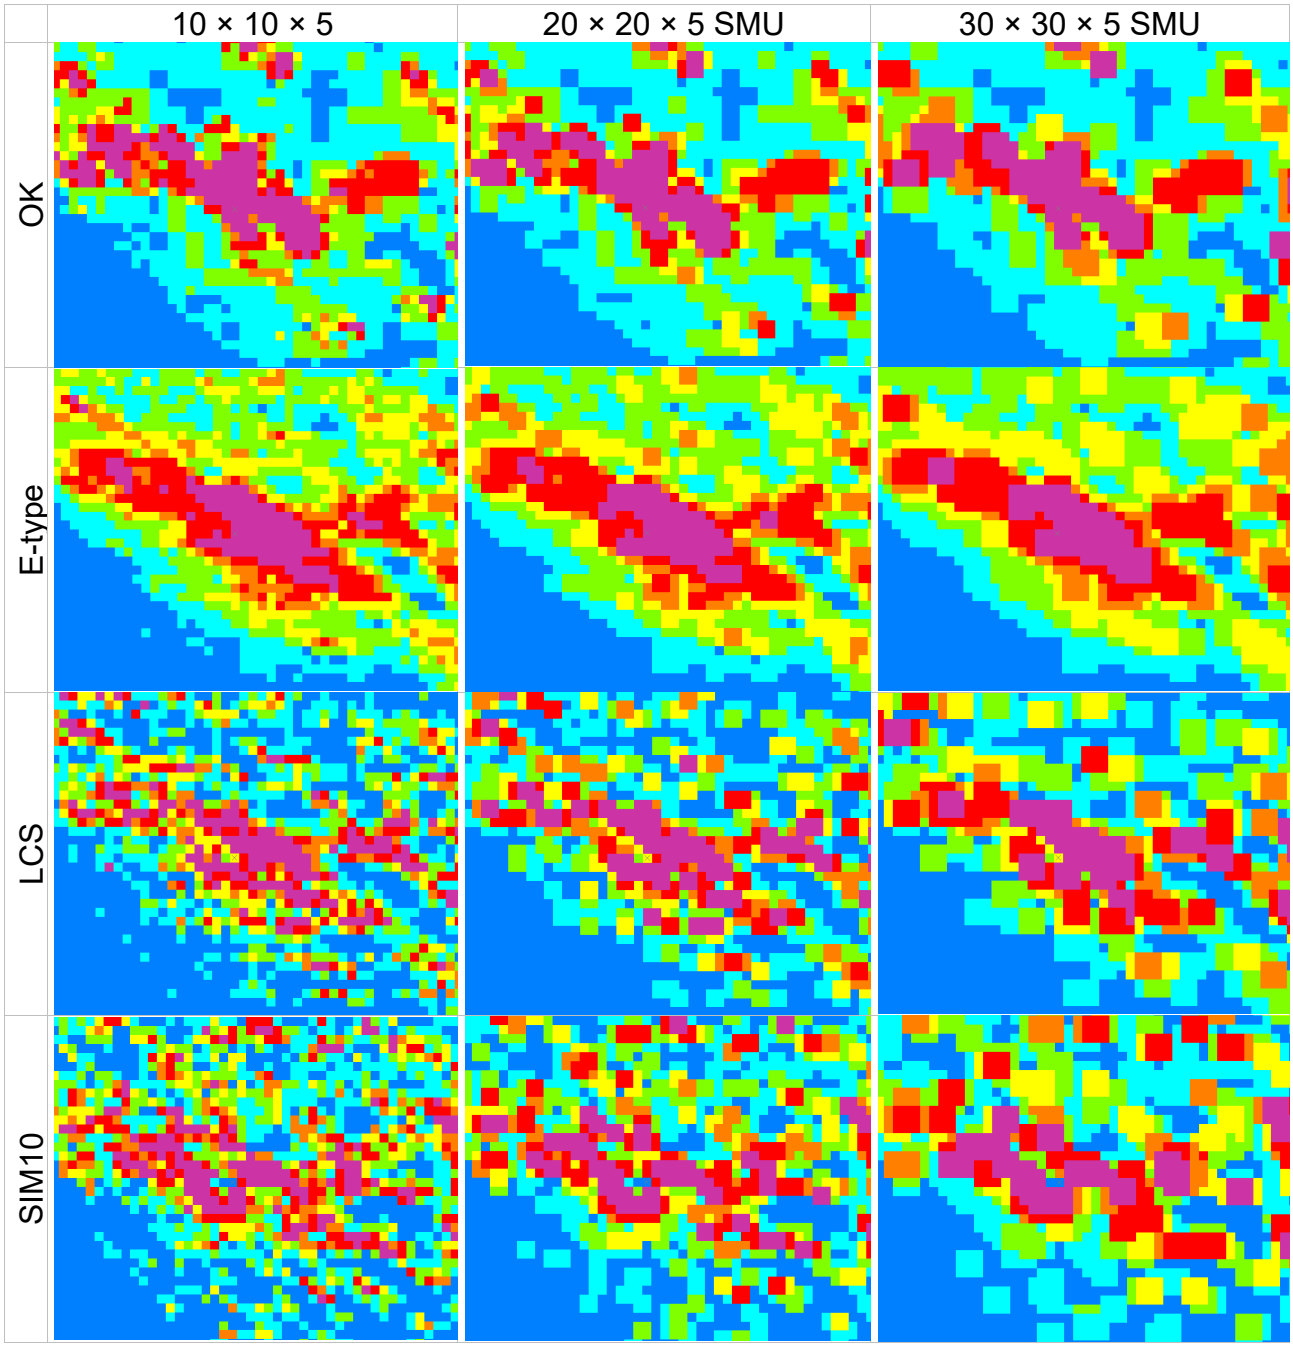
The floating SMU process was applied bench by bench to approximate the selectivity and dilution associated with a practical mining unit. A rectangular SMU window was moved across the original 10 × 10 × 5 m block model and the grades and tonnes of the blocks captured within the window were averaged to represent the material that would be mined together. The averaged value was then assigned back to the underlying blocks before the window progressed through the bench. The process therefore reduces local grade extremes, blends adjacent high- and low-grade material, and provides an estimate of the grade-tonnage response expected after applying the nominated mining support.

Importantly, the floating-SMU transformation applied in this study is distinct from the change-of-support procedure used during construction of the LCS resource model. The former is a mine planning approximation of modifying factors applied after resource estimation. Its purpose is not to generate a new recoverable resource estimate, but to approximate the additional dilution and loss of selectivity that occur when material represented at resource-model support is converted into practical mining units. It should therefore be interpreted as a mining-support transformation rather than a second geostatistical change-of-support procedure. The choice of mining support can

materially influence the grade-tonnage response and therefore the optimisation and design decisions that follow.

The term ‘floating SMU’ is used here to distinguish the procedure from conventional fixed-grid reblocking. The procedure is intentionally analogous to a spatial moving-window blending operation: the nominated mining unit is allowed to float relative to the underlying resource block grid so that the result is not dependent on an arbitrary fixed reblock origin. It is not intended to represent a formal geostatistical change of support. Rather, it provides a simple planning approximation of the blending and reduction in selectivity associated with mining at a support larger than the underlying 10 × 10 × 5 m blocks.

Two floating SMU supports were evaluated to demonstrate the sensitivity to mining support: 20 × 20 × 5 m and 30 × 30 × 5 m. Figure 4 compares the visual impact of the process on the OK, E-type, LCS and selected simulation models. The floating SMU has relatively little effect on the already-smoothed OK and E-type models. In contrast, the LCS and simulation models undergo substantial modification as local variability is progressively blended. As the SMU size increases, the differences between the models become less pronounced.



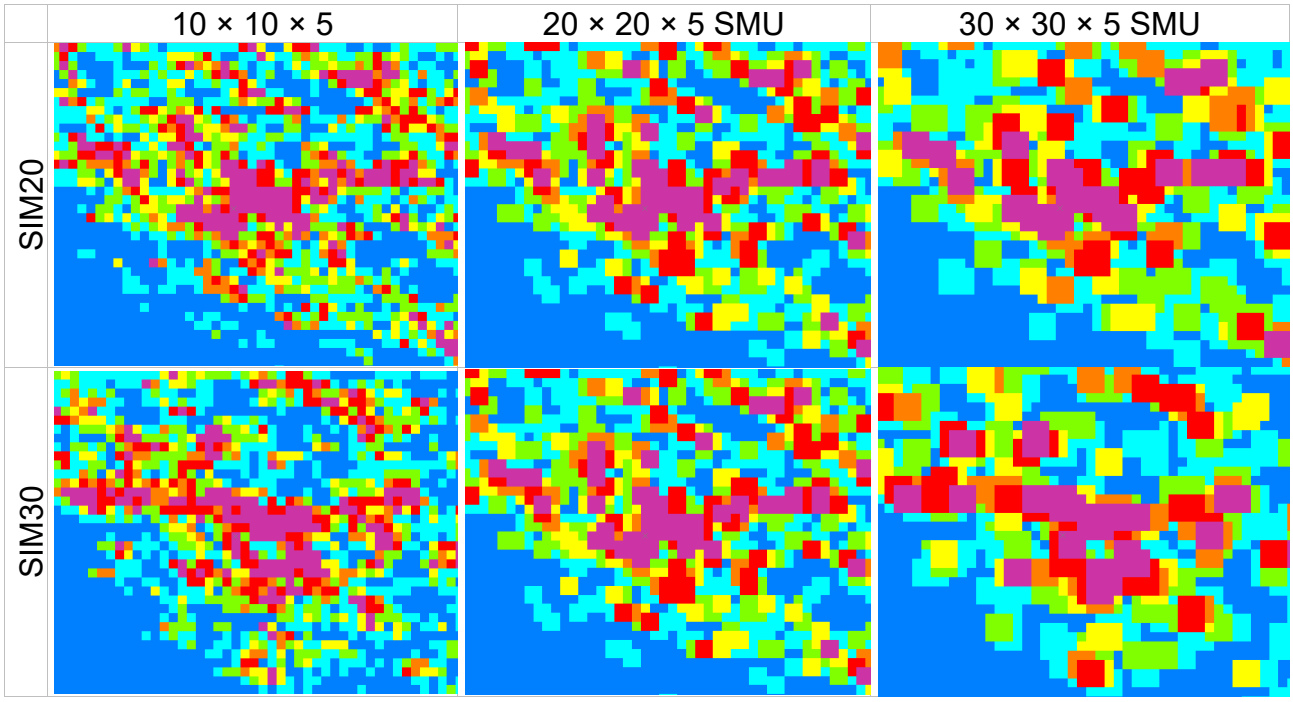


FIG 4 – Impact of floating SMU concept on OK, LCS, E-type and selected simulation models. Blue is low grade, purple is high grade.

Figure 5 shows the net change in tonnes, grade and contained metal when a common mining-support transformation is applied to each resource model. These changes should not be interpreted directly as physical dilution or ore loss. Each estimation approach provides a different starting representation of the spatial distribution of grade and, consequently, of the selectivity available at the original model support. For example, the comparatively smoothed OK and E-type models contain less local variability prior to application of the floating SMU, whereas the LCS and simulation models retain substantially greater local variability. The larger response of the latter models therefore reflects the greater amount of selectivity contained in their starting representations and the effect of reducing that selectivity at a common mining support.



FIG 5 – Relative change in tonnes, grade and contained metal from original resource model (10 × 10 × 5 m blocks).

Impact on pit optimisation and pit selection under uncertainty

Once mining support had been applied, pit optimisations were completed on each of the 33 candidate models (OK, E-type, LCS and 30 simulation realisations) using identical economic and geotechnical assumptions. The only difference between the optimisations was the grade model used, with each model derived from the 20 × 20 × 5 m floating SMU output. The 20 × 20 × 5 m support was carried forward for the subsequent analysis to demonstrate the mine planning workflow without unnecessarily repeating the optimisation analysis at multiple mining supports; it is not intended to represent an optimised or recommended SMU size. The objective was to assess how the choice of grade model influences pit design and whether different modelling approaches produce materially different outcomes under geological uncertainty.

Plan and section views of the resulting pit shells are shown in Figure 6. The differences between the resulting shells are substantial. The largest pit, generated from simulation realisation SIM20, is more than twice the volume of the smallest pit, generated from the E-type model. This demonstrates the sensitivity of optimal pit limits to the underlying representation of grade and geological continuity.

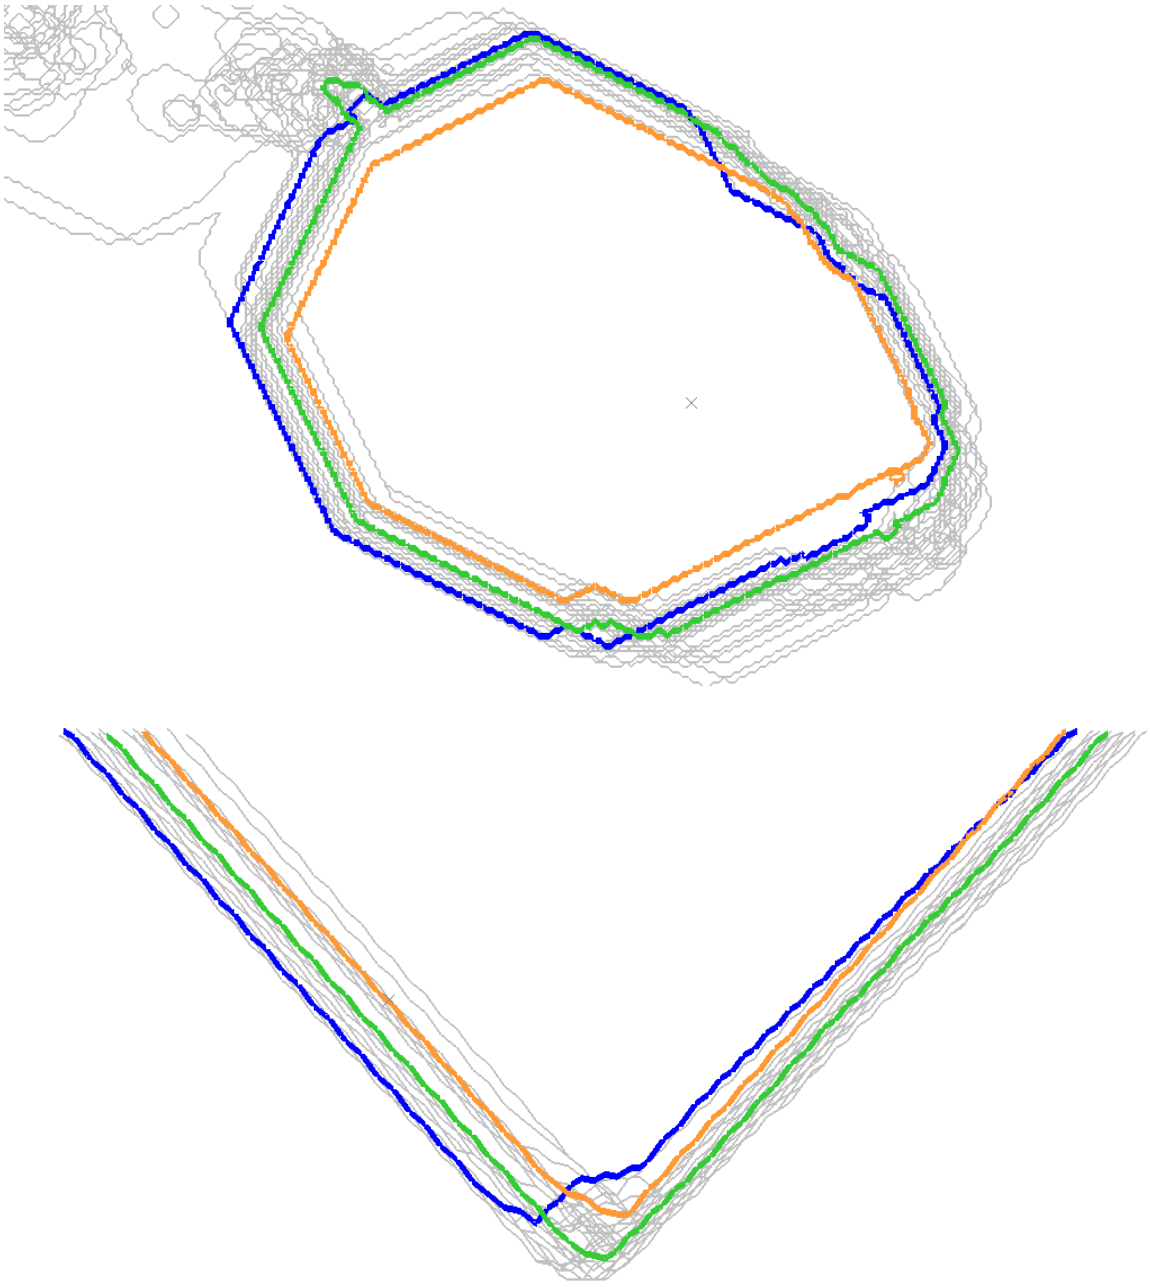


FIG 6 – Pit optimisation outlines for LCS (green), OK (blue), E-type (orange) and individual simulations (grey).

While each simulation-derived pit may be optimal for a particular geological outcome, it is impossible to know in advance which outcome will ultimately be realised. Consequently, selecting a pit solely because it performs well under a single geological outcome may not produce the most robust planning decision. A more meaningful assessment is obtained by evaluating each candidate pit shell across the full range of plausible geological outcomes represented by the simulation ensemble.

To achieve this, each of the 33 candidate pit shells was evaluated against all 30 simulation realisations. This cross-evaluation framework extends the concept proposed by Elkington and Prott (2013), whereby optimisation and uncertainty assessment are treated as separate activities. The key difference is that the candidate shells considered here include an LCS-derived planning model rather than being generated solely from individual simulation realisations. The pit shell defines the design decision, while the simulation realisations provide the uncertainty space against which that decision is tested. A shell optimised on one model may perform well when assessed against that model but less favourably when evaluated across the full suite of possible geological outcomes.

The results of the cross-evaluation are shown in Figure 7. The LCS-derived shell achieved the highest average value across the simulation ensemble while maintaining a coefficient of variation comparable with the alternative candidate shells. The E-type shell produced a slightly lower average value but also a lower coefficient of variation, reflecting its smaller pit size. The shells generated directly from individual simulation realisations achieved mean values between 2 per cent and 12 per cent lower than the LCS shell.

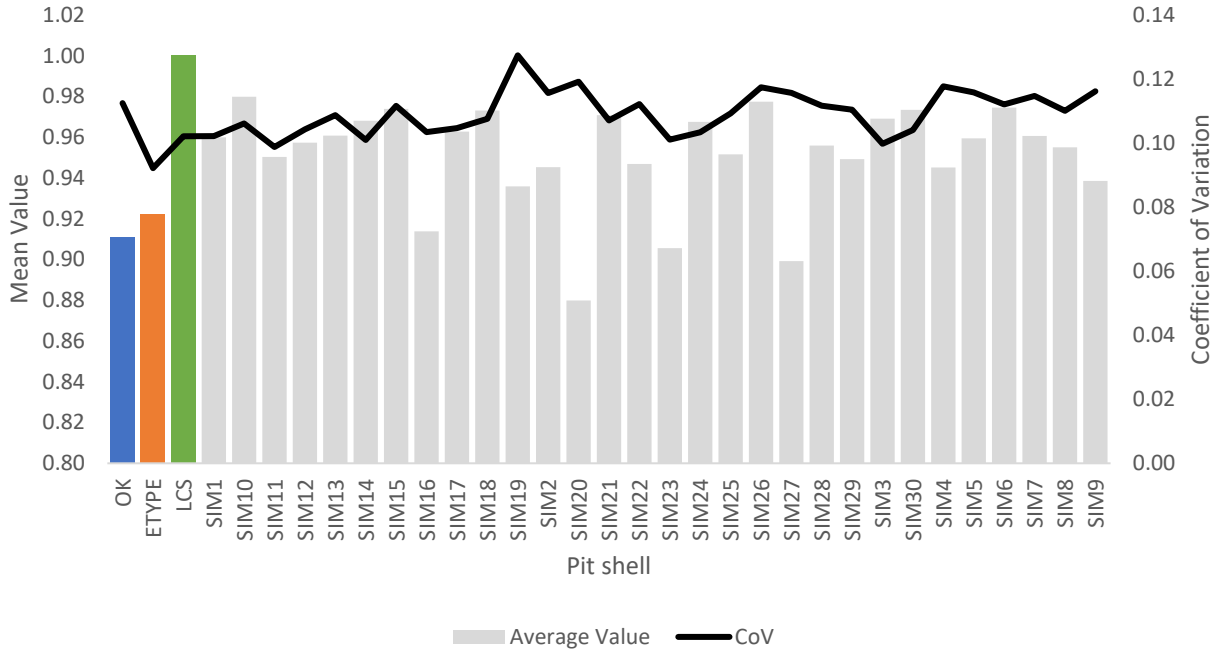


FIG 7 – Value cross-evaluation of candidate pit shells against the simulation realisations. The highest mean candidate is set to 100.

This outcome is consistent with the objective of localisation. The LCS model incorporates information from all simulation realisations and seeks to reproduce the recoverable characteristics of the simulation ensemble rather than any individual geological outcome. Consequently, it is not expected to produce the highest value when evaluated against a single realisation. Instead, its value lies in providing a balanced planning model that performs well across the full uncertainty space represented by the simulations.

The range of pit shells shown in Figure 6 also provides insight into the spatial distribution of uncertainty within the deposit. Areas that are consistently included across most pit shells represent relatively robust design decisions, whereas areas where pit boundaries diverge indicate locations where additional drilling may materially influence the final design. In this way, the simulation

ensemble provides information not only on economic uncertainty but also on where future data collection may have the greatest value.

The results suggest that pit selection under uncertainty should not focus solely on identifying the shell with the highest value for a single geological interpretation. Rather, the objective should be to identify a shell that performs robustly across a range of plausible outcomes. In this case study, the LCS shell provided the highest average value across the uncertainty space considered, demonstrating how localisation can be used to translate multiple simulation realisations into a single planning model for pit optimisation.

Future work could further test the robustness of the LCS-based pit optimisation by cross-evaluating the candidate pits against an independently generated set of conditional simulation realisations.

Extension to underground mine planning

Open pit boundaries are typically established well in advance of mining and are therefore particularly sensitive to geological uncertainty. Underground mines operate differently because additional geological information is acquired during development and production drilling. Consequently, the focus of early-stage underground planning is often not to define the exact location of future stopes, but rather to understand the likely inventory available for extraction and the associated uncertainty. The objective of this analysis was therefore to investigate whether the observations from the open pit workflow also apply to underground mine planning and to assess how alternative grade models influence mineable shape optimisation outcomes.

As demonstrated in the open pit examples, the OK, E-type and LCS models differ significantly in their representation of local continuity and grade variability. These differences are particularly important for underground mining because mineable stopes require sufficient continuity of economic mineralisation to satisfy geometric and operational constraints. Models that overstate continuity may generate overly optimistic inventories, while models that understate continuity may fail to identify potentially mineable material.

To investigate this effect, MSO analyses were completed on the same suite of 33 grade models using consistent stope dimensions, mining assumptions and a range of cut-off grades. The resulting inventory grade-tonnage curves are shown in Figure 8.

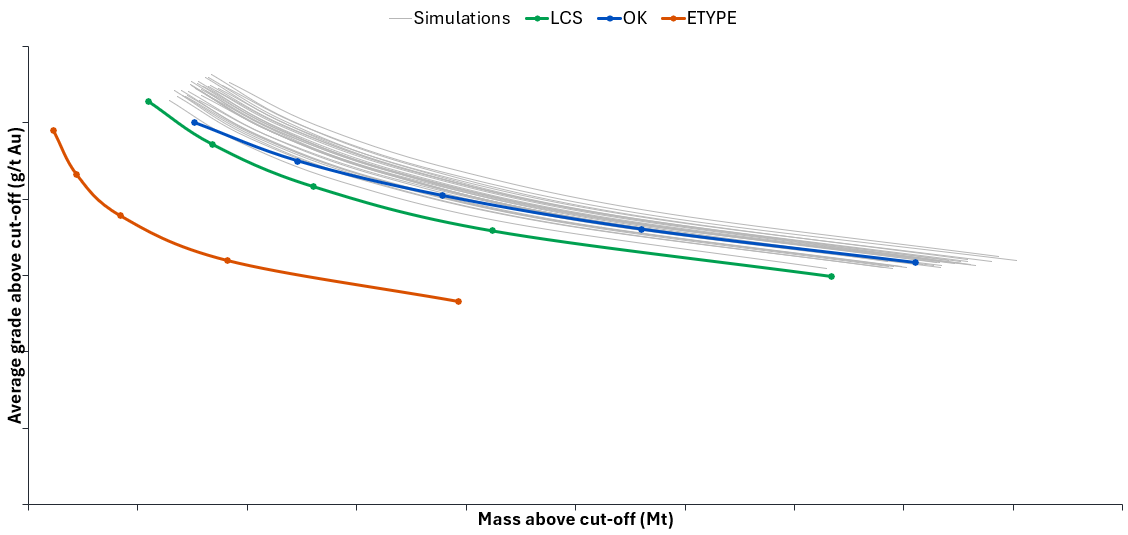


FIG 8 – Grade-tonnage curve of MSO inventories for each of the models.

The results indicate clear differences between the alternative modelling approaches. The E-type model consistently reports the lowest inventory, reflecting the smoothing effect associated with averaging across simulation realisations. High-grade responses are smoothed and local selectivity is reduced, resulting in fewer mineable shapes satisfying the optimisation criteria. The simulation realisations generally produce the largest inventories, as they preserve both local grade variability and geological continuity.

The LCS inventories lie between these two extremes. Although the localisation process is intended to reproduce the recoverable grade distributions of the simulations at the panel scale, it does not preserve the same degree of local spatial continuity as an individual simulation realisation. As a consequence, some mineable shapes that exist within individual simulations are not reproduced within the localised model, resulting in slightly lower inventory estimates. This behaviour is expected and reflects one of the fundamental trade-offs associated with localisation.

Plan and section views of the resulting stopes are shown in Figure 9. At the local scale, there is considerable variability in the location and geometry of individual stopes between the simulation realisations. This variability reflects the geological uncertainty present within the deposit and highlights the difficulty of defining definitive underground extraction shapes at an early planning stage.

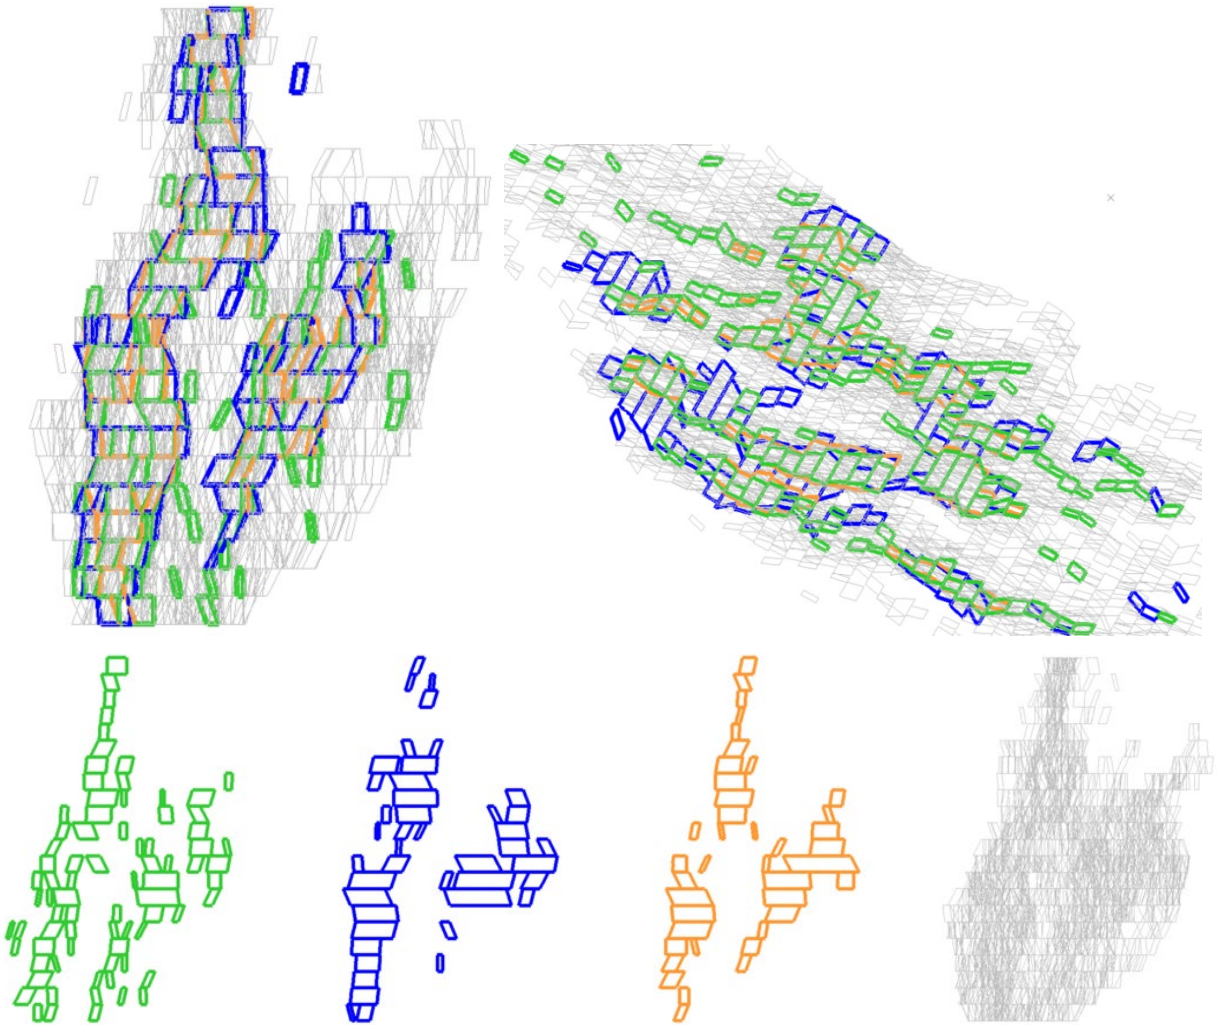


FIG 9 – Stope outlines in section and plan for LCS (green), OK (blue), E-type (orange) and individual simulations (grey).

However, when the results are aggregated to a level scale, a different picture emerges. Figure 10 shows the tonnes above cut-off by level for the lowest cut-off grade considered. While the precise location of stopes varies between models, the overall distribution of tonnes by level shows broad agreement between the simulation ensemble, the LCS model and the OK model. This suggests that although geological uncertainty may significantly affect the location of individual stopes, the larger-scale inventory available for extraction remains relatively stable.

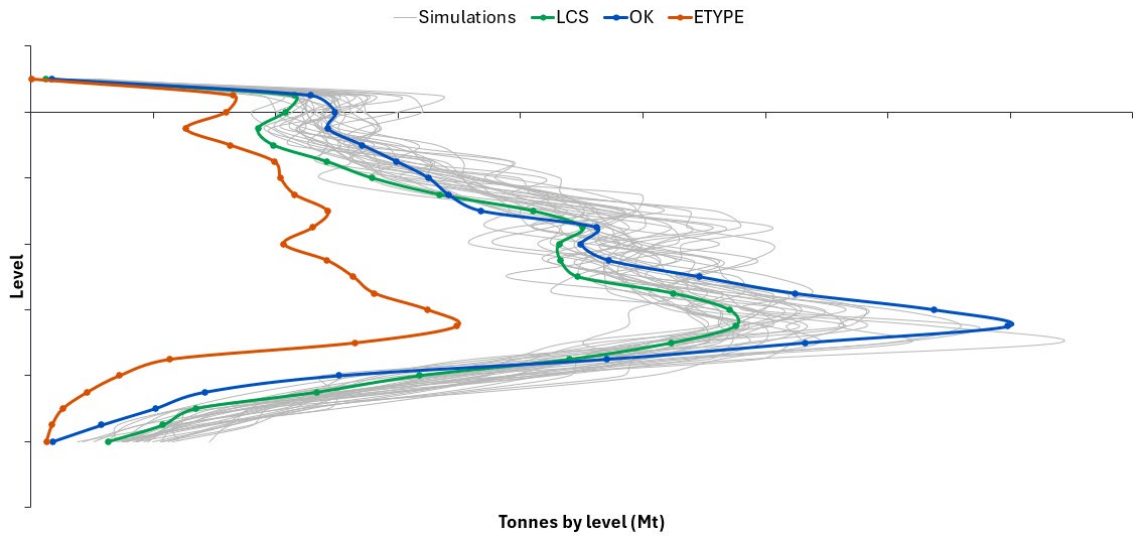


FIG 10 – Tonnes above cut-off (lowest cut-off) by level for each of the models.

From a planning perspective, this may be the more important outcome. Early-stage underground studies are often concerned with development requirements, infrastructure placement, production capacity and economic potential rather than the precise geometry of future stopes. The results indicate that LCS provides a realistic representation of recoverable inventory while avoiding some of the continuity assumptions inherent in deterministic models. Although the localised model produces slightly lower inventories than the simulations, it offers a practical basis for underground planning while retaining access to the simulation ensemble for uncertainty assessment.

Overall, the underground analysis demonstrates that the benefits of LCS extend beyond open pit optimisation. By linking a recoverable planning model with the underlying simulation realisations, the approach provides a framework for evaluating both the expected inventory and the uncertainty associated with underground development opportunities.

While the level-based inventory distributions are relatively stable, the spatial variability observed between the simulation realisations remains important for early-stage underground planning. Individual stope locations are expected to change as additional geological information becomes available through development and grade control drilling. Major development and infrastructure decisions, however, are comparatively difficult and costly to modify once implemented. The simulation ensemble can therefore provide insight into the likely spatial distribution of future mining areas and the locations of greater geological uncertainty, helping to inform infrastructure placement and identify where additional drilling may have the greatest planning value.

DISCUSSION

The case study demonstrates that the choice of grade model influences mine planning outcomes well beyond the resource estimation stage. While ordinary kriging, E-type and localised conditional simulation are often compared on the basis of grade-tonnage characteristics, the results presented here show that their differences become more pronounced once practical mining considerations such as selectivity, dilution, pit optimisation and mineable shape generation are introduced. The implications are particularly important because mine planning decisions are generally made at mining support rather than resource model support.

The mining support analysis highlights one of the key challenges associated with recoverable resource approaches. Models such as LCS are specifically designed to preserve local grade variability and recoverable distributions. As a consequence, they exhibit a greater degree of local heterogeneity than traditional kriged estimates. While this is desirable from a resource estimation perspective, it can produce patterns that are not directly mineable at the intended selective mining unit. The results show that applying a practical mining support through the floating SMU process materially alters the grade-tonnage response of the LCS and simulation models, whereas the impact on the already-smoothed OK and E-type models is comparatively minor. This demonstrates that the response to mining support depends on the degree of selectivity represented by the underlying

resource model. The appropriate support for the mining decision must be incorporated before optimisation; otherwise, the planning process may be based on a representation of selectivity that cannot realistically be achieved in practice.

The pit optimisation results demonstrate the value of evaluating candidate designs across multiple geological outcomes rather than assessing each design only against the model from which it was generated. Individual simulation optimisations produced a wide range of pit sizes and values, reflecting the uncertainty present within the deposit. None of these pits can be considered the definitive answer because each is conditioned on a single possible geological outcome. Cross-evaluation against all simulation realisations provides a more meaningful measure of robustness by assessing how each candidate shell performs across the full uncertainty space. Within this framework, the LCS-derived shell achieved the highest average value while maintaining a level of variability comparable with the alternative candidate shells. This outcome is consistent with the objective of localisation, which seeks to represent the recoverable characteristics of the entire simulation ensemble rather than any individual realisation. The results therefore suggest that LCS can provide a practical planning model that supports robust decision-making without requiring planners to optimise and evaluate every realisation individually.

The comparison between LCS and E-type is also informative. In this case study, the E-type model consistently produced the smallest pit shells and the lowest recoverable inventories. This reflects the well-known smoothing effects associated with averaging across simulation realisations, which tend to suppress local high-grade responses and reduce apparent selectivity. While E-type models provide a useful estimate of expected grade, they may under-represent the recoverable value available to mining operations. The contrast between E-type and LCS demonstrates the importance of preserving recoverable grade distributions when the objective is mine planning rather than resource reporting alone.

The underground results highlight a different consequence of localisation. Unlike open pit design, where the ultimate pit boundary is generally established well before mining occurs, underground mining benefits from additional information obtained through development and grade control drilling. Consequently, the primary objective of early-stage underground planning is not necessarily to define the exact location of future stopes, but rather to understand the likely inventory available for extraction and the associated uncertainty. The MSO analysis indicates that LCS generally produces inventories that are slightly lower than those obtained directly from the simulation realisations. This behaviour is expected because localisation is intended to reproduce recoverable distributions at the selected support but does not preserve the same degree of spatial continuity as individual simulations. From a planning perspective, this may be viewed as conservative, reducing the risk of overestimating continuity while still maintaining a realistic representation of recoverable grades.

The value of LCS is not that it replaces simulation, but that it provides a practical planning model derived from simulation that can be used within conventional mine planning workflows. The workflow presented in this study occupies a practical middle ground between traditional deterministic planning and fully stochastic optimisation. Fully stochastic approaches can explicitly incorporate uncertainty and risk preferences within the optimisation process, but they often require specialised software, significant computational effort and expertise that may not be available in routine planning environments. The workflow adopted here retains familiar optimisation and design tools while using simulation realisations to evaluate robustness after optimisation. Although this approach does not fully capture uncertainty within the optimisation process itself, it provides a transparent and accessible framework for incorporating geological uncertainty into planning decisions. For many operations, this may represent a more practical pathway toward uncertainty-aware mine planning than adopting fully stochastic optimisation methods.

More broadly, the results suggest that recoverable resource estimation and mine planning should not be viewed as separate activities. The benefits of recoverable resource methods are only realised when modifying factors, optimisation and design decisions are applied at an appropriate mining support. The combination of LCS with simulation-based uncertainty analysis provides a framework that links resource estimation, mine planning and risk assessment within a single workflow. Further work is required to test the approach within production scheduling and reserve estimation workflows,

and to assess its applicability across a broader range of commodities, mining methods and geological settings.

CONCLUSIONS

This case study demonstrates how localised conditional simulation can be integrated into conventional open pit and underground mine planning workflows while retaining geological uncertainty throughout the planning process. LCS provides a planning model that reproduces recoverable grade distributions at the selected support while maintaining the underlying simulation realisations required for uncertainty analysis.

The mining support analysis highlights the importance of applying modifying factors at an appropriate support when using recoverable resource models. The floating SMU approach produced substantially greater changes to the grade-tonnage response of the LCS and simulation models than to the already-smoothed OK and E-type models. This difference does not indicate greater physical dilution or ore loss, but instead reflects the greater local variability and selectivity represented by these models at the starting support.

The pit optimisation study showed that geological uncertainty can produce a wide range of candidate pit shells and economic outcomes. Cross-evaluation of all candidate shells against the simulation ensemble provided a practical measure of robustness. In this case study, the LCS-derived shell achieved the highest average value across the suite of geological outcomes considered, suggesting that localisation provides an effective mechanism for translating multiple simulations into a single planning model for pit optimisation.

The underground MSO analysis demonstrated that LCS produced inventories that were generally consistent with the simulation ensemble while avoiding some of the continuity assumptions inherent in deterministic models. Although the localised model produced slightly lower inventories than the simulations, it provided a realistic representation of recoverable mineralisation and a useful basis for evaluating underground development options.

Overall, the study demonstrates that LCS provides a practical and effective bridge between recoverable resource estimation and uncertainty-aware mine planning. By combining a single planning model with simulation-based uncertainty assessment, the approach enables conventional optimisation and design workflows to incorporate geological uncertainty without requiring fully stochastic optimisation methods.

Future work should focus on extending the workflow into production scheduling, reserve estimation and reconciliation studies. In particular, the ability to compare deterministic planning outcomes derived from LCS with fully stochastic scheduling approaches offers an opportunity to better understand the value of geological uncertainty throughout the mining value chain.

REFERENCES

Abzalov, M, 2006. Localised uniform conditioning: a new approach for direct modelling of small blocks, *Mathematical Geology,* 38(4):393–411.

Boisvert, J B and Deutsch, C V, 2012. A short note on a general framework and software for localization of mining reserves, *CCG Annual Report*, paper 313.

Dimitrakopoulos, R, Farrelly, C and Godoy, M, 2002. Moving forward from traditional optimisation: grade uncertainty and risk effects in open-pit design, *Transactions of the Institution of Mining and Metallurgy*, *Mining Technology,* 111:A82–A88.

Dimitrakopoulos, R, Martinez, L and Ramazan, S, 2007. Optimising open pit design with simulated orebodies and Whittle Four-X: a maximum upside/minimum downside approach, in *Orebody Modelling and Strategic Mine Planning*, 2nd edn, pp 201–206 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Elkington, T and Prott, R, 2013. Simple method to select a pit in the presence of grade uncertainty, *Mining Engineering*, 65(12):63–66.

Hardtke, W, Allen, L and Douglas, I, 2011. Localised indicator kriging, in *Proceedings of the 35th APCOM Symposium*, pp 141–147.

Leite, A and Dimitrakopoulos, R, 2007. Stochastic optimisation model for open pit mine planning: application and risk analysis at copper deposit, *Transactions of the Institution of Mining and Metallurgy*, *Mining Technology,* 116(3):A109–A118.

A kernel mean embedding method for high-order geostatistical simulation in orebody modelling

J Herrero[1] and R Dimitrakopoulos[2]

1. Postdoctoral researcher, Stochastic Mine Planning Laboratory, Department of Mining and Materials Engineering, McGill University, Montreal QC H3A 2A7, Canada. Email: julien.herrero@mcgill.ca
2. Professor, Stochastic Mine Planning Laboratory, Department of Mining and Materials Engineering, McGill University, Montreal QC H3A 2A7, Canada. Email: roussos.dimitrakopoulos@mcgill.ca

ABSTRACT

Orebody models must reproduce not only spatial grade variability but also the spatial connectivity and multi-scale geological structures that control mineable volumes, selectivity, and production forecasts. In practice, widely used simulation workflows based on two-point statistics and Gaussian assumptions such as sequential Gaussian simulation may generate diffuse realisations and fail to reproduce mineralisation geometries and high-grade connectivity, which can propagate into limiting effects in mine planning and production forecasting. Alternatively, multiple-point statistics can reproduce complex geological patterns through training images, but their performance may degrade when the spatial statistics of the training image differ from those of the conditioning data. High-order geostatistical approaches address these limitations by explicitly modelling multipoint spatial interactions beyond second-order statistics. Such statistics enable stochastic simulations to capture non-Gaussian behaviour and complex geological architectures. This contribution presents a novel high-order geostatistical simulation method based on kernel mean embedding with a Gaussian kernel function. By operating in a reproducing kernel Hilbert space, the method leverages a computationally tractable statistical learning formulation to estimate local joint distributions, where high-order spatial statistics are reproduced without requiring their explicit computation. Local spatial patterns are learned from both conditioning data and geological information extracted from a training image. To address potential inconsistencies between the spatial statistics of the training image and those of the conditioning data, the proposed approach introduces a tuning parameter controlling their relative influence. The proposed workflow is evaluated on 2D test cases for controlled validation and a 3D gold deposit application representative of operational orebody modelling challenges. Results show improved reproduction of geological continuity through high-order spatial statistics compared with other methods. These improvements support more reliable uncertainty quantification for downstream mine planning, with potential benefits for selectivity analysis and reserve sensitivity assessment. Future developments could extend the framework to multivariate simulation of correlated attributes and investigate its application to direct block-support simulation for mine planning workflows.

INTRODUCTION

Orebody modelling plays a central role in mineral resource evaluation and mine planning, as geological models directly influence estimates of grade distribution, mineable volumes, production schedules, and economic forecasts. Because only a limited number of drill hole samples are available, geostatistical techniques are commonly used to interpolate grades and quantify geological uncertainty throughout the deposit (Journel and Huijbregts, 1976).

Traditional deterministic estimation methods, such as kriging, provide smooth estimates of spatial grade distributions but do not adequately represent geological uncertainty (Asghari, Soltni and Amnieh, 2009). To address this limitation, stochastic simulation methods based on random field models were developed to generate multiple equally probable realisations of the orebody while reproducing the available data and prescribed spatial statistics. Among these methods, sequential Gaussian simulation (SGS) remains one of the most widely used approaches in geostatistical applications (Goovaerts, 1997; Chilès and Delfiner, 2012).

Despite their practical success, SGS and related simulation techniques are fundamentally based on second-order spatial statistics, typically represented through covariance or variogram models. While second-order statistics adequately describe Gaussian random fields, they are often insufficient to characterise the complex geological structures encountered in mineral deposits. Features such as curvilinear mineralisation trends, folded structures, nested geological events, and the connectivity of high-grade zones involve interactions among multiple spatial locations that cannot be fully described by pairwise correlations alone (Guardiano and Srivastava, 1993; Journel, 2005; Dimitrakopoulos, Mustapha and Gloaguen, 2010). As a consequence, SGS realisations may appear overly diffuse and may underestimate the continuity of economically significant mineralised domains.

To better reproduce complex geological architectures, multiple-point statistics (MPS) methods were introduced during the 1990s (Guardiano and Srivastava, 1993; Strebelle, 2002). Rather than relying exclusively on variograms, MPS methods use a training image (TI) representing geological concepts or analogues of the deposit under study. Spatial patterns extracted from the TI are reproduced during simulation, allowing realistic reconstruction of channels, lenses, vein systems, and other complex structures. However, because MPS methods are primarily driven by the TI, their performance depends strongly on the representativeness of the selected analogue. When the spatial statistics of the TI differ from those of the conditioning data, inconsistencies may arise and simulated realisations may fail to honour the actual geological information available from drilling (Osterholt and Dimitrakopoulos, 2007; Goodfellow *et al,* 2012).

An alternative framework is provided by high-order geostatistical simulation methods, which extend conventional geostatistical modelling beyond second-order statistics. Instead of describing spatial variability solely through pairwise relationships, these methods do not rely on restrictive distributional assumptions and explicitly model multipoint spatial interactions using high-order spatial moments and cumulants (Mustapha and Dimitrakopoulos, 2010, 2011). High-order spatial moments and cumulants may be viewed as extensions of covariance capable of characterising complex multipoint spatial structures. These statistics capture geological features such as connectivity, directional continuity, multiple-point spatial organisation, and non-Gaussian behaviour. Previous developments have demonstrated that high-order simulation methods can reproduce geological structures while remaining primarily driven by the spatial information contained in the available data (Mustapha and Dimitrakopoulos, 2010; Minniakhmetov, Dimitrakopoulos and Godoy, 2018; Yao, Dimitrakopoulos and Gamache, 2018).

More recently, statistical learning approaches based on reproducing kernel Hilbert spaces (RKHS) and spatial Legendre moment kernels have opened new possibilities for high-order geostatistical modelling (Yao, Dimitrakopoulos and Gamache, 2020, 2021). Unlike previous high-order simulation methods that may become numerically unstable at increasing orders, kernel methods provide a stable and flexible framework for representing probability distributions and complex spatial statistics without explicitly computing high-order moments or cumulants. Through a kernel mean embedding (KME), local spatial distributions can be compared and learned directly in a high-dimensional feature space, allowing complex geological patterns to be incorporated into stochastic simulation without restrictive distributional assumptions.

This paper presents a new high-order geostatistical simulation framework that extends previous RKHS/KME-based approaches through the use of Gaussian radial basis function (RBF) kernels. The proposed Gaussian kernel provides a universal representation in which high-order spatial statistics are reproduced implicitly through the kernel geometry rather than through predefined moment-based constructions. Additionally, to address potential statistical conflicts between conditioning data and the training image, a weighting parameter is introduced to control the relative influence of both information sources.

The methodology is evaluated using synthetic examples on the Stanford V Reservoir data set and a three-dimensional gold deposit representative of operational orebody modelling challenges.

METHODOLOGY

Kernel-based learning of local spatial distributions

The proposed methodology combines concepts from high-order geostatistics and kernel-based statistical learning to infer local conditional grade distributions during simulation. The proposed framework represents local probability distributions through KME in an RKHS (Smola *et al,* 2007; Song *et al,* 2008; Gretton *et al,* 2012).

A Gaussian RBF kernel is adopted to compare grade values and local spatial patterns:

$$k(x, y) = \exp\left(-\frac{(x-y)^2}{(2\sigma_k^2)}\right) \quad (1)$$

where x and y denote the values being compared and σ_k controls the similarity scale. Because Gaussian kernels are characteristic kernels, they implicitly encode information from moments of all orders through the geometry of the RKHS. Consequently, complex spatial structures and connectivity patterns can be reproduced without explicitly computing high-order spatial cumulants.

Rather than estimating a single interpolated value, the method infers a complete local conditional probability distribution at each simulation node. This distribution is represented as a convex mixture of Gaussian prototype distributions:

$$f(z) = \sum_i \alpha_i p_i(z) \quad (2)$$

where $p_i(z) = \mathbb{N}(\mu_i, \sigma^2{}_i)$, with μ_i and σ_i denoting the mean and standard deviation of prototype i, respectively, and the coefficient α_i is the mixture weight associated with the prototype. Prototype means μ_i are uniformly distributed over the range of observed grades, and the corresponding mixture weights α_i are estimated by solving a convex optimisation problem subject to non-negativity and unit-sum constraints. The Gaussian prototypes do not imply a Gaussian random field assumption. Instead, they provide a flexible basis capable of representing highly non-Gaussian and multimodal local distributions.

Learning from conditioning data and training images

The local distribution is inferred from two complementary information sources. The first source corresponds to the conditioning data, including drill hole samples and previously simulated nodes. Neighbouring observations are selected within a search template and weighted according to a conventional variogram model. This information source preserves consistency with the available data and ensures compatibility with established geostatistical modelling principles.

The second source corresponds to a TI representing geological concepts and spatial patterns. Local neighbourhood configurations are compared with TI patterns using kernel similarity measures, and the most compatible patterns are retained and used to infer complex spatial structures, connectivity, and multi-scale geological continuity that cannot be reliably reconstructed from sparse data alone.

The information provided by both sources is combined directly in RKHS through a weighting parameter:

$$b = \lambda b_{data} + (1 - \lambda) b_{TI} \quad (3)$$

where b_{data} and b_{TI} denote the embeddings derived from conditioning data and TI patterns, respectively. The tuning parameter λ controls the relative influence of both information sources. Values close to one favour conditioning data, whereas smaller values increase the influence of the TI. This mechanism provides a practical solution when statistical discrepancies exist between TI patterns and conditioning data.

The optimal local distribution is obtained by minimising the discrepancy between the candidate distribution (Equation 2) and the target RKHS embedding (Equation 3). In practice, this corresponds to minimising a Maximum Mean Discrepancy (MMD) criterion in feature space (Gretton *et al,* 2012). The resulting convex optimisation problem is solved under simplex constraints to infer a valid probability distribution.

SEQUENTIAL SIMULATION WORKFLOW

The methodology is implemented within a conventional sequential simulation framework. A randompath is first generated through the simulation grid. At each unsimulated node, a local neighbourhood is constructed and both the conditioning-data embedding b_{data} and the TI embedding b_{TI} are evaluated. The weighted embedding (Equation 3) is then used to infer a local conditional probability distribution by minimising the MMD criterion, from which a simulated value is sampled. The newly simulated node is subsequently added to the data set. The procedure is repeated until the simulation grid is fully informed. Multiple realisations can then be generated to quantify geological uncertainty and evaluate the impact of alternative spatial scenarios on grade-tonnage relationships, resource estimation, and mine-planning decisions.

Figure 1 summarises the overall workflow of the proposed kernel-based high-order simulation framework.

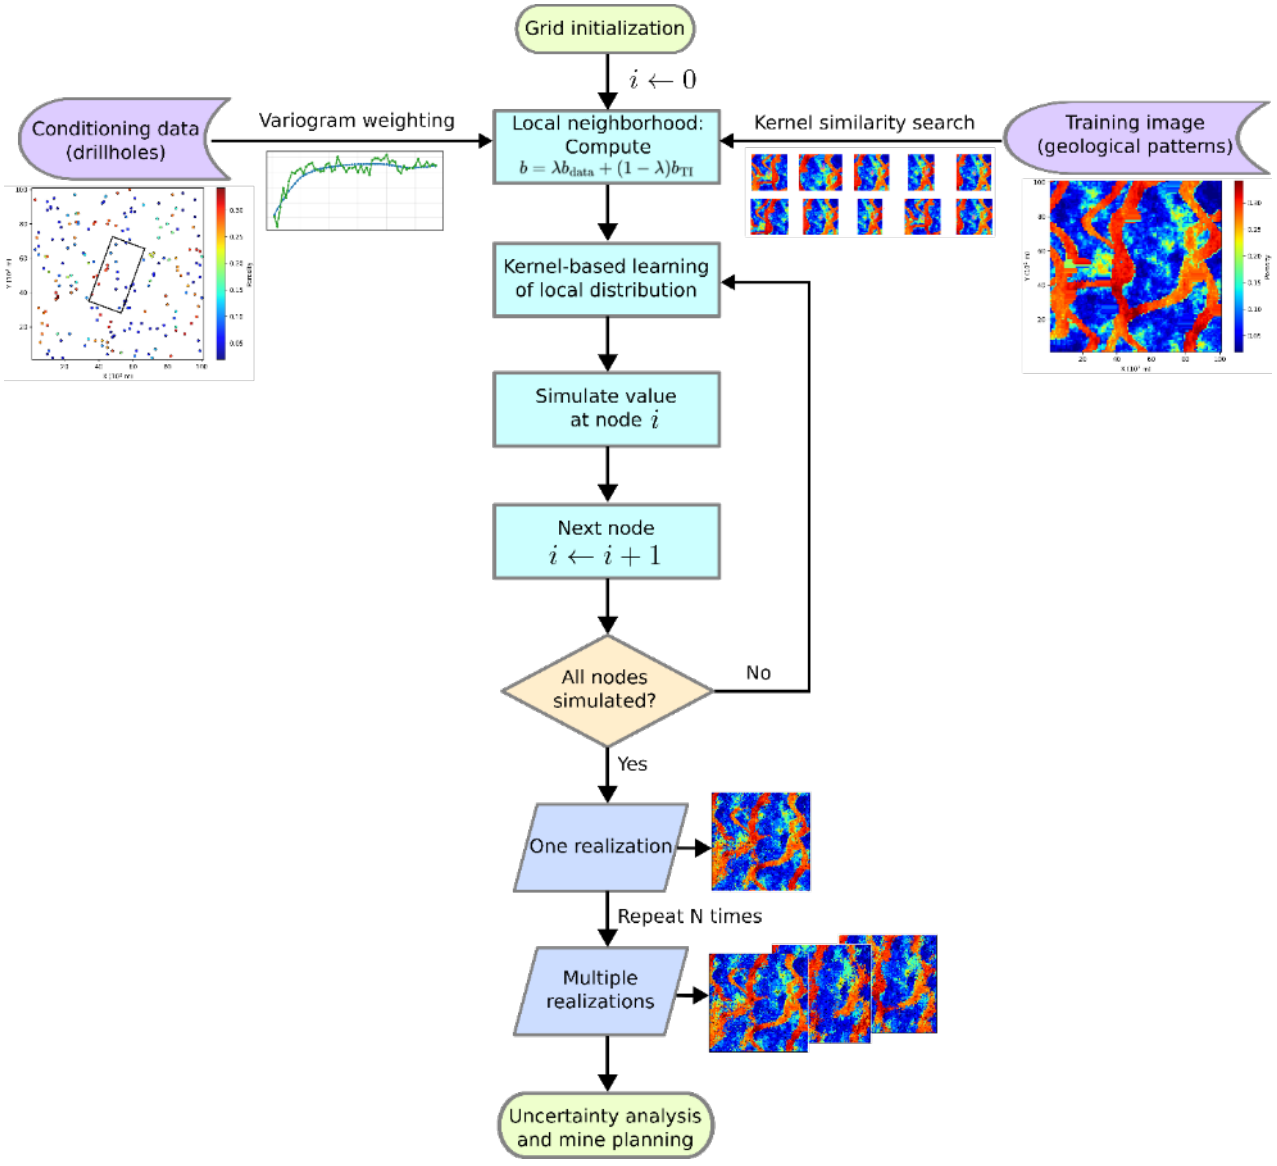


FIG 1 – Workflow of the proposed kernel-based high-order simulation method.

RESULTS

Synthetic Stanford V case study

Reference model and experimental set-up

A synthetic channelised reservoir extracted from the Stanford V benchmark data set (Mao and Journel, 1999) is used to evaluate the proposed kernel-based simulation framework. This data set, which has already been employed in previous studies on high-order geostatistical simulation (Yao, Dimitrakopoulos and Gamache, 2018, 2020), exhibits channelised structures characterised by complex curvilinear shapes, strong directional continuity, and non-Gaussian spatial organisation. Hence, it provides a challenging benchmark for stochastic simulation methods.

The reference model consists of a two-dimensional porosity field discretised on a 100 × 100 grid. This model is characterised by highly connected channels embedded within a lower-porosity background (Figure 2a). A data set of 200 conditioning points is randomly extracted from the exhaustive image and used as hard data throughout the experiments (Figure 2b).

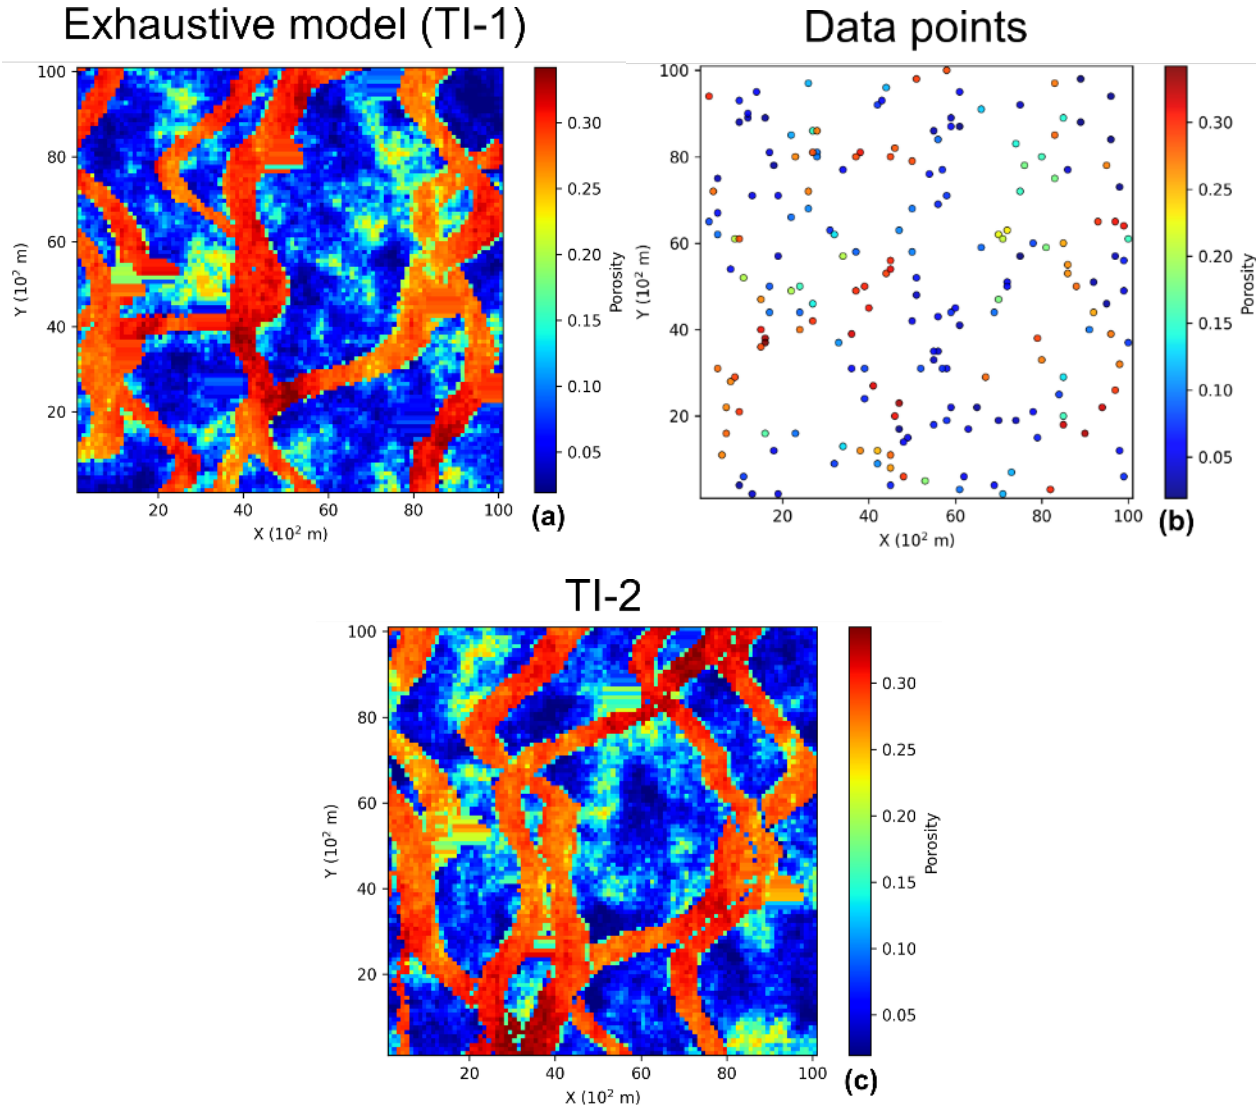


FIG 2 – (a) Exhaustive porosity model extracted from the Stanford V reservoir, also used as the reference training image TI-1; (b) Conditioning data extracted from the exhaustive image, consisting of 200 randomly distributed points; (c) Additional training image extracted from the Stanford V reservoir model, denoted TI-2.

Two training images are considered. The first training image (TI-1) corresponds to the exhaustive reference model itself and therefore represents an idealised case where the TI perfectly describes the target spatial structures. The second training image (TI-2) is extracted from a different layer of the Stanford V model and exhibits similar channelised patterns while remaining statistically distinct from the reference model (Figure 2c). This second configuration is more representative of practical applications, where the available geological analogue is only partially representative of the target deposit.

Comparison with conventional simulation methods

To evaluate the proposed framework, simulation results are first compared with several widely used geostatistical approaches. Ordinary kriging (OK) is first considered as a deterministic baseline, as it remains one of the most commonly used interpolation methods in industrial mineral resource estimation and orebody modelling (Journel and Huijbregts, 1976). SGS is then used as a representative two-point stochastic simulation approach based on covariance modelling (Goovaerts, 1997). Finally, results are compared with an MPS method called Direct Sampling (DS) (Mariethoz, Renard and Straubhaar, 2010), which is applied using both TI-1 and TI-2. The proposed kernel-based approach, hereafter referred to as the KME approach, is then evaluated using both TI-1 and TI-2. For the KME simulations, a weighting coefficient of $\lambda = 0.1$ is selected in order to emphasise the influence of the TI and facilitate comparison with the DS results.

Figure 3 presents representative realisations obtained with each method. The ordinary kriging estimate (Figure 3a) produces a smooth interpolation of the conditioning data and significantly attenuates the spatial variability observed in the reference model. The channelised structures are largely blurred and the resulting model fails to reproduce the complex connectivity of the underlying geological system. SGS (Figure 3b) restores part of the missing variability while matching both the conditioning data and the variogram model. However, the realisation remains relatively diffuse and does not recover the characteristic geometry of the channelised structures. This behaviour is consistent with the limitations of two-point geostatistics, where spatial continuity is described exclusively through pairwise relationships. DS realisations generated using TI-1 and TI-2 are shown in Figure 3c and 3d, respectively. In both cases, the simulated models exhibit realistic channelised patterns inherited from the TIs, illustrating the ability of the tested DS implementation to reproduce complex geological architectures when the TI is consistent with the conditioning data. As expected, the realisation generated with TI-1 provides the closest visual agreement with the reference model since the TI corresponds directly to the target geological reality. Nevertheless, both realisations show locally noisy patterns, with an imperfect reproduction of the small-scale porosity variations observed within the channelised structures.

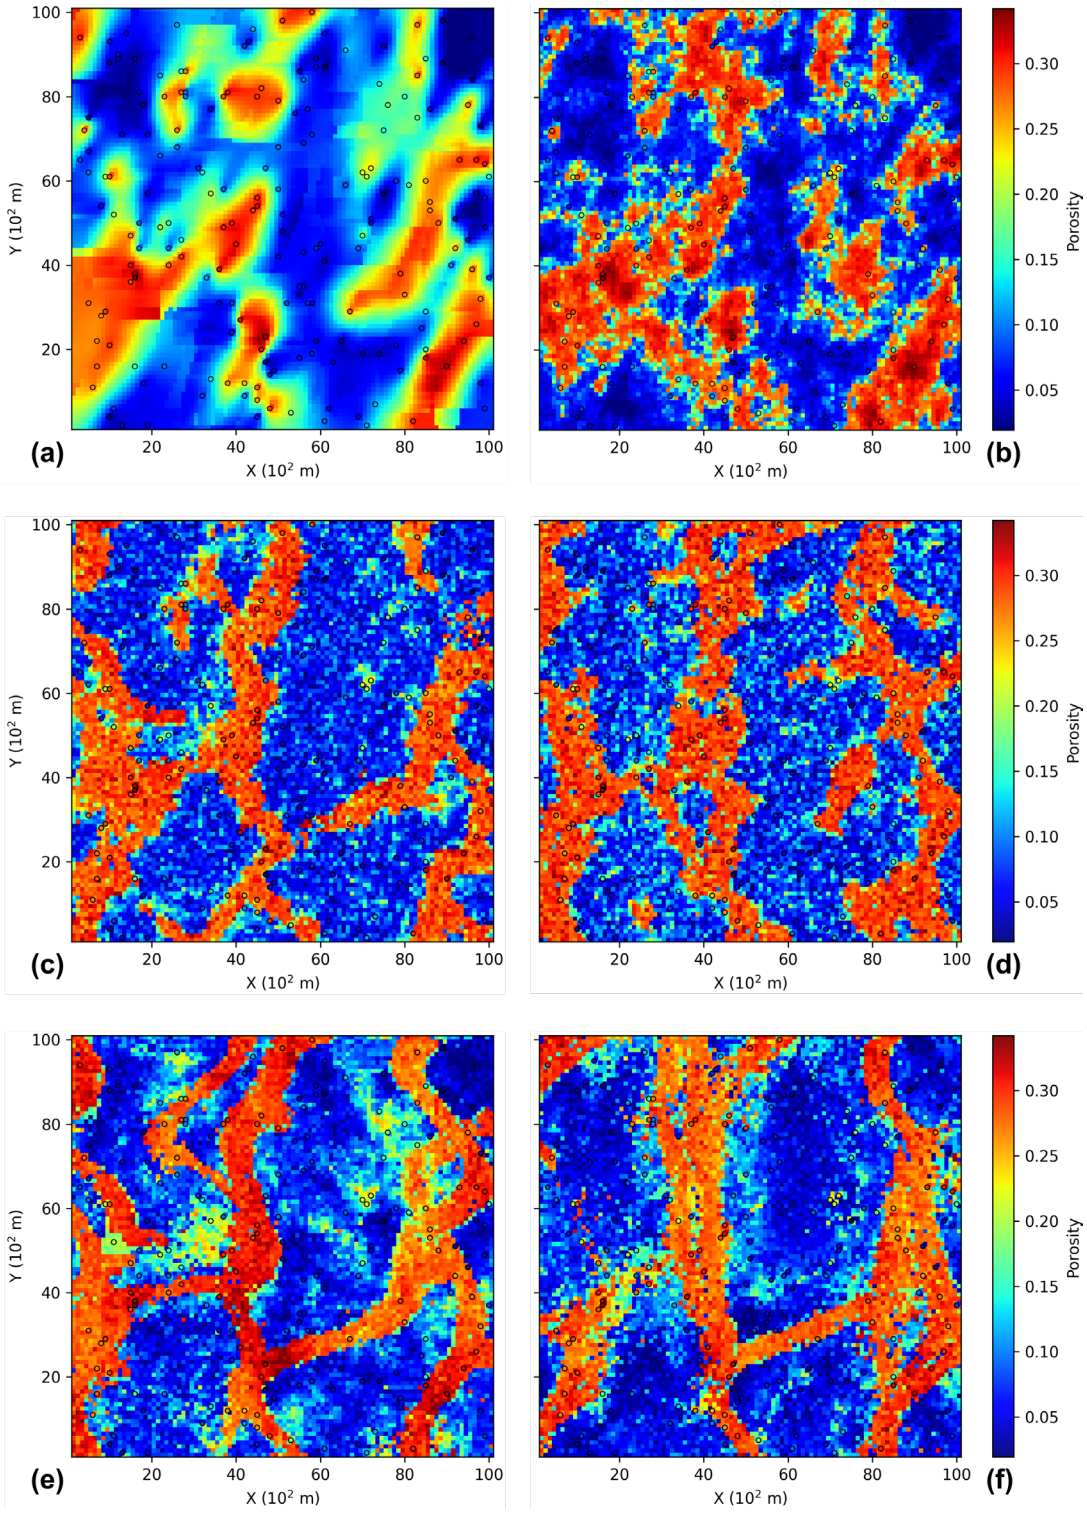


FIG 3 – (a) Ordinary kriging result; (b) SGS realisation; (c) and (d) DS realisations obtained using TI-1 and TI-2, respectively; (e) and (f) KME realisations obtained using TI-1 and TI-2, respectively.

The KME realisations obtained using TI-1 and TI-2 are shown in Figure 3e and 3f. Similar to the DS results, the proposed approach successfully reproduces the dominant channelised structures and geological continuity. However, the resulting realisations appear slightly more coherent in the preservation of local variability and fine-scale structures. In particular, the realisation obtained using TI-1 closely reproduces both the large-scale channel geometry and the fine-scale porosity heterogeneity observed in the reference model (Figure 2a). Even when using TI-2, which differs from the reference model, the resulting realisation preserves the main channel organisation and internal spatial variability.

Figure 4 presents histograms (top) and directional variograms (bottom) computed from one kriging and ten stochastic realisations generated using each approach. The OK estimate (Figure 4a) exhibits the well-known smoothing effect. The histogram is strongly concentrated around the mean value and underestimates the variability observed in the conditioning data distribution, while the variogram indicates an overestimation of spatial continuity relative to the reference model. The SGS realisations (Figure 4b) reproduce both the global histogram and the target variogram, confirming the consistency of the method with the imposed two-point statistics. Similar behaviour is observed for the DS realisations (Figure 4c and 4d) and KME realisations (Figure 4e and 4f), which also preserve the global probability distribution and second-order spatial continuity.

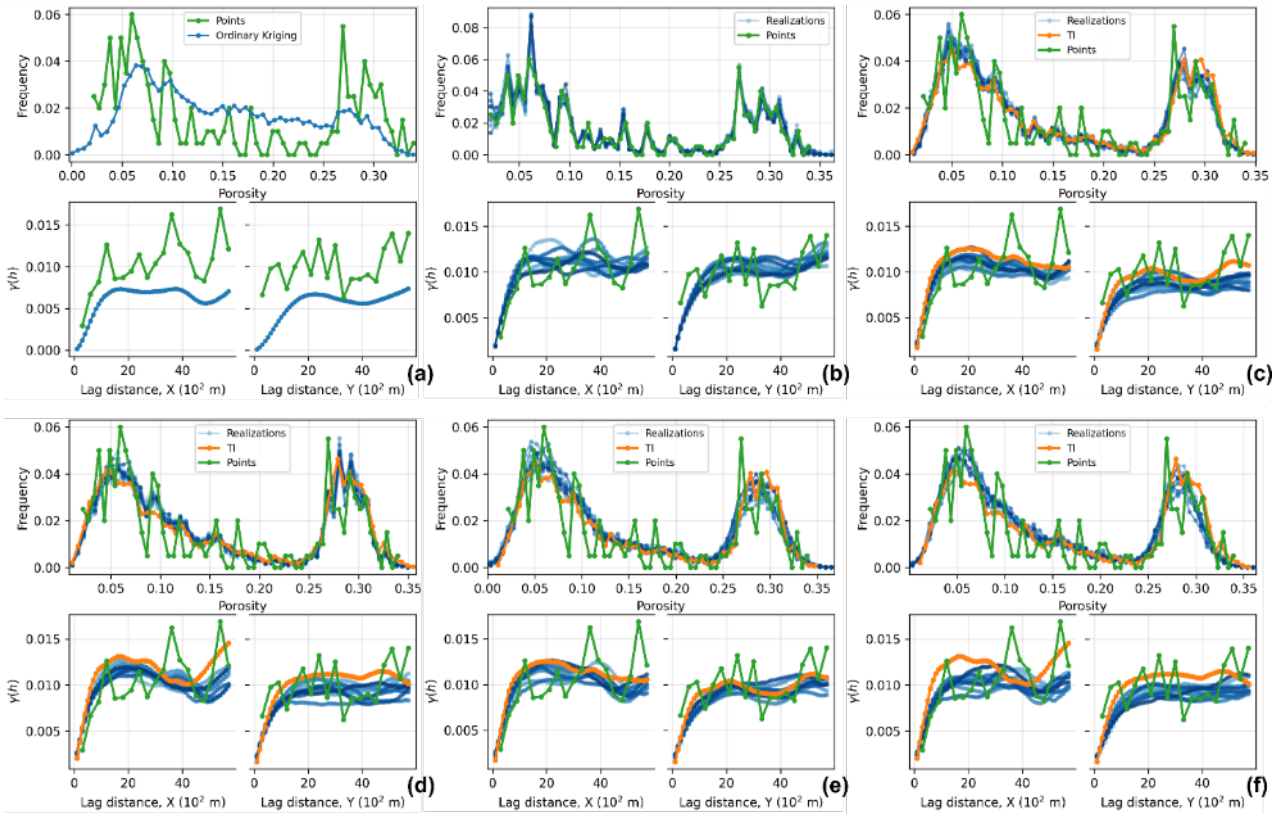


FIG 4 – Histograms (top) and directional variograms along the X and Y directions (bottom) for: (a) ordinary kriging; (b) 10 SGS realisations; (c) and (d) 10 DS realisations obtained using TI-1 and TI-2, respectively; and (e) and (f) 10 KME realisations obtained using TI-1 and TI-2, respectively.

Although the variogram responses are broadly comparable among the stochastic simulation methods, visual inspection of the realisations indicates substantial differences in the spatial organisation of the simulated structures. These differences are further investigated using a qualitative comparison of high-order spatial statistics. Third-order cumulant maps are shown in Figure 5, computed from the reference model or TI-1 (Figure 5a), TI-2 (Figure 5b), conditioning data (Figure 5c), OK (Figure 5d), and the average cumulant map from the ten realisations for each method (Figure 5e–5i). Third-order cumulants characterise asymmetric three-point spatial relationships and provide information about directional connectivity and channel organisation. The different methods are expected to reproduce, as closely as possible, the third-order cumulant map of the exhaustive reference model TI-1 (Figure 5a).

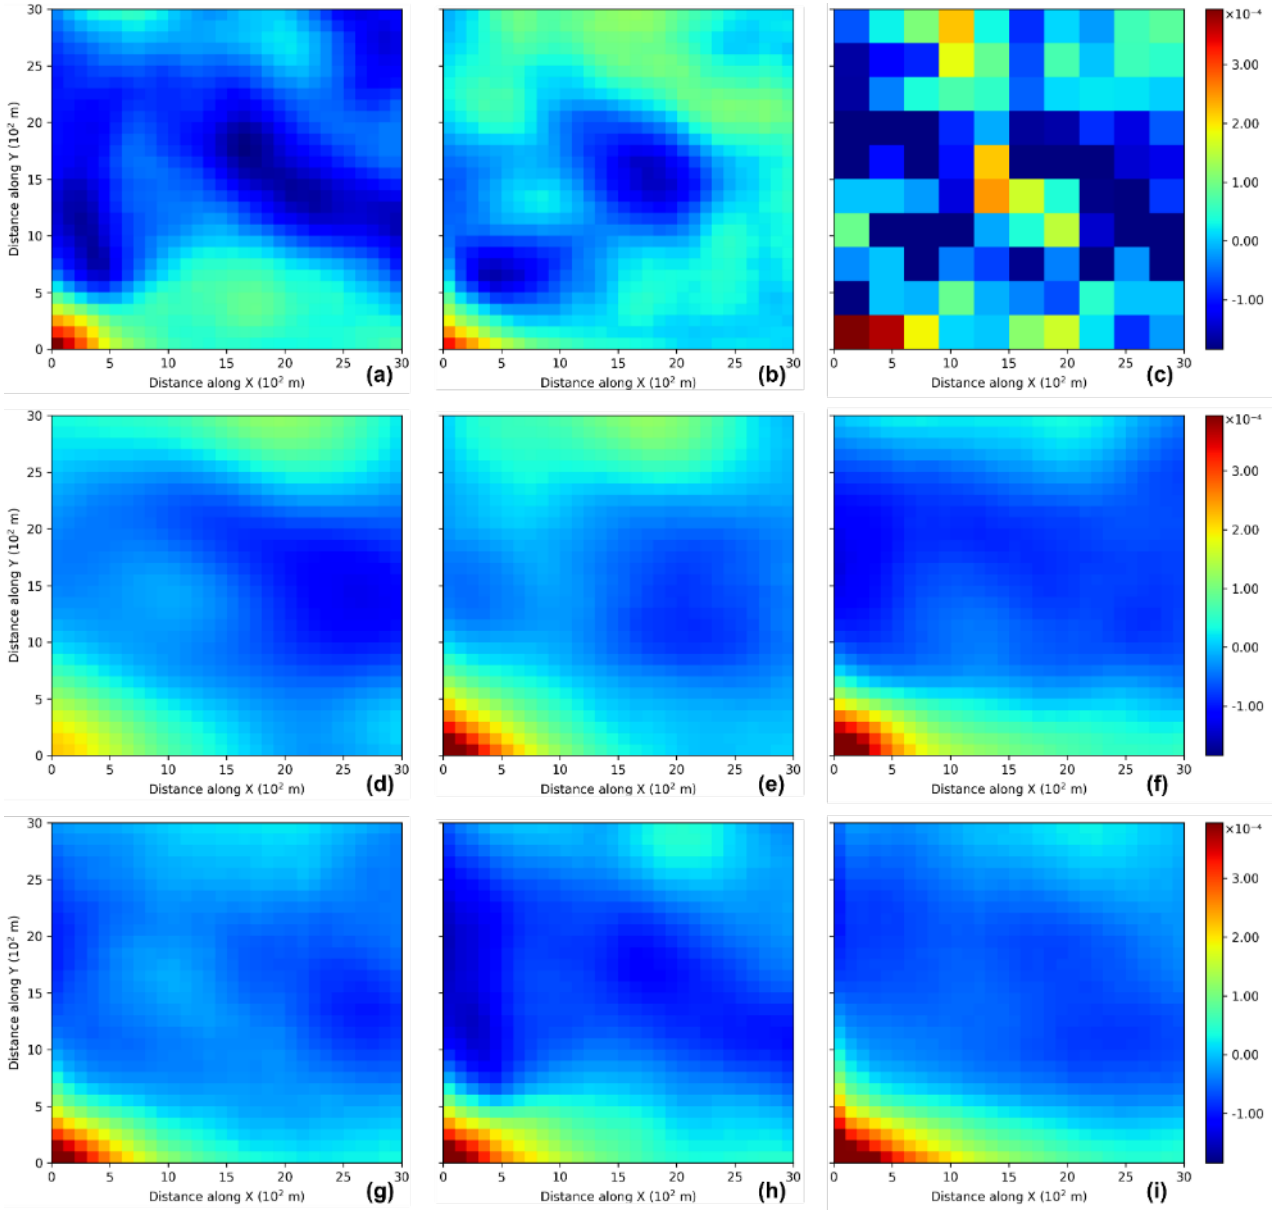


FIG 5 – Third-order cumulant map of: (a) TI-1; (b) TI-2; (c) conditioning data; (d) ordinary kriging; (e) the ten SGS realisations; (f) and (g) the ten DS realisations obtained using TI-1 and TI-2, respectively; and (h) and (i) the ten KME realisations obtained using TI-1 and TI-2, respectively.

OK fails to properly reproduce short-range and long-range three-point correlations as the result is far from the cumulant map of the reference model. The SGS realisations reproduce only part of these spatial signatures, and several characteristic features of the reference cumulant map remain attenuated or absent. The DS realisations reproduce the cumulant structures contained in the TI, which demonstrates their ability to reconstruct multipoint geological patterns. The KME realisations also produce cumulant maps that remain consistent with the geological structures encoded in TI-1. This indicates that the proposed kernel-based method successfully reproduces complex spatial relationships without explicitly computing these high-order cumulants during simulation.

Figure 6 presents the corresponding fourth-order cumulant maps, computed along the X and Y directions and the diagonal direction between the two axes. Fourth-order cumulants characterise higher-order connectivity and complex spatial organisation beyond pairwise and triplet interactions. As observed for third-order cumulants, OK and SGS reproduce only a limited portion of these structures. The DS realisations preserve the patterns contained in the TI. The KME realisations also show good agreement with the reference fourth-order cumulant structures.

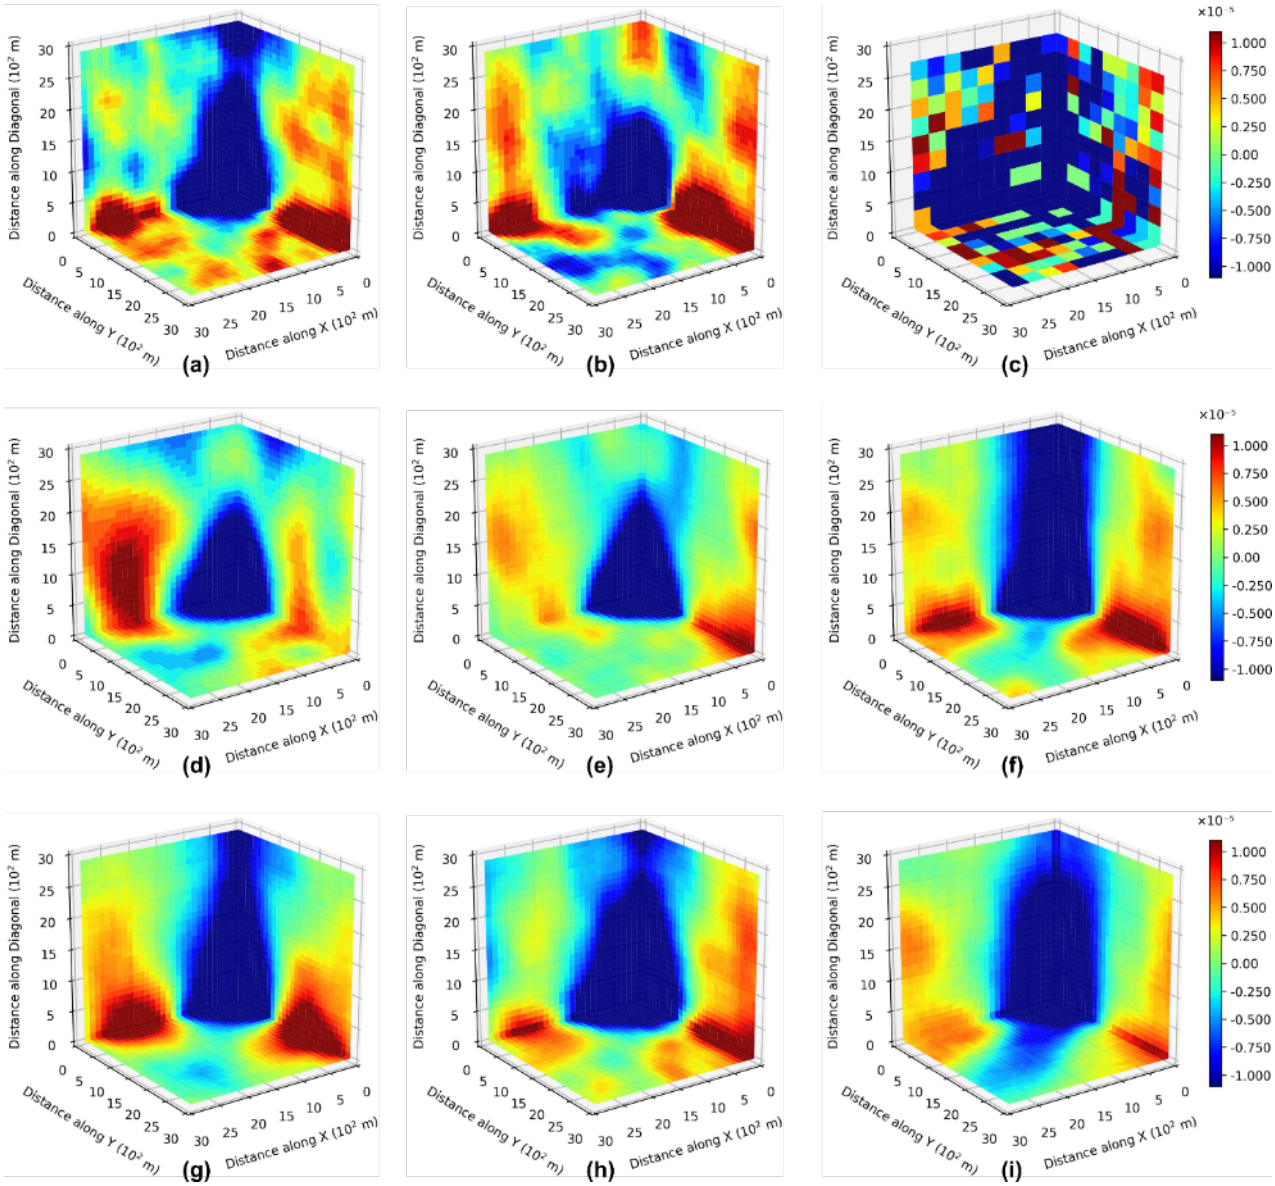


FIG 6 – Fourth-order cumulant map of: (a) TI-1; (b) TI-2; (c) conditioning data; (d) ordinary kriging; (e) the ten SGS realisations; (f) and (g) the ten DS realisations obtained using TI-1 and TI-2, respectively; and (h) and (i) the ten KME realisations obtained using TI-1 and TI-2, respectively.

Overall, the results demonstrate that, unlike conventional two-point methods such as OK and SGS, the proposed KME framework preserves complex multipoint spatial interactions and higher-order statistical structures. At the same time, the generated realisations exhibit a more coherent reproduction of geological patterns and fine-scale heterogeneity than the DS realisations considered here, while remaining consistent with both the conditioning data and the information contained in the TI.

Influence of the weighting parameter

In geostatistical simulation, MPS methods can reproduce complex multipoint spatial statistics through the use of TIs, but their performance may be affected when the spatial patterns encoded in the TI are not fully consistent with the available conditioning data. In contrast, high-order statistical approaches aim to reproduce complex spatial relationships while remaining primarily driven by the information contained in the data (Dimitrakopoulos, Mustapha and Gloaguen, 2010; Mustapha and Dimitrakopoulos, 2010). One of the key features of the proposed framework is the ability to explicitly control the relative contribution of conditioning data and TI information through the weighting coefficient λ. In practical applications, the degree of confidence assigned to a TI may vary considerably depending on the geological context and the quality of the available analogue. It is

therefore important to evaluate how the simulated realisations evolve as the balance between both information sources changes, and whether this additional flexibility can help reconcile potential conflicts between conditioning data and TI statistics.

Figure 7 illustrates the influence of λ by presenting a KME realisation for both TI-1 and TI-2 using three representative values: $\lambda = 0.1$, $\lambda = 0.5$, and $\lambda = 0.9$. For $\lambda = 0.1$ (Figure 7a and 7b), as already seen in Figure 3e and 3f with different realisations, the simulation is primarily driven by the TI. Consequently, the resulting realisations closely reproduce the channelised structures and connectivity patterns encoded in the selected TI. Increasing λ progressively increases the influence of the conditioning data. For $\lambda = 0.5$ (Figure 7c and 7d), both information sources contribute comparably to the local distribution learning process, resulting in a compromise between conditioning consistency and geological realism. Compared with $\lambda = 0.1$, the fine-scale heterogeneity inherited from the TI is reproduced less accurately, although most of the channelised structures are preserved. For $\lambda = 0.9$ (Figure 7e and 7f), the simulation process becomes predominantly guided by the data and the influence of the TI is substantially reduced. As a result, conditioning to the hard data is improved. However, the realisations obtained using TI-1 and TI-2 become increasingly similar, and progressively lose the fine-scale heterogeneity and detailed channel structures inherited from the TI. The simulated models therefore appear coarser and less geologically detailed than those obtained for smaller values of λ, as the contribution of the multipoint information provided by the TI is reduced.

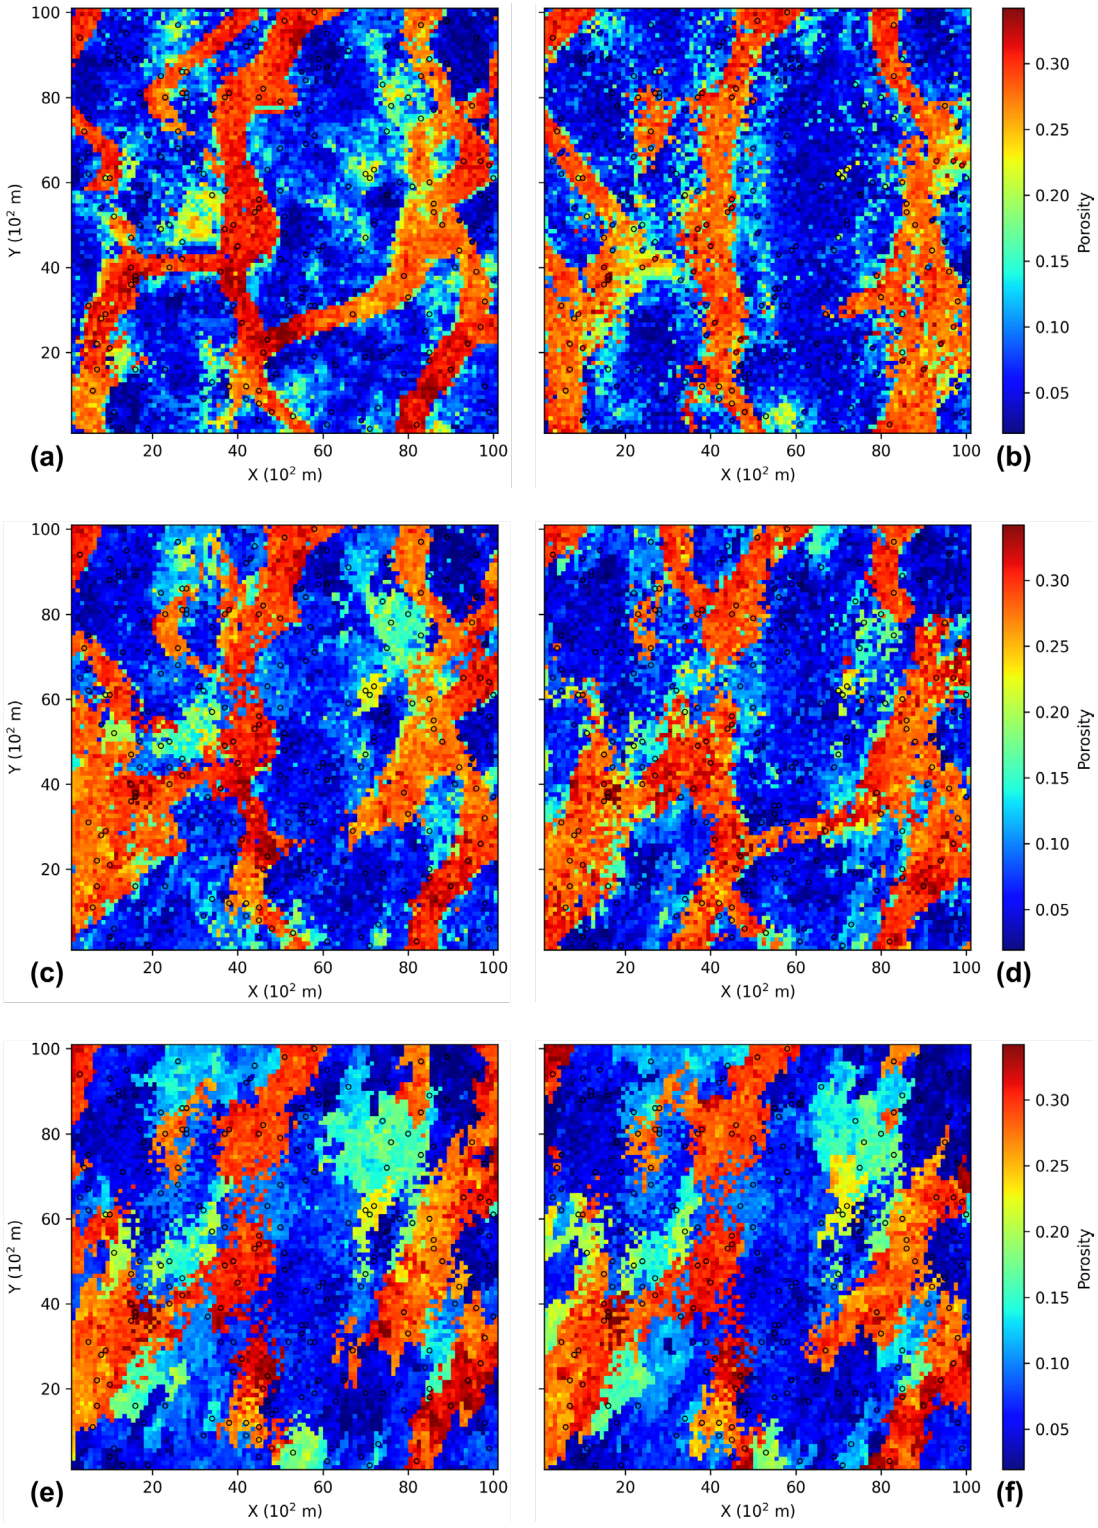


FIG 7 – KME realisations obtained using (a) and (b) $\lambda = 0.1$; (c) and (d) $\lambda = 0.5$; and (e) and (f) $\lambda = 0.9$, for TI-1 and TI-2, respectively.

Figure 8 presents the corresponding histograms (top) and directional variograms (bottom) for each weighting value, based on a set of ten realisations. The histograms remain largely unchanged for all values of λ and remain consistent with both the conditioning data and the TI. Similarly, the variogram responses remain generally stable across the tested configurations. However, increasing λ leads to a progressive reduction in the spread of the variograms computed from the ten realisations, which indicates a decrease in stochastic variability as the simulations become increasingly influenced by the hard data. Figures 9 and 10 show average third- and fourth-order cumulant maps computed from the ten realisations for each weighting coefficient. In these results, a progressive transition between

TI and data spatial statistics can be observed. For $\lambda = 0.1$, the reproduced cumulants remain close to those of the TI, which suggests a strong influence of the geological patterns encoded in the TI. As λ increases, the cumulant structures progressively converge towards similar results regardless of the TI used and gradually shift towards the cumulant maps of the conditioning data, but they remain relatively close to the high-order statistics of the exhaustive model.

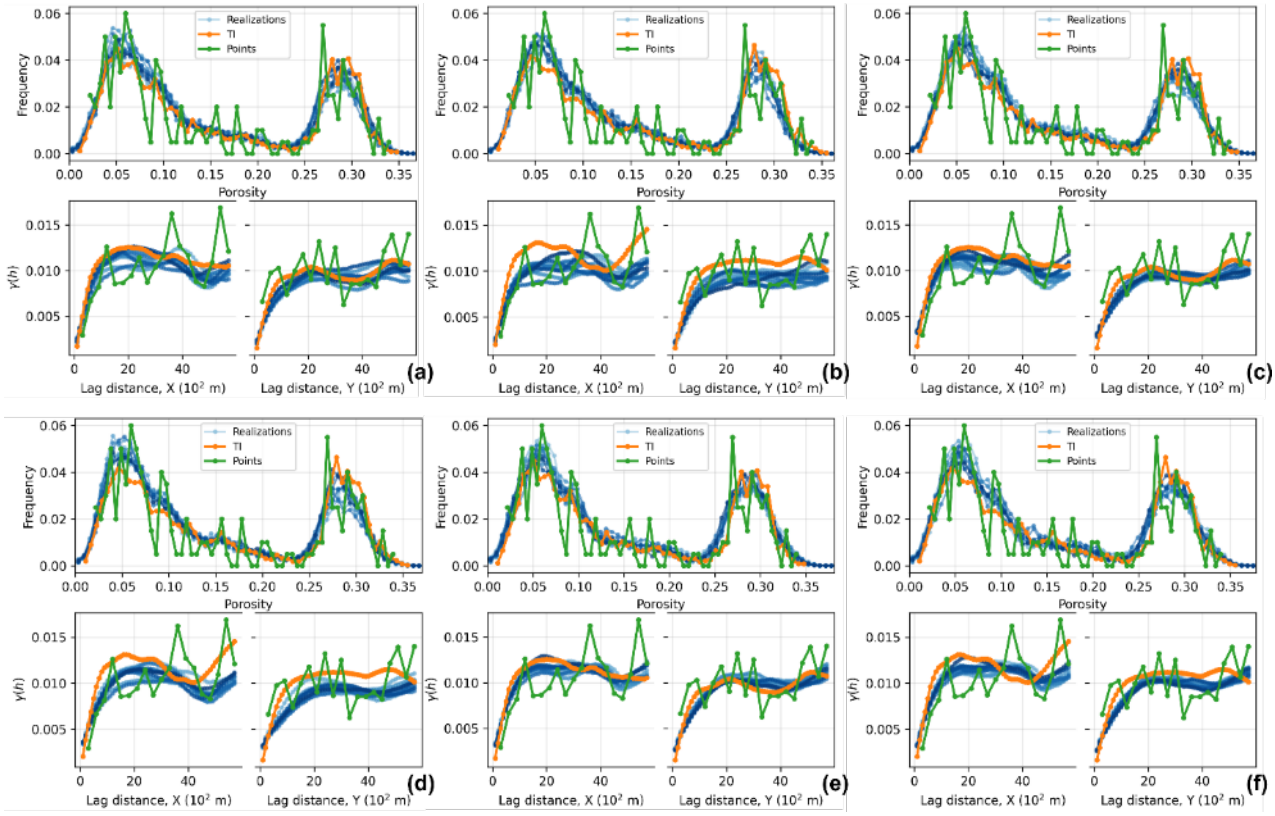


FIG 8 – Histograms (top) and directional variograms along the X and Y directions (bottom) for ten KME realisations obtained using (a) and (b) $\lambda = 0.1$; (c) and (d) $\lambda = 0.5$; and (e) and (f) $\lambda = 0.9$, for TI-1 and TI-2, respectively.

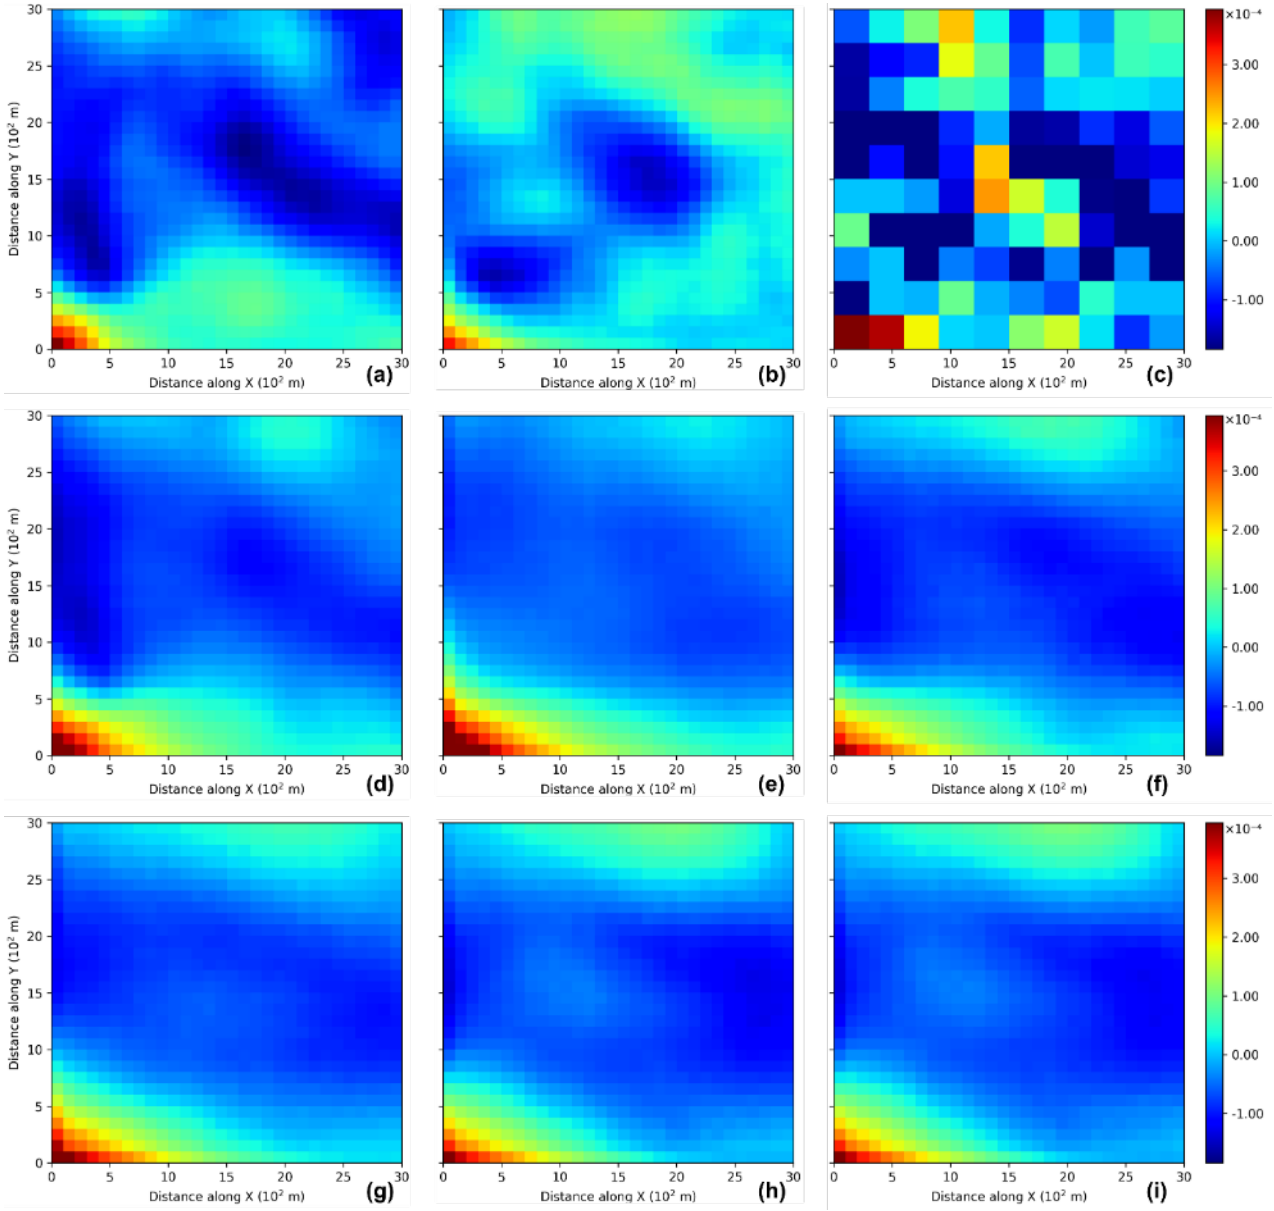


FIG 9 – Third-order cumulant map of (a) TI-1; (b) TI-2; (c) conditioning data; (d) and (e) the ten KME realisations obtained using $\lambda = 0.1$ and TI-1 and TI-2, respectively; (f) and (g) the ten KME realisations obtained using $\lambda = 0.5$ and TI-1 and TI-2, respectively; and (h) and (i) the ten KME realisations obtained using $\lambda = 0.9$ and TI-1 and TI-2, respectively.

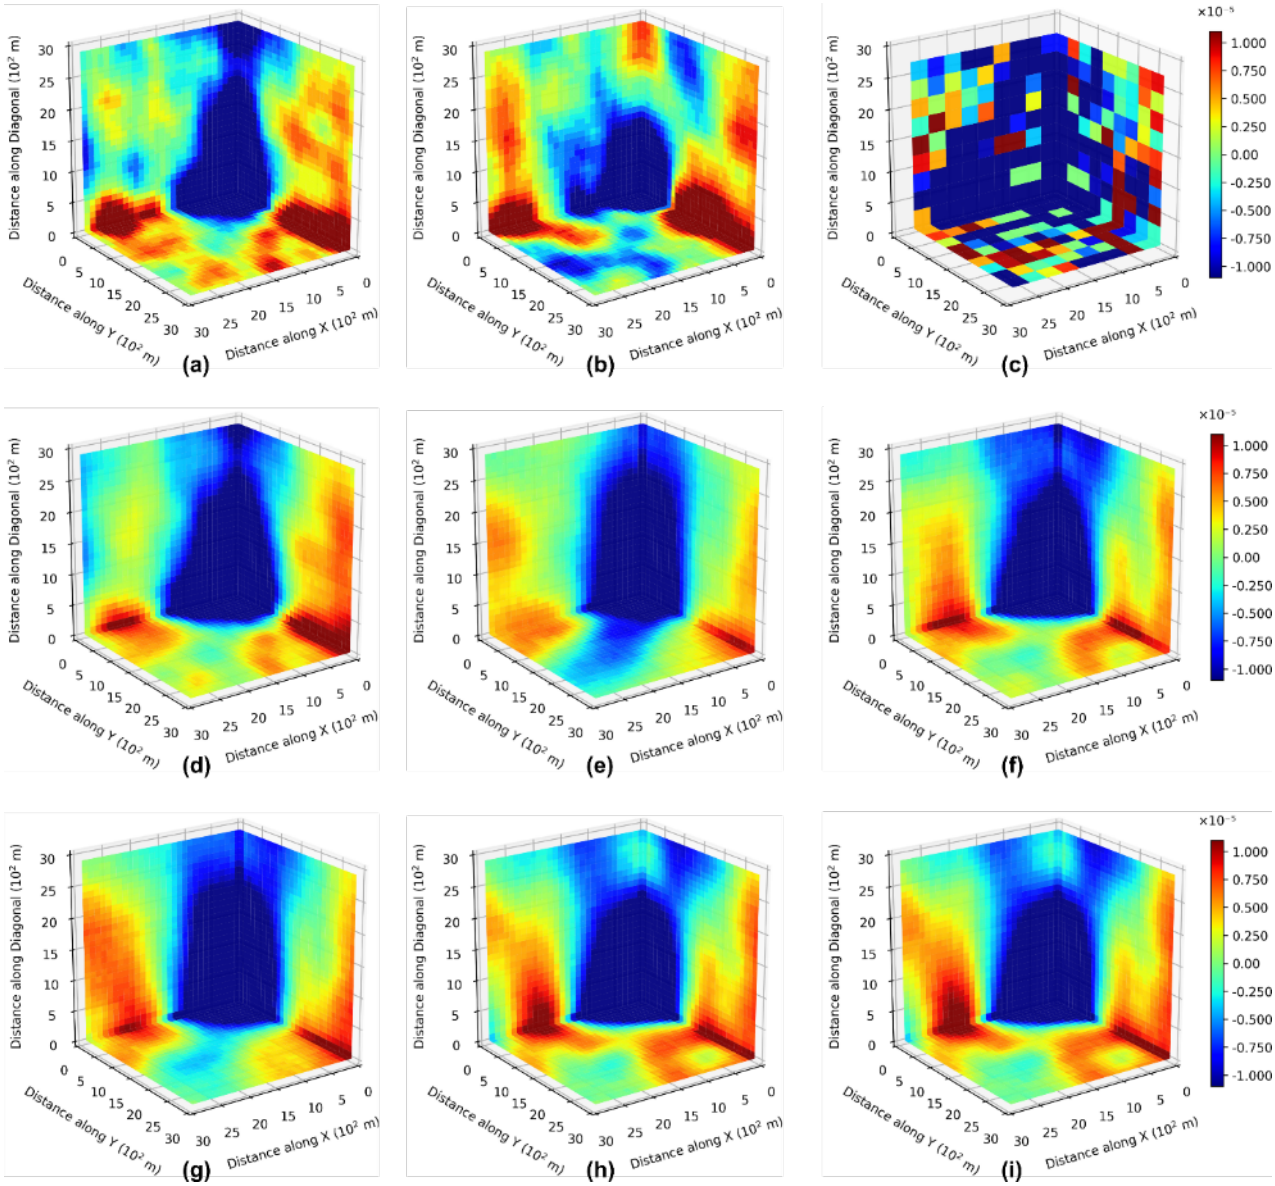


FIG 10 – Fourth-order cumulant map of (a) TI-1; (b) TI-2; (c) conditioning data; (d) and (e) the ten KME realisations obtained using $\lambda = 0.1$ and TI-1 and TI-2, respectively; (f) and (g) the ten KME realisations obtained using $\lambda = 0.5$ and TI-1 and TI-2, respectively; and (h) and (i) the ten KME realisations obtained using $\lambda = 0.9$ and TI-1 and TI-2, respectively.

These results demonstrate that λ provides a practical mechanism for managing the trade-off between conditioning consistency and the reproduction of high-order spatial structures. When the TI is considered reliable and representative of the target system, smaller values of λ are probably preferable because they preserve the multipoint statistics and complex geological architectures encoded in the TI. Conversely, when the TI is uncertain or only partially representative, larger values of λ increase the influence of the conditioning data and help mitigate potential statistical inconsistencies, albeit at the expense of the reproduction of some fine-scale heterogeneity and high-order spatial information. Hence, the choice of λ should reflect the confidence assigned to the training image relative to the conditioning data. In the present study, λ is specified by the user to illustrate its influence. Future research will investigate objective and adaptive strategies for selecting λ automatically from the spatial configuration of the conditioning data and the consistency of these data with the TI. For example, such strategies would increase the influence of conditioning data in well-informed regions while giving greater weight to the TI in sparsely informed areas.

Application on a 3D gold deposit

Data and training image

A three-dimensional gold deposit data set is used to evaluate the proposed framework under realistic orebody modelling conditions. The data set consists of 2671 composited gold-grade samples extracted from 407 exploration drill holes, with a constant composite length of 10 m. The spatial distribution of the conditioning data is shown in Figure 11a. The simulation domain is discretised on a Cartesian grid of 100 × 92 × 13 cells. Although the drill hole coverage is relatively dense, some portions of the simulation grid remain unsampled, requiring the simulation method to extrapolate gold grade values away from conditioning locations.

A TI derived from blasthole data collected in a mined-out portion of the orebody is used to provide geological information on the spatial organisation and continuity of mineralisation (Figure 11b). Following previous studies (Yao, Dimitrakopoulos and Gamache, 2020), the mined-out area is assumed to be geologically analogous to the simulation domain. SGS and the proposed KME approach are compared. For the KME simulations, a weighting coefficient of $\lambda = 0.1$ is adopted to allow for the reproduction of multipoint spatial structures from the TI.

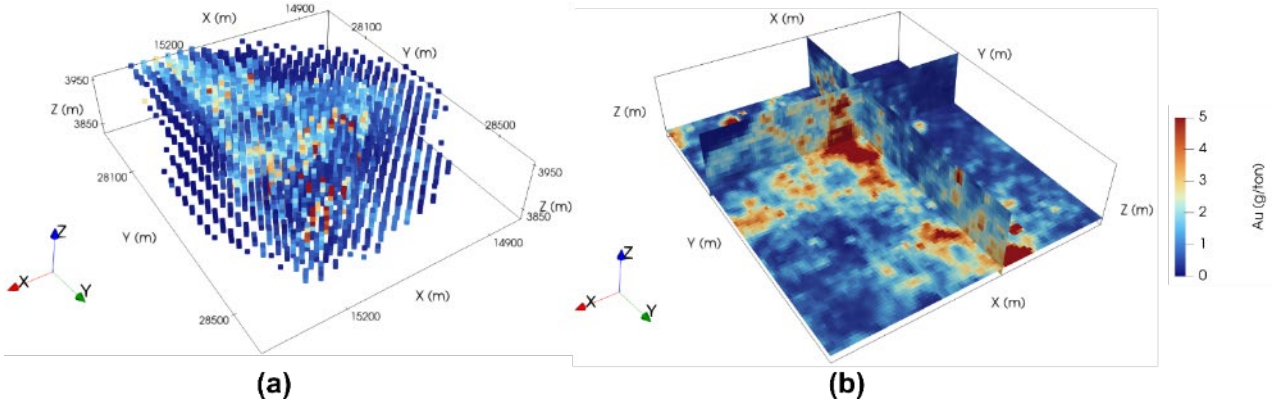


FIG 11 – (a) Conditioning gold-grade samples; (b) Cross-section of the TI.

Simulation results

Figure 12 presents cross-sections extracted from two SGS realisations (Figure 12a and 12b) and two KME realisations (Figure 12c and 12d) of the three-dimensional gold deposit model. Figure 13 compares the corresponding histograms (top) and directional variograms (bottom) of ten SGS realisations (Figure 13a) and ten KME realisations (Figure 13b) for the three axes. Both SGS and KME reproduce the global grade distribution satisfactorily. The variograms also remain globally consistent with those computed from the conditioning data and the TI, although some discrepancies are observed at larger lag distances, particularly in the vertical direction and, to a lesser extent, along the horizontal directions. These discrepancies may reflect residual non-stationarity in the conditioning data set and the TI.

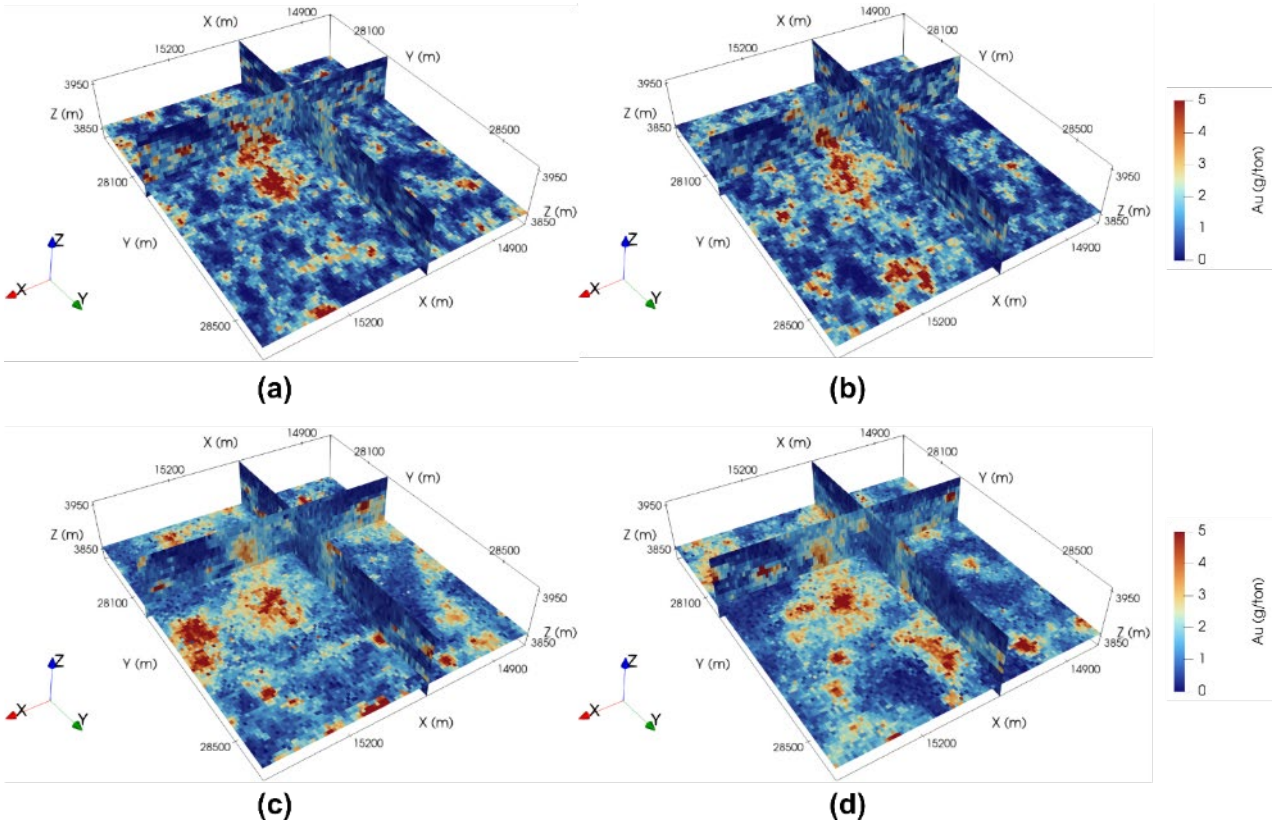


FIG 12 – Cross-sections of the gold-grade model extracted from (a) and (b) two SGS realisations; and (c) and (d) two KME realisations.

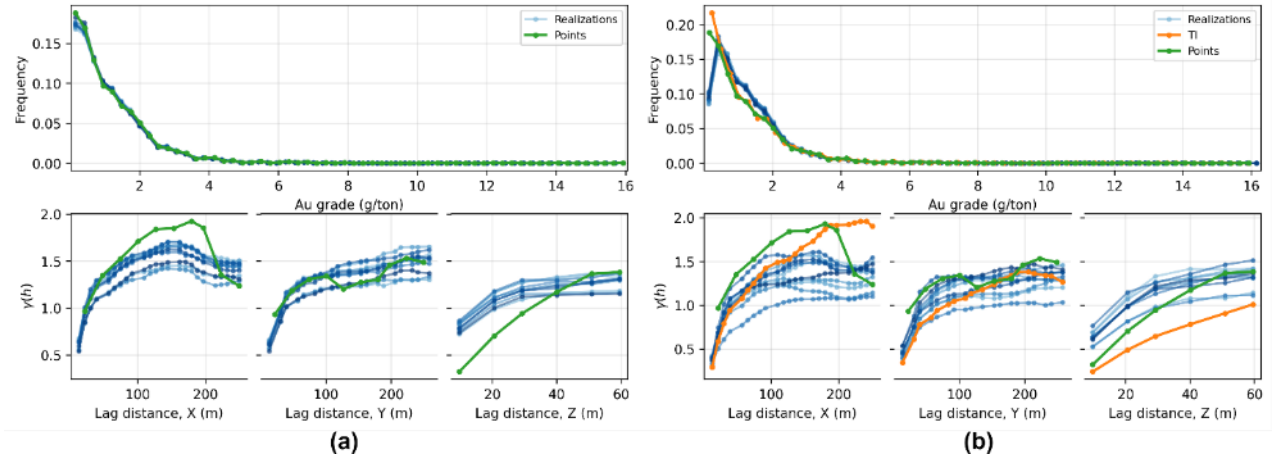


FIG 13 – Histograms (top) and directional variograms along the X, Y and Z directions (bottom) computed from: (a) ten SGS realisations; and (b) ten KME realisations.

The reproduction of high-order spatial statistics is investigated through third- and fourth-order cumulant maps. Figure 14 presents the third-order cumulants computed in the X–Y plane from the conditioning data (Figure 14a), the TI (Figure 14b), and the average third-order cumulants from the 1 ten 0 realisations generated by SGS (Figure 14c) and KME (Figure 14d). Figure 15 presents the corresponding fourth-order cumulants. Compared with the KME realisations, the SGS realisations reproduce only part of the observed spatial signatures and fail to recover several characteristic multipoint interactions, especially at short distances. In the KME realisations, the cumulant maps remain close to those derived from the hard data, indicating that, despite the contribution of the TI, the simulated high-order statistics remain consistent with those derived from the conditioning data.

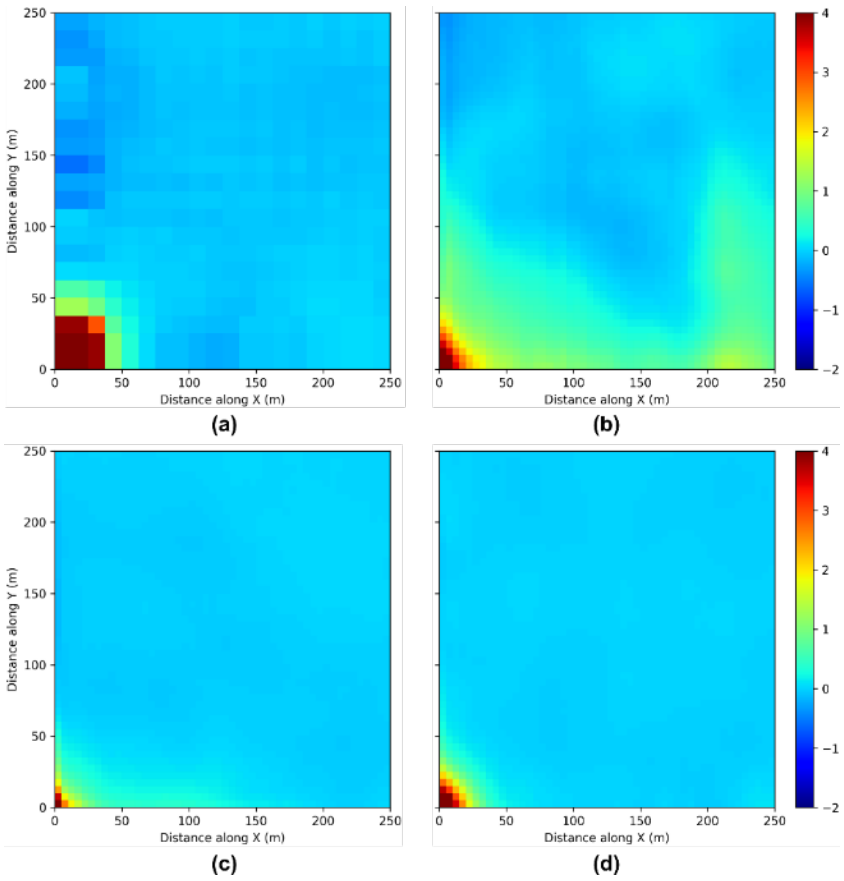


FIG 14 – Third-order spatial cumulant maps computed from: (a) the conditioning data; (b) the TI; (c) the ten SGS realisations; and (d) the ten KME realisations.

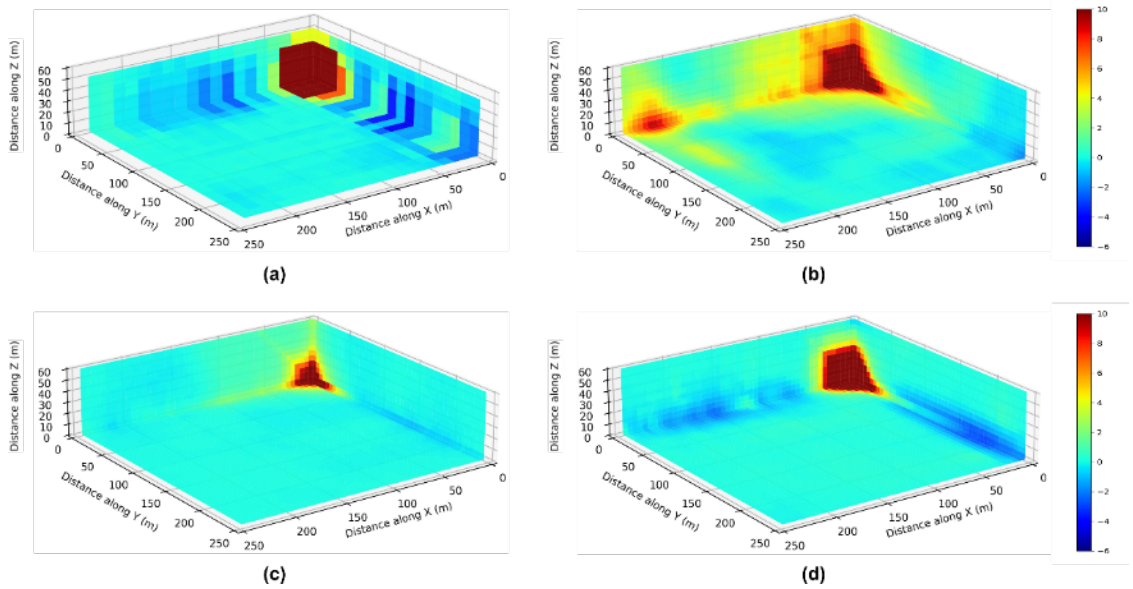


FIG 15 – Fourth-order spatial cumulant maps computed from: (a) the conditioning data; (b) the TI; (c) the ten SGS realisations; and (d) the ten KME realisations.

These results are consistent with the observations obtained from the synthetic case study and demonstrate that the proposed framework can reproduce complex three-dimensional geological continuity and preserve high-order spatial statistics under realistic orebody modelling conditions.

CONCLUSIONS

This paper presents a new high-order geostatistical simulation framework based on kernel mean embedding and Gaussian reproducing kernels. By operating in a reproducing kernel Hilbert space, the proposed approach infers local conditional probability distributions through statistical learning and reproduces high-order spatial statistics without requiring their explicit computation. The framework combines information from conditioning data and training images through a tuning parameter that allows their relative influence to be controlled during simulation.

The method was evaluated on both a synthetic Stanford V benchmark and a three-dimensional gold deposit case study. The synthetic experiments showed that the proposed approach reproduces complex channelised structures and high-order spatial statistics more effectively than conventional two-point geostatistical methods and the tested DS implementation of MPS. The results also showed that the proposed framework remains flexible with respect to the relative contribution of conditioning data and training-image information, which allows the simulation strategy to be adapted to the level of confidence assigned to each source. Application to the gold deposit confirmed that the method remains applicable under realistic orebody modelling conditions and preserves high-order spatial signatures.

These results suggest that kernel-based high-order geostatistical simulation provides a practical way for modelling complex geological structures and quantifying uncertainty in mineral resource applications. Future work could investigate alternative kernel functions, extend the framework to multivariate simulation of correlated attributes, and explore direct block-support simulation for mine planning and production forecasting workflows.

ACKNOWLEDGEMENTS

This work was funded by the Natural Sciences and Engineering Research Council (NSERC) of Canada NSERC Discovery Grant 239019, the COSMO Mining Industry Consortium (AngloGold Ashanti, BHP, De Beers, AngloAmerican, IAMGOLD, Kinross Gold, Newmont Mining, Agnico Eagle, and Vale), and the Canada Research Chairs Program.

REFERENCES

Asghari, O, Soltni, F and Amnieh, H B, 2009. The comparison between sequential Gaussian simulation (SGS) of Choghart ore deposit and geostatistical estimation through ordinary kriging, *Australian Journal of Basic and Applied Sciences*, 3(1):330–341.

Chilès, J-P and Delfiner, P, 2012. *Geostatistics: Modeling Spatial Uncertainty*, 2nd edn, Wiley Series in Probability and Statistics, 699 p (John Wiley and Sons).

Dimitrakopoulos, R, Mustapha, H and Gloaguen, E, 2010. High-order statistics of spatial random fields: exploring spatial cumulants for modeling complex non-Gaussian and non-linear phenomena, *Mathematical Geosciences*, 42(1):65–99.

Goodfellow, R, Albor Consuegra, F, Dimitrakopoulos, R and Lloyd, T, 2012. Quantifying multi-element and volumetric uncertainty, Coleman McCreedy deposit, Ontario, Canada, *Computers and Geosciences*, 42:71–78.

Goovaerts, P, 1997. *Geostatistics for Natural Resources Evaluation*, Applied Geostatistics Series, 483 p (Oxford University Press: New York).

Gretton, A, Borgwardt, K M, Rasch, M J, Schölkopf, B and Smola, A, 2012. A kernel two-sample test, *Journal of Machine Learning Research*, 13(25):723–773.

Guardiano, F B and Srivastava, R M, 1993. Multivariate geostatistics: beyond bivariate moments, in *Geostatistics Tróia '92*, pp 133–144 (Springer).

Journel, A G and Huijbregts, C J, 1976. *Mining Geostatistics* (Academic Press: London).

Journel, A G, 2005. Beyond covariance: the advent of multiple-point geostatistics, in *Geostatistics Banff 2004* (eds: O Leuangthong and C V Deutsch), pp 225–233 (Springer).

Mao, S and Journel, A G, 1999. Generation of a reference petrophysical/seismic data set: the Stanford V reservoir, SCRF Report, Stanford University.

Mariethoz, G, Renard, P and Straubhaar, J, 2010. The direct sampling method to perform multiple-point geostatistical simulations, *Water Resources Research*, 46(11).

Minniakhmetov, I, Dimitrakopoulos, R and Godoy, M, 2018. High-order spatial simulation using Legendre-like orthogonal splines, *Mathematical Geosciences*, 50(7):753–780.

Mustapha, H and Dimitrakopoulos, R, 2010. High-order stochastic simulation of complex spatially distributed natural phenomena, *Mathematical Geosciences*, 42(5):457–485.

Mustapha, H and Dimitrakopoulos, R, 2011. HOSIM: A high-order stochastic simulation algorithm for generating three-dimensional complex geological patterns, *Computers and Geosciences*, 37(9):1242–1253.

Osterholt, V and Dimitrakopoulos, R, 2007. Simulation of wireframes and geometric features with multiple-point techniques: application at Yandi iron ore deposit, *Orebody Modelling and Strategic Mine Planning*, 2nd edn, Spectrum Series, 14:95–124 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Smola, A, Gretton, A, Song, L and Schölkopf, B, 2007. A Hilbert space embedding for distributions, in *Algorithmic Learning Theory*, pp 13–31 (Springer: Berlin Heidelberg).

Song, L, Zhang, X, Smola, A, Gretton, A and Schölkopf, B, 2008. Tailoring density estimation via reproducing kernel moment matching, in *Proceedings of the 25th International Conference on Machine Learning*, pp 992–999.

Strebelle, S, 2002. Conditional simulation of complex geological structures using multiple-point statistics, *Mathematical Geology*, 34(1):1–21.

Yao, L, Dimitrakopoulos, R and Gamache, M, 2018. A new computational model of high-order stochastic simulation based on spatial Legendre moments, *Mathematical Geosciences*, 50(8):929–960.

Yao, L, Dimitrakopoulos, R and Gamache, M, 2020. High-order sequential simulation via statistical learning in reproducing kernel Hilbert space, *Mathematical Geosciences*, 52(5):693–723.

Yao, L, Dimitrakopoulos, R and Gamache, M, 2021. Learning high-order spatial statistics at multiple scales: a kernel-based stochastic simulation algorithm and its implementation, *Computers and Geosciences*, 149:104702.

Dynamic structural modelling for near real time inputs for geotechnical risk management and optimisation

A Jani[1]

1. Specialist Structural Geologist, Rio Tinto Iron Ore, Perth WA 6000.
   Email: akash.jani@riotinto.com

ABSTRACT

Structural complexities such as folding, faulting, jointing, and shear zones heavily influence the stability of open pit slopes in iron ore mining. As orebodies become more complex and business goals of de-carbonisation and heritage responsibility grows, technology-driven solutions are essential for increasing predictability and adaptability in mining operations. Achieving optimal pit design requires a thorough understanding and modelling of geological structures. However, managing geological uncertainty and variations in orebody geometry remains challenging. Pit face mapping generally employed to provide input into the uncertainty has very long lead times and is high risk. This often results in a reactive approach as mining progresses, leading to unforeseen geotechnical risks or ore sterilisation.

Traditional approaches that rely primarily on drill hole data and resource modelling fall short of capturing the structural detail needed to identify uncertainty and complex structural domains. As such, a pit-scale structural model is necessary for effective slope design and management.

By employing automated ‘drone in a box’ (DIB) UAVs (unmanned aerial vehicles) for data capture and advanced implicit modelling, detailed, local-scale pit models can be created to better address these uncertainties using predictive analysis. Pit face mapping, conducted safely and efficiently using UAVs, provides data for dynamic structural models, which work in near real-time and can integrate orientation data to predict kinematic behaviours. This enhances safety and operational efficiency.

Structural domains informed by this pit mapping, along with resource model data, allow for accurate predictions of rock behaviour under varying extraction conditions. These models are continuously updated as new mapping and drill hole data becomes available, reflecting real-time geology and reducing the risk of encountering unexpected ground conditions, which boosts safety and ore recovery.

Incorporating these technologies, alongside predictive analytics, significantly reduces operational uncertainties, optimises resource extraction, and lowers costs by automating repetitive tasks and reducing human exposure to hazardous areas.

INTRODUCTION

This study focuses on enhancing the stability of open pit slopes in iron ore mining by addressing the challenges of geological uncertainty, structural complexities, and the need for technology-driven solutions. The principal objective is to improve pit design, safety, and ore recovery by utilising automated UAVs (unmanned aerial vehicles / ‘drone in a box’) and advanced implicit modelling. Methods include UAV-based pit face mapping, real-time structural modelling, and predictive analysis. Results show that this approach reduces operational uncertainties, enhances safety, optimises resource extraction, and minimises risks such as unforeseen geotechnical issues and ore sterilisation, thereby lowering costs and improving operational efficiency.

METHODS AND RESULTS

A large part of data collection is done in the pit including pit face mapping and reverse circulation (RC) grade control drilling which feed into the resource and mining models respectively. However, this study largely focuses on remote work techniques and implicit modelling which assist in providing near real time information for structural geological uncertainty.

Structure from motion (SfM) photogrammetry is a commonly adopted technique in mining operations to remotely capture data from inaccessible or high-risk locations, Dey, Roy and Matin (2021). A drone (or UAV) is typically deployed to photograph the highwall, and through computer vision techniques,

a 3D mesh is generated. This mesh can then be used for geological interpretation and orientation measurements. Recent advancements in UAV technology have led to the implementation of a ‘drone in a box’ (DIB) solution, which allows for automated, regular data capture as mining progresses.

An implicit model is set-up in Leapfrog Geo using the latest resource model data including all mapping and drilling data sets. The model build should include a structural surface which can then be used to inform changes in the overall geometry of the stratigraphy in the pit. Any structural uncertainty or gaps can be identified at this stage and flagged for data capture using photogrammetry.

As described by Thiele *et al* (2017), the 3D mesh is interpreted using a combination of automated facet extraction and manual plane picking techniques to identify structural features. The orientation and geometry information gathered from this interpretation is then integrated back into the implicit model to update the structural model, adding explicit controls as needed.

The next step is to address any probabilities of failure based on various kinematic failure mechanisms. This process can be employed directly in the implicit model process using the workflow from the webinar on Seequent (2020). This allows us to assess in near real time any probability of failure based on the latest information available (Figure 1). This process also allows to predict areas of uncertainty or risk based on the current data set and where the highest priority of data capture will be as mining progresses.

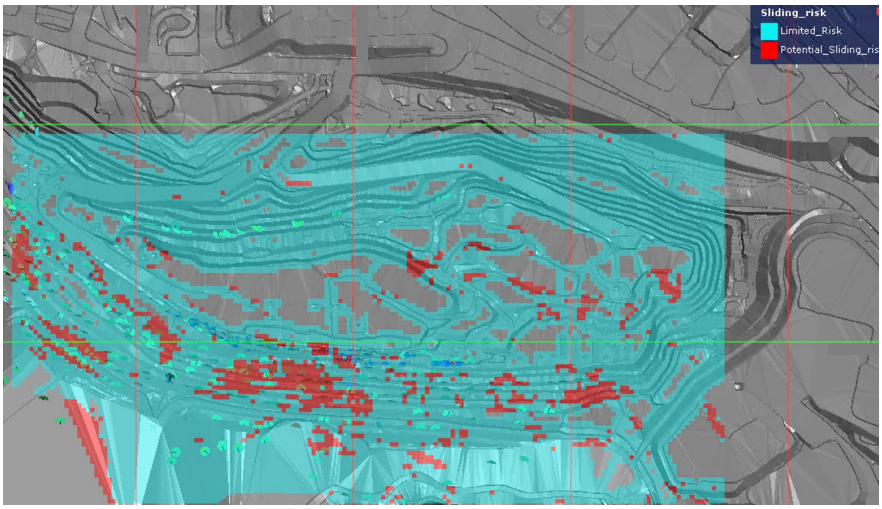


FIG 1 – Areas of pit highlighted based on probability of sliding failure risk where red is potential risk and blue is limited risk.

CONCLUSIONS

Dynamic updates to the structural model provide an efficient way to define and assess geological uncertainty, while also contributing valuable input to geotechnical stability analysis. As a remote method, it eliminates the need for personnel in high-risk environments, minimising disruption to ongoing mining operations. Additionally, it aids in proactively identifying data gaps and potential risks based on the most current data sets. By incorporating probabilistic kinematic failure mechanisms, this approach further assists in predicting potential failure modes. This input is essential for geotechnical risk assessments, as it helps identify and flag areas of high risk and uncertainty.

REFERENCES

Dey, J, Roy, S and Matin, A, 2021. Drone photogrammetry: a structural data gathering tool for open pit mining geotechnics, SRK Mining Services (India) Pvt Ltd, India, University of Calcutta, India.

Seequent, 2020. Improving geotechnical stability analysis using Leapfrog Geo: Edge Partner Webinar, SRK Consulting [online]. Available from: <https://www.seequent.com/improving-geotechnical-stability-analysis-using-leapfrog-geo-edge-partner-webinar-srk-consulting>

Thiele, S T, Grose, L, Cui, T, Cruden, A R and Micklethwaite, S, 2017. Extraction of high-resolution structural orientations from digital data: A Bayesian approach, School of Earth, Atmosphere and Environment, Monash University, Melbourne, Australia.

Generative geostatistical simulation using Restricted Boltzmann Machine, comparison with VAEs and GANs

Y Mu[1] and F Ferrie[2]

1. Postdoctoral Fellow, McGill University, Montreal QC H3A 2A7, Canada. Email: yanyan.mu@mcgill.ca
2. Professor, McGill University, Montreal QC H3A 2A7, Canada. Email: ferrie@cim.mcgill.ca

ABSTRACT

While modern generative models such as Variational Autoencoders (VAEs) and Generative Adversarial Networks (GANs) have been introduced to the mining industry for simulating complex geological patterns, they often lack the precision required for rigorous orebody modelling. These frameworks were originally designed for general image processing, where the objective is to produce visually plausible results rather than adhering to strict, coordinate-specific data. In a mining context, however, the simulation must move beyond visual similarity to achieve robust 'conditioning'—the critical requirement of ensuring the model strictly honours every borehole measurement and local observation. Because VAEs and GANs prioritise global patterns over local data rigidity, they frequently struggle to anchor their realisations to specific spatial coordinates. To force these models to match physical data, engineers must often resort to complex auxiliary networks or extensive extra training, resulting in a workflow that is both time-consuming and prone to error in data-constrained environments.

This paper presents RBMSim, a specialised simulation approach based on Restricted Boltzmann Machines (RBMs) that overcomes these operational hurdles. Unlike standard deep learning models, RBMSim utilises a unified architecture that handles both large-scale geological patterns and local borehole constraints simultaneously. This is achieved by creating a 'Digital Twin' of the physical site data directly within the model's input layer. By 'clamping' known borehole measurements into the network, the system naturally honours every piece of hard data as an untouchable physical constraint, without requiring the 'tricks,' specialised tuning, or secondary training processes common in VAEs and GANs. This architecture ensures that the physical reality of the deposit remains the 'source of truth,' providing a more reliable and streamlined workflow for geologists and mining engineers.

To demonstrate its practical effectiveness, we compared RBMSim against VAEs and GANs using a training image from an active copper mine. The results confirm that RBMSim is significantly more effective at reproducing high-order spatial statistics, specifically the connectivity and physical shape of orebodies as measured via spatial cumulants and moments. By providing a more accurate representation of geological continuity while strictly honouring local conditioning data, RBMSim offers a robust and user-friendly standard for quantifying geological uncertainty. This research provides mining teams with a dependable tool for resource estimation that bridges the gap between sophisticated generative modelling and the practical realities of mine site data.

INTRODUCTION

Stochastic geostatistical simulation plays a paramount role in modern mine planning and mineral resource estimation (Dimitrakopoulos, 2011). Accurately quantifying geological uncertainty is critical for mitigating financial risks associated with the spatial distribution, grade variability, and connectivity of orebodies (Journel and Huijbregts, 1978). Traditionally, two-point geostatistical methods, such as sequential Gaussian simulation (SGS) and sequential indicator simulation (SIS) (Deutsch and Journel, 1998), have served as the industry standards. These classical approaches rely heavily on variograms, which are spatial covariance functions limited to capturing only linear correlations between pairs of spatial locations. Consequently, they are inherently incapable of reproducing complex, non-Gaussian geological features, such as continuous curvilinear structures, sharp lithological boundaries, and complex sinusoidal ore deposit geometries (Caers, 2011).

To overcome the rigid limitations of variogram-based techniques, Multiple-Point Geostatistics (MPS) was introduced as a major paradigm shift in mathematical geosciences, pioneered by the

development of the SNESIM algorithm (Strebelle, 2002). Instead of relying on analytical spatial functions, MPS utilises a training image (TI) as a conceptual, non-parametric standard to explicitly provide a prior repository of geological patterns (Guardiano and Srivastava, 1993). While SNESIM successfully captured complex discrete categories using an internal search-tree data structure (Strebelle, 2002), it suffered from severe RAM and computational overhead when scaling to large, continuous, or multi-variate data templates. To mitigate this computational bottleneck, subsequent pattern-based MPS methods, such as FILTERSIM (Zhang, Switzer and Journel, 2006) and WAVESIM (Chatterjee and Dimitrakopoulos, 2011) were developed to compress pattern configurations into low-dimensional feature spaces using filters. Despite these algorithmic refinements, traditional MPS frameworks often suffer from a fundamental limitation: they lack a rigorous, unified probability density function to govern the stochastic sampling process, making them highly sensitive to the ad hoc configuration of simulation paths and random seed inputs.

Concurrently, within the domain of stochastic geomodelling, researchers recognised the need for high-order spatial simulation frameworks that could quantify complex spatial continuity beyond simple two-point configurations. High-order simulation methods seek to mathematically characterise non-linear continuous spatial phenomena by utilising spatial cumulants and multi-point moments (Dimitrakopoulos, Mustapha and Gamafi, 2010). Spatial cumulants act as the mathematical extension of variograms to three, four, or N spatial points, explicitly measuring the multi-point connectivity and geometric topology of orebodies. Although these high-order formulations provide a mathematically rigorous, distribution-free framework for capturing non-Gaussian curvilinear attributes (Dimitrakopoulos, Mustapha and Gamafi, 2010; Mustapha and Dimitrakopoulos, 2011), the empirical estimation from sparse field data presented immense computational challenges. To accommodate highly complex curvilinear patterns under realistic, non-stationary earth environments, subsequent advancements introduced calibrated pattern importance weights within multi-grid approaches (Haji Abolhassani *et al,* 2022).

Furthermore, to establish a formal learning set-up, high-order sequential simulation was elevated to a statistical learning framework within a Reproducing Kernel Hilbert Space (RKHS), systematically incorporating spatial Legendre moments (SLM) via specialised kernels (Yao and Dimitrakopoulos, 2020). This learning-based trajectory eventually enabled training-image-free high-order spatial inference through aggregated kernel statistics, allowing sparse field data to directly infer the conditional probability distribution (Yao, Dimitrakopoulos and Gamache, 2021). More recently, these high-order statistics have been deployed to dynamically update multi-attribute orebody models across varying production scales, integrating point-support blasthole assays and non-linear mill performance data into production planning (Findlay and Dimitrakopoulos, 2026). Despite the mathematical elegance of these framework updates, the combinatorial complexities and the heavy reliance on complex kernel functions or multi-point configuration matching still present high computational thresholds.

In recent years, the rapid advancement of deep learning has introduced powerful generative deep learning architectures to the mining and geostatistical communities, promising an implicit, data-driven alternative to pattern learning. Prominent among these are Variational Autoencoders (VAEs) and Generative Adversarial Networks (GANs) (Laloy *et al,* 2017, 2018; Moss, Srinivasan and Caers, 2020). VAE simulation frameworks optimise an explicit lower bound on the log-likelihood of the data by mapping complex, high-dimensional geological images onto a continuous, lower-dimensional latent space governed by a standard Gaussian distribution (Kingma and Welling, 2013). This approach was successfully adapted for airborne electromagnetic data inversion and spatial pattern replication by Laloy *et al* (2017). Once trained, new geological realisations can be generated by sampling from this continuous latent manifold. However, the latent structure of VAEs introduces a pervasive ‘blurry limitation’ (Laloy *et al,* 2017). The Gaussian smoothing objective function within VAEs acts as a spatial low-pass filter, which systematically erases sharp geological contacts, abrupt faults, and narrow, high-grade vein structures.

Conversely, GAN simulation architectures bypass explicit likelihood estimation by framing the generative process as a minimax game between a generator network and a discriminator network (Goodfellow *et al,* 2014). As demonstrated by Laloy *et al* (2018) through Spatial GANs, and more recently extended to deep convolutional networks for modelling categorical variables and lithological facies (Hadavand and Deutsch, 2026), GANs are highly proficient at capturing crisp, visually

stunning, and continuous macro-scale geological features. However, GANs are notoriously unstable during training and are highly susceptible to 'mode collapse.' In a mining context, mode collapse causes the generator to selectively reproduce only a narrow subset of dominant spatial patterns from the training imagery (Laloy *et al,* 2018; Moss, Srinivasan and Caers, 2020), completely erasing the stochastic variety and the vital, low-probability but high-impact scenarios required for an unbiased conditional simulation ensemble.

Furthermore, a shared, catastrophic deficiency in both VAE (Laloy *et al,* 2017) and GAN frameworks—including state-of-the-art applications (Hadavand and Deutsch, 2026)—is their inability to naturally ingest sparse, coordinate-specific physical constraints (hard data), such as borehole samples. Because these networks prioritise continuous global topologies over local data rigidity, they frequently struggle to anchor their realisations to specific spatial coordinates. To force GANs or VAEs to honour drill hole configurations, researchers must resort to integrating conditioning data directly into specialised adversarial training regimes, using complex auxiliary loss terms or secondary optimisation loops (Hadavand and Deutsch, 2026; Dupont *et al,* 2018). This ad hoc intervention inherently distorts the learned underlying probability space, yielding a 'pseudo-conditional' or 'surrogate' distribution that fails to reflect the true geostatistical posterior distribution.

To bridge the gap between sophisticated generative modelling and the rigid, practical data constraints of mine sites, this paper proposes an updated and streamlined generative geostatistical framework called RBMSim, rooted in the mathematics of Restricted Boltzmann Machines (RBMs) (Hinton, 2012). Originally conceptualised as a two-phase learning model for sparse image reconstruction in mine planning (Mu, Ferrie and Dimitrakopoulos, 2015), where hidden layer biases were localised to match sparse core samples, we extend this foundational work into a unified, single-hidden-layer generative model. Unlike black-box deep networks (VAEs/GANs) that isolate features within continuous latent spaces, the RBM is inherently an energy-based model whose discrete mathematical formulation shares an elegant equivalence with the classic Ising model in statistical physics (Ackley, Hinton and Sejnowski, 1985). This unique property provides unparalleled structural interpretability, mapping complex geological connectivity directly to global energy states.

Crucially, because the RBM characterises and preserves the full joint distribution over the discrete spatial pattern configuration space, transitioning from an unconditional to a conditional simulation requires zero architectural modifications or secondary optimisation training loops—a significant advancement over the two-phase implementation in Mu, Ferrie and Dimitrakopoulos (2015). By partitioning the visible layer into conditioned spatial locations v_{cond} and unconditioned simulation locations v_{sim}, known borehole measurements can be directly 'clamped' into the input layer as an untouchable physical constraint. The conditional distribution $p(v_{sim}|v_{cond})$ thus operates as a mathematically exact slice (cross-section) of the learned joint distribution $p(v,h)$, ensuring that the physical reality of the deposit remains the uncompromised 'source of truth.'

Using a real-world training image from an active copper mine, we demonstrate that RBMSim significantly outperforms VAEs and GANs in reproducing high-order spatial statistics—specifically the connectivity and physical shape of orebodies as measured via spatial cumulants and moments—while establishing a robust, reproducible, and unbiased standard for Geological Uncertainty Quantification (UQ).

MODEL AND LEARNING ALGORITHM

To establish a rigorous mathematical foundation for generative geostatistical modelling, this section formalises the architectures, probability distributions, and optimisation algorithms of the three primary generative frameworks evaluated in this study: Restricted Boltzmann Machines (RBMs), Variational Autoencoders (VAEs), and Generative Adversarial Networks (GANs).

- Notation and terminology:

  Activation of visible units v_i, activation of the hidden units h_i, bias terms associated with the visible and hidden units respectively b_i, b_j,weights of connections between the visible and hidden units w_{ij}, joint energy of the network for the configuration of visible and hidden units $E(v,h)$, joint distribution of the network for the configuration of visible and hidden units $p(v,h)$,

marginalised distribution of visible/hidden units respectively $p(v), p(h)$, learning rate of gradient descent ϵ.

- Restricted Boltzmann Machine:

  Restricted Boltzmann Machine (RBM) is a generative model with a bipartite structure. RBM is analogous to a Markov Random Field (MRF), it has two layers, a visible layer and a hidden layer. The units between the two layers are fully connected. but different from MRF that RBM has no lateral connections within each layer. RBM is also analogous to regular feed-forward neural network with one layer of binary stochastic hidden units and tied weights.

An RBM is also an energy-based model. The energy over the configuration over all visible and hidden units is defined by Equation 1; the probability distribution which is associated with the energy is defined by Equation 2; and the normalising factor Z is the sum over all configurations and called the partition function (Equation 3).

$$E(v,h) = -\sum_{i,j} w_{i,j}\, v_i h_j - \sum_i b_i\, v_i - \sum_j b_j\, h_j \tag{1}$$

$$p(v,h) = \frac{e^{-E(v,h)}}{Z} \tag{2}$$

$$Z = \sum_{v,h} e^{-E(v,h)} \tag{3}$$

The marginalised distribution of visible units (input layer) is factored by Equations 4 and 5, to simplify the equation, temporally ignoring the two bias vectors b_i, b_j. Given the absence of lateral connections, an RBM is a 'product of expert' model (Hinton, 2012), the joint distribution of all visible units is combined with each individual visible unit multiplicatively.

$$p(v) = \frac{1}{Z} \prod_i e^{b_i v_i} \prod_j \left(1 + e^{\sum_i (b_j + w_{ij} v_i)}\right) \tag{4}$$

Similarly,

$$p(h) = \frac{1}{Z} \prod_j e^{b_j h_j} \prod_i \left(1 + e^{\sum_j (b_i + w_{ij} h_j)}\right) \tag{5}$$

The conditional distribution of hidden units (hidden layer) given the visible units (input layer) is factored by Equations 6 and 7. Due to the absence of lateral connections, hidden units are independent given the state of the visible units and *vice versa*.

$$p(v|h) = \prod_i p(v_i|h) \tag{6}$$

Meanwhile

$$p(h|v) = \prod_j p\left(h_j \middle| v\right) \tag{7}$$

Activation probability for each individual unit is as shown in Equations 8 and 9. This activation probability (conditional probability of a single unit) is also analogous to a regular feed-forward neural network.

$$p(v_i|h) = Sigmoid\left(b_i + \sum_j h_j\, w_{i,j}\right) \tag{8}$$

$$p\left(h_j \middle| v\right) = Sigmoid\left(b_j + \sum_i h_i\, w_{i,j}\right) \tag{9}$$

- Variational Autoencoder:

  A VAE is a directed probabilistic graphical model comprising a recognition model (encoder) q(z|v) and a generative model (decoder) p(v|z), where z ∈ ℝ^d denotes a continuous lower-dimensional latent space. The marginal log-likelihood of the spatial data log p(v) is optimised by maximising the Evidence Lower Bound (ELBO), derived via Jensen's inequality in Equation 10:

$$\log p_\theta\,(v) \geq \mathcal{L}(\theta, \phi; v) = E_{q_\phi(Z|V)}[\log p_\theta\,(v|z)] - D_{KL}\left(q_\phi(z|v), p(z)\right) \tag{10}$$

where p(z) is the prior distribution over the latent variables, and D denotes the Kullback-Leibler divergence. The encoder parameterises a multivariate Gaussian distribution in Equation 11:

$$q_\phi(z|v) = \mathcal{N}\left(z; \mu_\phi(v), \text{diag}\left(\sigma_\phi^2(v)\right)\right) \tag{11}$$

To enable gradient backpropagation through stochastic nodes, the reparameterisation trick is employed by Equation 12:

$$z = \mu_\phi(v) + \sigma_\phi(v) \odot \epsilon, \quad \epsilon \sim \mathcal{N}(0, I) \tag{12}$$

The VAE is trained by balances a reconstruction error term and the Kullback-Leibler (KL) regularisation term. In geostatistical simulation, the minimisation objective acts as a spatial low-level smoothing filter, leading to the well-known 'blurry limitation' that washes out sharp geological boundaries and high-grade vein structures. Furthermore, incorporating sparse coordinate conditioning requires auxiliary objective tweaks that disrupt the true posterior distribution.

- Generative Adversarial Network:

  A GAN formulates generative modelling as a zero-sum minimax game between two subnetworks: a generator G(z; g) that maps latent noise z ~ p_z(z) to the spatial data space, and a discriminator D(v; d) that estimates the probability that a given spatial realisation is drawn from the true geological training images rather than the generator. The optimisation objective is expressed as the value function V(D, G) in Equation 13:

$$\min_G \max_D V(D, G) = E_{v \sim p_{data}(v)}[\log D(v)] + E_{z \sim p_z(z)}\left[\log\left(1 - D\left(G(z)\right)\right)\right] \tag{13}$$

During training, the discriminator parameters are updated to maximise the classification accuracy of real versus synthetic geological realisations, while the generator parameters are optimised to fool the discriminator by minimising log(1 – D(G(z))).

While Convolutional GANs (Spatial GANs) excel at rendering sharp, visually striking macro-scale patterns, they are notoriously vulnerable to mode collapse, where the generator captures only a restricted mode subset of the training image variability. This suppresses low-probability, high-impact geological scenarios essential for resource risk assessment. Moreover, enforcing hard conditioning via conditional GANs (cGANs) or auxiliary loss networks introduces structural artifacts that distort the true spatial probability space.

COMPARISON

In mining operations, simulation models must strictly honour local borehole observations (hard data) as absolute physical constraints. To evaluate how different generative models achieve this while balancing stochastic variability, we compare the conditional simulation results of RBMSim, VAEs, and GANs at an 8 per cent sub-sampled hard data density, as illustrated in Figures 1 and 2, with comparisons in Table 1.

- RBMSim (Exact Joint Slicing): As shown in Figures 1(b) and 2(b), RBMSim successfully balances stochastic diversity and structural sharpness. Because the RBM characterises the full joint distribution, conditional simulation requires zero architectural modifications or secondary optimisation loops. By partitioning the visible layer into conditioned locations v_{cond} (borehole data) and unconditioned simulation locations v_{sim}, the conditional distribution $p(v_{sim}|v_{cond})$ is obtained as a mathematically exact slice of the joint distribution p(v,h). The physical data acts as an uncompromised source of truth. While minor local fluctuations may occur due to discrete grid sampling, these do not negatively impact subsequent mine planning workflows.
- VAEs and GANs (Auxiliary Approximations): As depicted in Figures 1(c) and 2(c), VAE simulations are overly smoothed and suffer from a severe loss of geological detail. Conversely, while GAN simulations render visually striking macro-scale patterns, independent realisations often appear nearly identical, exhibiting a critical lack of stochastic variability that poses significant challenges for risk-based mine planning. This operational bottleneck arises because VAEs and GANs prioritise global continuous topologies over local data rigidity, they cannot naturally ingest sparse spatial coordinates. Forcing these models to honour drill hole data requires complex auxiliary loss functions, conditional GAN (cGAN) architectures, or secondary optimisation retraining. These ad hoc interventions distort the underlying probability space,

resulting in a pseudo-conditional surrogate distribution that deviates from the true geostatistical posterior.

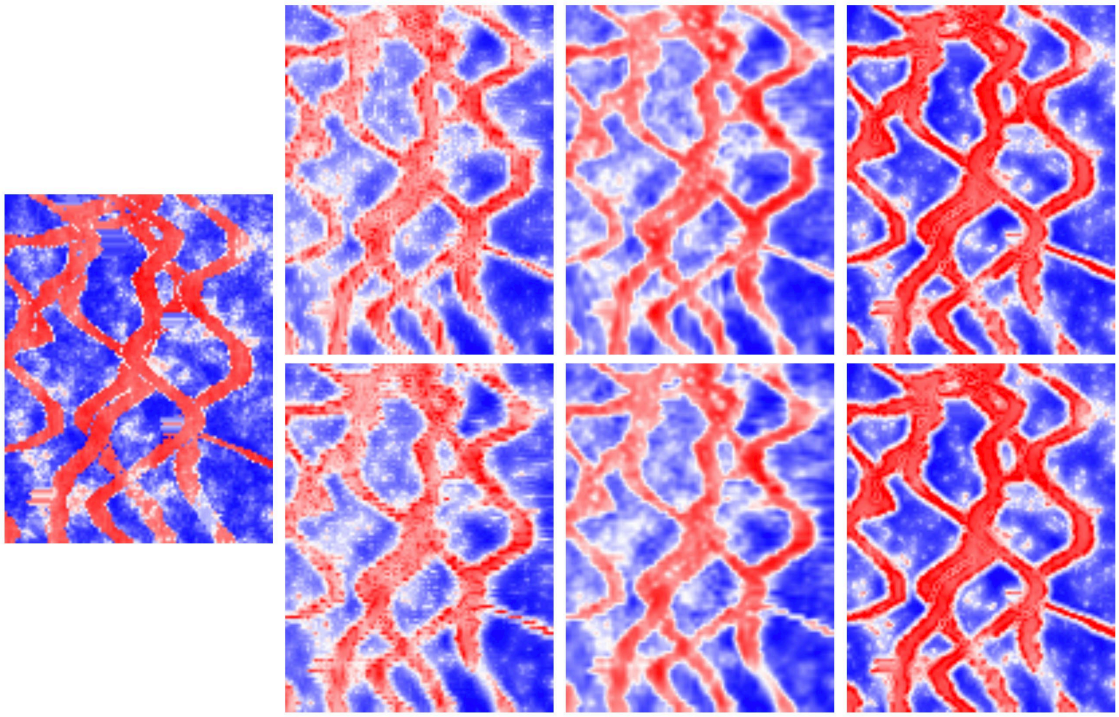


FIG 1 – Simulation results 1: (a) Testing image for conditional simulation. Simulation results: From left to right, each column contains two independent conditional simulation on hard data sub-sampled at densities of 8 per cent. (b) RBMSim results, (c) VAEs results, and (d) GANs results.

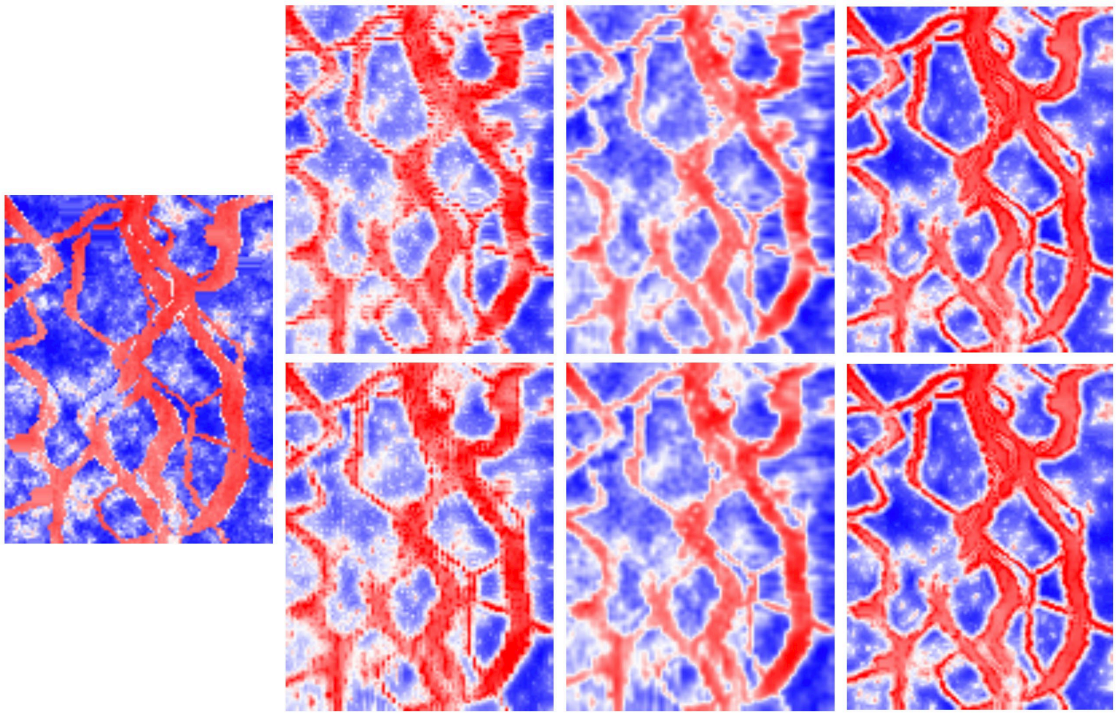


FIG 2 – Simulation results 2: (a) Testing image for conditional simulation. Simulation results: From left to right, each column contains two independent conditional simulation on hard data sub-sampled at densities of 8 per cent. (b) RBMSim results, (c) VAEs results, and (d) GANs results.

TABLE 1

Model comparison.

Evaluation criteria	RBMSim	VAEs	GANs
Mathematical foundation	Energy-based model (Ising equivalence)	Variational Bayesian inference	Minimax zero-sum game
Spatial continuity representation	Sharp boundaries with artifices	Continuous low-pass approximation	Crisp macro-scale visual patterns
Stochastic diversity	Complete coverage of configuration space	Smooth Gaussian latent space	High susceptibility to Mode Collapse
Hard data conditioning	Exact joint distribution slicing $p(v_{sim}\|v_{cond})$	Requires auxiliary loss/retraining	Requires cGAN/ secondary optimisation
Structural interpretability	High (weights map to physical energy)	Low (black-box latent space)	Low (black-box generator/ discriminator)

CONCLUSIONS

While the mainstream AI community consistently drives toward deeper and more complex architectures, the field of mathematical geosciences demands a distinct paradigm prioritised by structural interpretability and comprehensive probability coverage of uncertainty spaces. In rigorous orebody modelling, deep generative frameworks like Variational Autoencoders (VAEs) and Generative Adversarial Networks (GANs) encounter fundamental operational limitations that our proposed single-hidden-layer RBM effectively resolves. This superiority is rooted in three critical architectural advantages.

1. Equivalence to Physical Mechanisms (Energy-Based Interpretability):

   Unlike black-box deep networks, the RBM is inherently an energy-based generative model. Its underlying mathematical formulation shares a direct, elegant equivalence with the classic Ising model in statistical physics, where spatial configurations of binary states are governed by a global energy function. In a geological context, this framework allows geo-statisticians to explicitly map intricate spatial connectivity and curvilinear structures directly to the system's energy states. Every weight and bias within the RBM interaction matrix directly reflects the physical likelihood of geological continuity, providing an intuitive and mathematically rigorous baseline that geologists can interpret and trust—a stark contrast to the uninterpretable latent dimensions of deep architectures.

2. Completeness in Manifold Sampling for Unconditional Simulation:

   When evaluating geological resources prior to establishing spatial constraints, unconditional simulation must capture the full diversity of the configuration space to accurately quantify geological risk. However, the latent structure of modern deep architectures inherently distorts this stochastic variety:

   - The Blurry Limitation of VAEs: VAEs typically enforce a continuous, single-mode Gaussian distribution on their latent spaces. When applied to highly non-linear, sharp geological boundaries—such as narrow vein-type deposits, faults, or abrupt lithological contacts—this Gaussian smoothing objective acts as a low-pass filter, yielding 'blurry' realisations that systematically fail to replicate exact physical continuity.
   - The Mode Collapse of GANs: GANs are notorious for mode collapse. The generator frequently optimises by selectively reproducing only a handful of dominant, visually pleasing geological patterns. For mining teams attempting to evaluate operational financial risk, this lack of stochastic variety is catastrophic, as it completely omits the low-probability, high-impact scenarios required for a representative simulation ensemble.

3. Mathematically Exact Conditioning via Joint Distribution Slicing:

   The divergence between RBMs and deep learning models becomes even more pronounced in conditional simulation, where the realisation must strictly honour coordinate-specific borehole measurements (hard data).

   - Because VAEs and GANs prioritise global continuous topologies over local data rigidity, they cannot naturally ingest sparse, location-specific physical constraints. To force these models to honour local observations, engineers must rely on auxiliary loss functions, secondary optimisation loops, or specialised model re-training. This ad hoc intervention distorts the underlying probability space, yielding a pseudo-conditional or surrogate distribution that no longer reflects the true posterior.
   - Our RBMSim framework completely circumvents these operational and mathematical hurdles. Because the network preserves the full joint distribution $p(v,h)$, transitioning from an unconditional to a conditional framework requires zero architectural modifications or secondary training loops. By partitioning the visible layer into conditioned spatial locations v_{cond} and unconditioned simulation locations v_{sim}, we can directly clamp the known borehole measurements into v_{cond}.

Consequently, the conditional distribution $p(v_{sim}|v_{cond})$ operates as a mathematically exact slice (or cross-section) of the learned full joint distribution $p(v,h)$. The locked physical data naturally dictates the conditional posterior as a 'source of truth,' ensuring that the generated realisations are drawn from a mathematically 'true' conditional distribution inherently embedded within the model. By generating independent, fully conditioned realisations from this exact equilibrium slice, RBMSim establishes a structurally sound, reproducible, and unbiased standard for Geological Uncertainty Quantification (UQ).

REFERENCES

Ackley, D H, Hinton, G E and Sejnowski, T J, 1985. A learning algorithm for Boltzmann machines, *Cognitive Science*, 9(1):147–169.

Caers, J, 2011. *Modeling Uncertainty in the Earth Sciences* (John Wiley and Sons).

Chatterjee, S and Dimitrakopoulos, R, 2011. Multi-scale stochastic simulation of mineral deposits using wavelets, *Mathematical Geosciences*, 43(6):623–642.

Deutsch, C V and Journel, A G, 1998. *GSLIB: Geostatistical Software Library and User's Guide*, 2nd ed (Oxford University Press).

Dimitrakopoulos, R, 2011. Stochastic optimization for strategic mine planning: A modern perspective on mitigating risk and maximizing economic value, *Journal of the Southern African Institute of Mining and Metallurgy*, 111(2):77–88.

Dimitrakopoulos, R, Mustapha, H and Gamafi, M, 2010. High-order spatial simulation of complex non-Gaussian geological features with spatial cumulants, *Mathematical Geosciences*, 42(5):473–495.

Dupont, E, Laloy, E, Jacques, D and Linde, N, 2018. Conditioning of generative adversarial networks (GANs) to pixel-level data with applications to subsurface characterization, *arXiv preprint*, arXiv:1811.11186.

Findlay, L and Dimitrakopoulos, R, 2026. Updating Simulated Multi-Attribute Orebody Models Using High-Order Statistics Considering Production Data of Multiple Support Scales, *Mathematical Geosciences*, https://doi.org/10.1007/s11004-025-10257-7

Goodfellow, I, Pouget-Abadie, J, Mirza, M, Xu, B, Warde-Farley, D, Ozair, S and Bengio, Y, 2014. Generative adversarial nets, *Advances in Neural Information Processing Systems*, 27:2672–2680.

Guardiano, F B and Srivastava, R M, 1993. Multivariate geostatistics: beyond bivariate distributions, *Geostatistics Troon*, 1:133–144.

Hadavand, M and Deutsch, C V, 2026. Application of Generative Adversarial Networks (GAN) to Geostatistical Modeling of Categorical Variables, *Mathematical Geosciences*. https://doi.org/10.1007/s11004-026-10280-2

Haji Abolhassani, A A, Dimitrakopoulos, R, Ferrie, F P and Yao, L, 2022. A New Non-stationary High-order Spatial Sequential Simulation Method, *Mathematical Geosciences*, 54(6):1097–1119.

Hinton, G E, 2012. A practical guide to training restricted Boltzmann machines, *Neural Networks: Tricks of the Trade*, pp 599–619 (Springer: Berlin).

Journel, A G and Huijbregts, C J, 1978. *Mining Geostatistics* (Academic Press: London).

Kingma, D P and Welling, M, 2013. Auto-encoding variational Bayes, *arXiv preprint*, arXiv:1312.6114

Laloy, E, Hérault, R, Jacques, D and Linde, N, 2018. Training image-based geostatistical inversion using a spatial generative adversarial network, *Water Resources Research*, 54(1):381–406.

Laloy, E, Hérault, R, Linde, N, Jacques, D and Vrugt, J A, 2017. Inversion of airborne electromagnetic data using a generative model based on a variational autoencoder, *IEEE Transactions on Geoscience and Remote Sensing*, 56(2):705–718.

Moss, R, Srinivasan, S and Caers, J, 2020. Generative adversarial networks for conditional simulation of complex geological structures, *Computers and Geosciences*, 141:104523.

Mu, Y, Ferrie, F and Dimitrakopoulos, R, 2015. Sparse Image Reconstruction by Two Phase RBM Learning: Application to Mine Planning, in *Proceedings of the 14th IAPR International Conference on Machine Vision Applications (MVA)*, pp 122–125.

Mustapha, H and Dimitrakopoulos, R, 2011. High-order stochastic simulation of geological patterns using spatial cumulants and a simulation path, *Mathematical Geosciences*, 43(3):323–345.

Strebelle, S, 2002. Conditional simulation of complex geological structures using multiple-point statistics, *Mathematical Geology*, 34(1):1–21.

Yao, L and Dimitrakopoulos, R, 2020. High-Order Sequential Simulation via Statistical Learning in Reproducing Kernel Hilbert Space, *Mathematical Geosciences*, 52(5):693–723.

Yao, L, Dimitrakopoulos, R and Gamache, M, 2021. Training Image Free High-Order Stochastic Simulation Based on Aggregated Kernel Statistics, *Mathematical Geosciences*, 53(7):1100–1125.

Zhang, T, Switzer, P and Journel, A, 2006. Filter-based classification of training images for multi-point geostatistics, *Mathematical Geology*, 38(1):63–91.

Geometallurgy and Mine to Mill integration

Strategic Mine to Mill as a risk-informed value system

F Faramarzi[1]

1. Technical Lead, GEOVIA Dassault Systèmes, Brisbane Qld 4000.
   Email: farhad.faramarzi@3ds.com

ABSTRACT

Mining is inherently probabilistic. Orebody variability, geo-metallurgical uncertainty, time-dependent economics, organisational incentives, and integration choices interact to produce outcomes that cannot be managed through operational adjustments alone. Value losses arising from disconnected upstream and downstream strategic planning, coupled with silo-driven performance metrics, cannot be recovered downstream; instead, they compound over time through path dependency and are often obscured by short-term performance improvements.

Short-term, siloed gains rarely support long-term profitability. Conventional Mine to Mill (M2M) practices have historically focused on short-term operational optimisation, typically targeting blast fragmentation, milling throughput, and recovery within production time frames. While such initiatives can temporarily improve downstream overall performance, they often fail to maintain or enhance value over the life-of-mine (LOM) because irreversible upstream decisions made under uncertainty dominate LOM value realisation. This paper introduces Strategic Mine to Mill as a new methodology, reframing Mine to Mill as a risk-informed LOM value system, rather than a downstream optimisation practice with limited sustainability. Strategic Mine to Mill is an approach that establishes the missing link between pit-level strategic decisions and downstream operational performance.

Integration strategy defines the bottleneck. Two integration strategies are discussed: one explores the what-ifs of adapting the pit to process plant constraints (Pit → Plant), and the other considers designing the process plant to accommodate orebody variability (Plant → Pit). Strategic Mine to Mill focuses on evaluating how early strategic choices may shift bottlenecks, reshape risk exposure, and fundamentally alter long-term value trajectories.

INTRODUCTION

What defines the *best possible value* in the minerals industry? While the '*best possible value*' in mining is a moving target, still deterministic methods maintain a strong base of supporters. It seems that the notion of 'best possible value' is both essential and elusive. While it guides decision-making, value is fluid and continuously shaped by changing economics, operational bottlenecks and constraints across the life-of-mine (LOM), evolving technologies, improved understanding and practices, and the inherent uncertainties of orebodies and mining operations. Mining and processing ore can easily be misinterpreted as creating value, when in reality, these activities alone may treat ore without maximising economic return. The confusion arises because the physical act of extracting ore is tangible and measurable, whereas true value is an economic outcome that depends on how 'efficiently and effectively' that ore is transformed into revenue within a certain time frame. Even mining and processing very high-grade ore at high rates can potentially result in profit lost over time if strategies and practices are poorly set and implemented. Moreover, value is inherently time-dependent. What constitutes the 'best possible value' today is shaped by the current time value of money, commodity prices, operational costs, and available technology. Decisions that are optimal under today's understandings and constraints may not remain so in the future. Value is therefore a moving target, inherently impacted by uncertainties and their associated risks. In other words, the measures defining an optimised strategic mine plan frequently evolve as new information becomes available, whether on orebody characteristics, operational, or market conditions which ultimately govern the expansion or contraction of mining projects. This highlights the importance of risk incorporation in strategic mine planning (SMP) as a factor that has received increasing attention in recent years through the adoption of stochastic approaches, which improve plan robustness by explicitly accounting for geological, geo-metallurgical, and economic uncertainties. That is, scenario-based analysis has been widely used in mining and processing with exemplified applications in optimising upstream (Bye, 2006, 2011; Dimitrakopoulos, 2018; Godoy, 2018; Kojovic, Michaux and McKenzie, 1995; Morales *et al,* 2019; Poblete *et al,* 2016; Rimélé, Dimitrakopoulos and Gamache,

2020; Smith, Faramarzi and Poblete, 2022) and downstream KPIs (Carrasco *et al,* 2017; Faramarzi *et al,* 2018, 2020; Faramarzi, Kanchibotla and Morrison, 2019).

A bad day for a silo can be a good day for the entire mining value chain. When it comes to interconnected systems such as mining, a common pitfall is confusing performance with profitability. Improved performance of individual silos such as blasting, loading, haulage, crushing, comminution, and beneficiation does not necessarily guarantee improved value. This is a fundamental understanding on which conventional Mine to Mill (M2M) thinking is built, accurately targeting the integrated nature of the mining industry as a lever for improved overall value. Hard rock mining is an interconnected size reduction–mineral liberation process, aiming for maximum possible profitability. Since the early 1980s, a number of M2M studies conducted across the globe have shown that for an improved overall value, tuning blast designs to produce run-of-mine (ROM) fragmentation in favour of the downstream ie crushing and grinding improves overall performance and hence profitability (McKee, 2013). Documented M2M projects have reported typical milling throughput improvements of 5–20 per cent, generally achieved through intensified blasting that generates a greater proportion of -10 mm fragmented ore in the ROM (Faramarzi *et al,* 2022; Hart *et al,* 2001; Kanchibotla, Valery and Morrell, 1999; Grundstrom *et al,* 2001; Valery, Duffy and Jankovic, 2019). Additionally, number of case studies from the Asia Pacific region indicate measurable improvement potentials that serve as the proof-of-concept basis for this study. Conventional Mine to Mill practice demonstrates the potential to improve milling capacity by up to 20 per cent through optimised blast fragmentation, increase pit loading productivity by up to 10 per cent, and improve metal recovery by up to 5 per cent through feed quality management (Burger *et al,* 2006; Lam *et al,* 2001; Strohmayr and Valery, 2001; Kanchibotla *et al,* 1998, 2015; Young *et al,* 1997). While these values are site specific, they are adopted in this paper to estimate the impact of such improvements on LOM value at both strategic and tactical levels. However, referring to the published M2M studies, conventional M2M practices/optimisation campaigns have traditionally focused on short-term performance improvements, principally targeting blast fragmentation, milling throughput, and metallurgical recovery. That means conventional M2M has been an optimisation approach implemented at the operational/tactical levels. While such initiatives can yield measurable periodic gains, they systematically fall short of preserving or enhancing value over the LOM for in majority of cases, M2M objectives are disconnected from the pit realities nor supported by strategic and tactical mine plans (Figure 1). This deficiency is fundamentally attributable to the dominant influence of irreversible pit-level decisions on downstream activities.

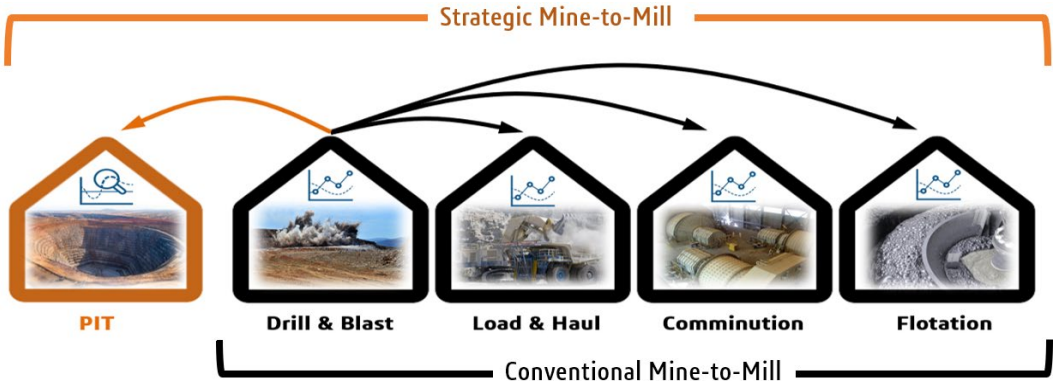


FIG 1 – The basic difference between the two M2M optimisation approaches.

This challenge motivated the author to explore possible solutions for establishing a more meaningful link between the pit and the process plant through a different concept, namely ‘Strategic Mine to Mill’ (Faramarzi and Smith, 2023a, 2023b). The Strategic M2M as a ‘risk-informed value system’ aims at:

- Making decisions where value is evaluated together with the risks involved.
- Improving the robustness of M2M campaigns for sustained value over the LOM.

Accordingly, this paper examines two contrasting Strategic Mine to Mill integration approaches: *pit-adapting-to-plant* (Pit → Plant), where process plant stability is prioritised at the potential expense of economic flexibility; and *plant-adapting-to-pit* (Plant → Pit), where net present value (NPV)

maximisation is pursued at the cost of milling stability, driven by exposure to more variable and less favourable ore feed characteristics. Both approaches present distinct trade-offs between operational robustness and economic value realisation. The paper critically discusses these trade-offs, highlighting their implications for LOM decision-making and long-term value realisation.

THE STRATEGIC MINE TO MILL APPROACH

For most operations it would be difficult if not impossible to optimise their way out of a bad strategic plan of mining. In the context of mining value chain optimisation, specifically M2M, an expensive day in blasting can translate into a highly profitable outcome for the entire value chain through improved milling performance. Assuming optimal M2M performance is achieved at the operational level and sustained consistently across the LOM, the question then becomes *whether such operational excellence alone is sufficient to maximise long-term value realisation*?

SMP is among the most consequential determinants of an operation's financial performance over its LOM. Suboptimal decisions regarding pit limits, cut-off grades, mining direction, pushback sequencing, and production scheduling can result in value losses of hundreds of millions of dollars; losses that downstream operational improvements are rarely able to recover. Regardless of the quality of metallurgical characterisation, drill-and-blast execution, Mine to Mill optimisation, or process plant performance, these gains cannot compensate for a fundamentally flawed strategic mine plan as errors at this stage propagate irreversibly through the entire mining value chain. With that proviso in mind, Strategic M2M offers the opportunity to further explore the interaction between the upstream (pit) versus downstream (plant) as response to an adopted integration strategy.

In the minerals industry, an integration strategy defines where the bottlenecks reside within the system. Moving beyond the conventional M2M paradigm, depending on the integration strategy (Pit → Plant or Plant → Pit) an operation may adopt either deliberately or by default, each carrying fundamentally different implications for performance, bottlenecks, and value realisation. The question is: *which approach truly drives value in an operation: the pit adapting to the process plant, or the process plant adapting to the pit?*

The two integrated optimisation approaches which fall under the Strategic M2M framework are discussed below to further explore this distinction and its likely practical outcomes

Pit adapting to process plant constraints (Pit → Plant)

Pit adapting to the process plant integration strategy comes with the risk of the pit becoming a bottleneck. In Pit → Plant integration strategy, the mine plan is configured to align with the operational constraints and specifications of the existing process plant. Rather than designing the plant around the variability of the orebody, the pit shell optimisation and extraction scheduling are adjusted so that the mine feed consistently satisfies processing limits. This represents a deliberate strategic shift in which processing stability, rather than ore value maximisation, becomes the primary driver of mine planning decisions.

From the planning and operational perspective, under Pit → Plant approach, pit shell optimisation and mine scheduling prioritise ore domains to ensure their metallurgical characteristics are compatible with plant design parameters. Feed attributes such as hardness, particle size distribution (PSD), grade variability, and moisture should be managed within the pit to reduce processing uncertainty. Scheduling decisions are therefore made with the explicit objective of optimising plant throughput, minimising feed variability, milling stability and controlling operational risk. This strategy delivers measurable operational benefits such as improved process efficiency, reduced unit operating costs ($/t processed), more predictable metallurgical recoveries, and high plant utilisation. These outcomes reflect an inherent prioritisation of short-term operability and process stability over value-driven mining flexibility. However, this stability comes at a cost. High-value ore domains whose geo-metallurgical characteristics fall outside plant constraints (eg too hard, too variable, fine-grained) may be deferred or excluded from the mining schedule. This selective exclusion, driven by the Pit → Plant integration approach, can potentially reduce NPV and overall pit utilisation/reserve recovery and representing a remarkable economic trade-off.

From the implementation perspective, this approach is generally easier to implement and carries lower operational risk, as it minimises complexity in plant operations and does not require significant overdesign or operational flexibility in the processing facility. The conventional M2M framework, as reflected in implemented projects and published literature, is consistent with this paradigm, which aims at optimising the interface between mining and processing within the bounds of fixed plant design.

Process plant adapting to pit constraints (Plant → Pit)

Process plant adapting to the pit integration strategy comes with the risk of the process plant becoming a bottleneck. In this more aggressive integration strategy, the process plant should be designed and operated to follow the production schedule of the pit, rather than constraining it. The mining team targets full orebody exploitation, prioritising extraction of high-economic-value domains in accordance with the LOM schedule, even where those ore domains introduce significant variability in feed grade, hardness, or mineralogy as these will be downstream issues to deal with. Capital and infrastructure decisions are therefore driven by the need for processing flexibility, rather than operational simplicity.

Scheduling and operational decisions under Plant → Pit integration approach are oriented towards NPV maximisation, accepting temporary process instability or throughput variability as an 'understood and accepted' consequence of value-driven mining. The plant must be capable of adapting to continuously changing feed conditions arising from multiple ore domains which exhibit varying comminution characteristics. This requires flexible and often overdesigned infrastructure, advanced process control systems, robust metallurgical response capabilities, and a highly skilled operational workforce.

What about implementation? Theoretically, this strategy maximises NPV through closer alignment with LOM economic objectives and more complete reserve exploitation. However, Plant → Pit integrated optimisation approach is considerably more difficult to implement than the Pit → Plant approach. The operational complexity introduced by variable feed conditions increases both capital expenditure (CAPEX) and operating expenditure (OPEX). Where plant adaptability proves insufficient, the consequences include reduced throughput, increased downtime, lower metallurgical recoveries, and accelerated equipment wear. More importantly, retrofitting an existing plant to accommodate this level of variability is both costly and operationally disruptive.

Plant → Pit strategy prioritises long-term economic value but carries a materially higher risk and implementation complexity profile relative to the Pit -to- Plant approach. It should be considered as a '*Design-Stage Integration*' strategy, and aligns with an advanced M2M philosophy in which processing flexibility is engineered into the plant from the outset to deal with inherent variabilities within an orebody.

Strategic M2M integrated optimisation strategies – Pit → Plant versus Plant → Pit

Depending on which integrated optimisation strategies is prioritised by an operation, it should be noted that:

> *The selection of an integration strategy creates a path dependency that fundamentally reshapes the trajectory of project development.*

This dependency propagates across mine planning and design, plant design and performance, scheduling priorities, resource allocation, and risk management frameworks. Consequently, even if two operations begin with identical orebody characteristics, the chosen integration strategy causes their development paths and long-term outcomes to diverge progressively over time. This divergence is not inherently favourable or unfavourable. This *path dependency* is an intrinsic feature of an integrated system ie the mining value chain, not a flaw in either strategy.

What is not intrinsic, however, and what is entirely within the control of the project team, is the misalignment of Key Performance Indicators (KPIs) and organisational objectives between the mining and processing functions. Unlike path dependency which is more system-driven, conflicting KPIs are a human artefact, which can be identified, managed, and resolved through deliberate

integration planning. Left unaddressed, however, this misalignment represents one of the most consequential sources of value loss across the two strategies.

Last but not least, the adopted integration strategy fundamentally determines where contingency provisions should be allocated across the mining value-chain. Under a Pit → Plant approach, contingency should be mainly concentrated within the mining (upstream) covering ore selectivity, extraction sequencing flexibility, and geotechnical risk since the pit carries the adaptive burden of satisfying plant constraints. Under a Plant → Pit strategy, contingency should be re-directed toward the process plant (upstream), addressing equipment overdesign margins, metallurgical variability buffers, and process control redundancy, as the plant operational stability and continuity is at risk. Therefore, misaligning contingency allocation against the adopted integration philosophy may build a systemic planning error with material consequences for NPV, reserve recovery, and operational continuity. The following section presents a comparative analysis of the two approaches as applied to an open pit copper deposit, evaluated through what-if scenario analysis using GEOVIA Whittle® software (version 2022 Refresh 3, Dassault Systèmes).

BASE-CASE MODEL DEVELOPMENT

Input data

A block model of an open pit copper (Cu) deposit used to develop the base case model for quantifying the impact of strategic integration choices on overall LOM value, measured by NPV. For this purpose, GEOVIA Whittle® was used as the pit optimisation and strategic scheduling platform, which enables the integration of open pit mining parameters with downstream process plant KPIs ie mill throughput and copper flotation recovery that are the focus of this paper.

Both base cases share identical input parameters: a fixed mining capacity of 100 Mt/a, milling capacity of 30 Mt/a, copper recovery of 88 per cent, and a long-term copper price of $6.56/lb. It is worth to be noted that mining, processing and selling costs were considered fixed. The two base cases are differentiated solely by mining direction, with Base Case A mining eastward and Base Case B mining northward, corresponding to the Plant → Pit and Pit → Plant strategies respectively. To mimic the consequence of the two different integration strategies, it is assumed that deposit introduces differential geo-metallurgical behaviour in different mining directions that underpins the distinction between the Pit → Plant and Plant → Pit integration strategies evaluated in the scenario analysis, as each rock type responds differently to comminution and flotation, thereby influencing both the milling rate and metal yield/recovery efficiency at which value is extracted from the orebody over the LOM. Pushbacks were identified to maximise value based on pit shells generated by the Pseudoflow algorithm. Figure 2 shows the ultimate pit, direction of mining toward north and east as well as Cu distribution within the block model.

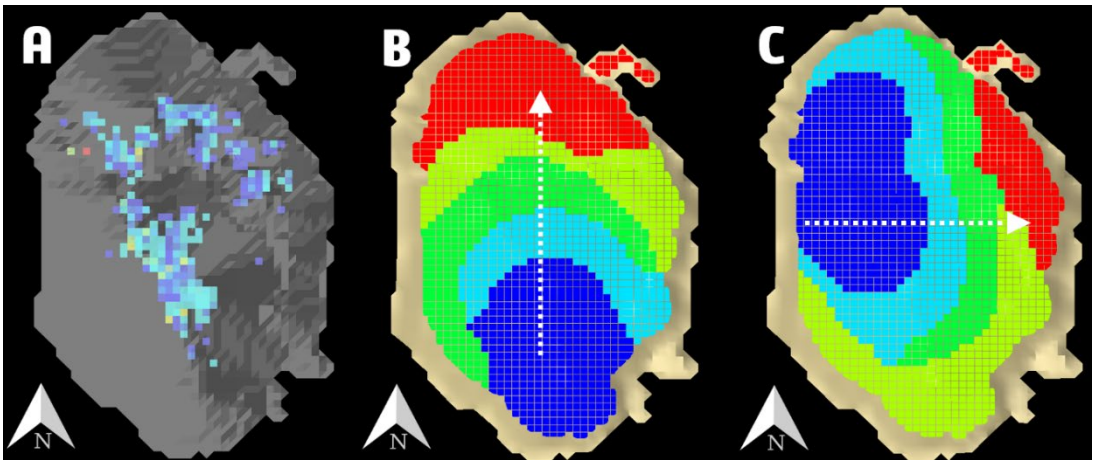


FIG 2 – Distribution of CU>0 within the block model (A), Development of pushbacks toward North as Pit → Plant (B), Development of pushbacks toward east as Plant → Pit integration approach.

The average life of each pushback ranges between four and six years across both pit scenarios, ensuring consistent ore feed availability to the downstream process plant. Production scheduling was performed using the Milawa NPV algorithm within GEOVIA Whittle®, which maximises LOM discounted cash flow. Figures 3 and 4 show pushbacks tonnages for the two approaches.

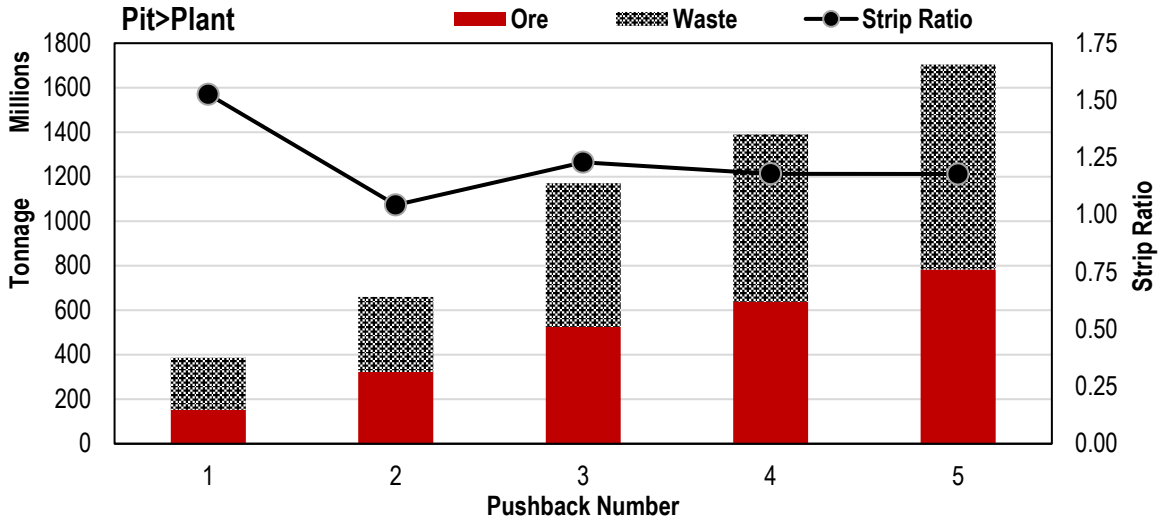


FIG 3 – Tonnage of pushbacks toward North as Pit → Plant.

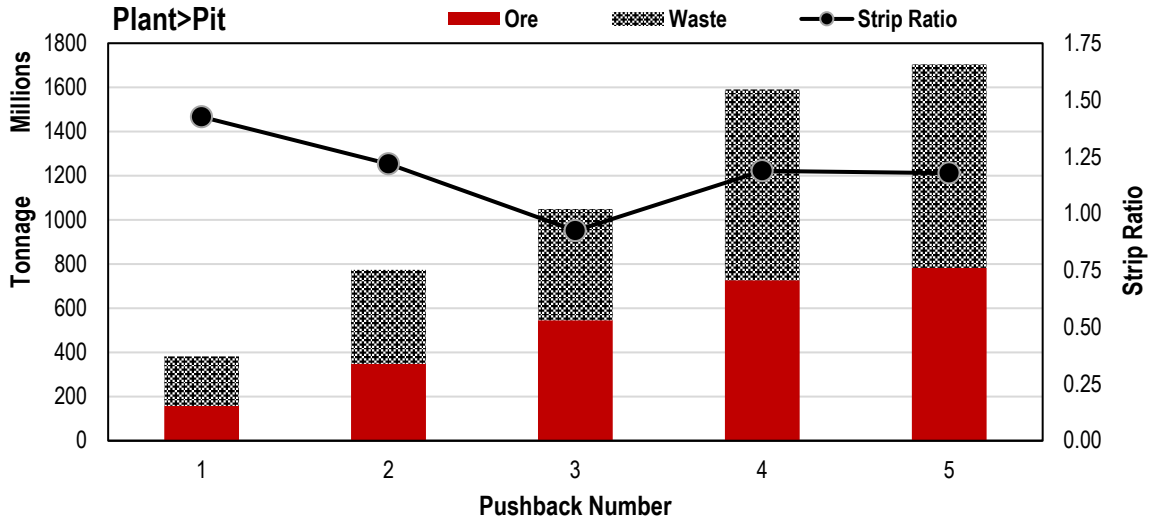


FIG 4 – Tonnage of pushbacks toward East as Plant → Pit.

Experimental design, assumptions and considerations

The experimental design evaluates two strategic integration approaches of Plant → Pit and Pit → Plant, each benchmarked against a base case. Mining capacity is assumed fixed at 100 Mt/a for both the two strategies which are differentiated by mining direction, with East assigned to Plant → Pit and North assigned to Pit → Plant.

In this study, mining direction determines what ore types should be prioritised for mining depending on their geo-metallurgical characteristics (ie grade + hardness) over the LOM, and is therefore the primary mechanism through which each strategy expresses its underlying planning philosophy. In other words:

- Pit → Plant: The mining sequence is planned in favour of process plant feed requirements, prioritising feed grade and hardness consistency to ensure maximum operational stability and predictable downstream performance over the LOM.

- Plant → Pit: The mining sequence is planned to maximise NPV over the LOM, irrespective of how varying ore characteristics, particularly feed grade and hardness may affect downstream process plant performance.

It is assumed that mining toward North prioritises ore domains with relatively uniform grade and hardness, delivering consistent feed to the process plant and supporting stable mill throughput and flotation recovery. This operational stability, however, comes at a cost: higher-value ore domains are deferred, reducing LOM NPV relative to the value-maximising alternative.

Mining eastward targets higher-grade, economically superior ore domains, maximising NPV over the scheduling periods. This mining direction, however, exposes the process plant to greater variability in feed grade and hardness, introducing throughput and recovery uncertainty that generally cannot be fully managed through operational adjustments alone.

The North direction therefore represents process stability at the expense of value, while the East direction represents value maximisation at the expense of process stability, which is a trade-off central to the Strategic Mine to Mill framework examined in this paper. Table 1 details scenario-based simulation runs for sensitivity analysis of individual key variables under different strategies.

TABLE 1

Ranges and levels of key variables used to run 60 scenarios in GEOVIA Whittle®.

Scenario	Mining direction	Mining capacity (Mt/a)	Milling capacity (Mt/a)	Copper recovery (%)	Copper price ($/lb)
Base Case A	**East**	**100 (fixed)**	**30 (fixed)**	**88 (fixed)**	**6.56 (fixed)**
Plant → Pit	East	100 (fixed)	30 (-5, -10, -20%)	88 (fixed)	6.56 (fixed)
Plant → Pit	East	100 (fixed)	30 (fixed)	88 (-3, -5%)	6.56 (fixed)
Plant → Pit	East	100 (fixed)	30 (-5, -10, -20%)	88 (-3, -5%)	6.56 (fixed)
Plant → Pit	East	100 (fixed)	30 (-5, -10, -20%)	88 (-3, -5%)	6.56 (-10, -20, -30%)
Base Case B	**North**	**100 (fixed)**	**30 (fixed)**	**88 (fixed)**	**6.56 (fixed)**
Pit → Plant	North	100 (fixed)	30 (+5, +10, +20%)	88 (fixed)	6.56 (fixed)
Pit → Plant	North	100 (fixed)	30 (fixed)	88 (+1, +3%)	6.56 (fixed)
Pit → Plant	North	100 (fixed)	30 (+5, +10, +20%)	88 (+1, +3%)	6.56 (fixed)
Pit → Plant	North	100 (fixed)	30 (+5, +10, +20%)	88 (+1, +3%)	6.56 (-10, -20, -30%)

In total, 60 scenario simulations are run, comprising 30 scenarios for each integration strategy. There is a base case for each mining strategy (North versus East): fixed milling capacity of 30 Mt/a, fixed copper recovery of 88 per cent, and a fixed long-term copper price of $6.56/lb. Four scenarios are then built within each strategy, each isolating or combining specific operational and economic variables:

- The first scenario varies milling throughput only by applying typical reductions of 5 per cent, 10 per cent, and 20 per cent for Plant → Pit to reflect throughput loss arising from unmanaged feed variability, and increases of 5 per cent, 10 per cent, and 20 per cent for Pit → Plant to reflect throughput gains achievable through planned desirable ore sequencing and feed grade control.
- The second scenario varies copper flotation recovery only by applying reductions of 3 per cent and 5 per cent for Plant → Pit to reflect the metallurgical penalty associated with variable feed, and increases of 1 per cent and 3 per cent for Pit → Plant to reflect the recovery benefit of consistent, grade-controlled feed delivery.
- The third scenario combines both throughput and recovery variations at the same time, representing the realistic integrated operational outcome of each strategy.

- The fourth scenario introduces copper price stress of 10 per cent, 20 per cent, and 30 per cent below the base case price, applied to the combined operational scenario, to evaluate the impact and test the robustness of each strategy (Plant → Pit versus Pit → Plant) under market volatility.

Assumption 1 – Pit → Plant: zero incremental CAPEX

Adopting a Pit → Plant integration strategy inherently involves minor operational cost adjustments, most notably increased drilling and blasting expenditure associated with achieving a target feed particle size distribution through intense blasting suited to existing mill design parameters. While these OPEX increments are real and acknowledged, they are considered negligible at the LOM scale evaluated in this study. For simplicity and comparative clarity, no incremental CAPEX or OPEX adjustment is applied to the Pit → Plant scenarios given in this paper.

Assumption 2 – Plant → Pit: zero incremental CAPEX

Engineering a process plant to follow pit variability requires remarkable capital investment. The magnitude of this investment scales with the inherent orebodies complexities as well as their geo-metallurgical variability particularly feed hardness/competence ranging from moderate upgrades to substantial circuit modifications depending on the deposit characteristics. These capital consequences are real, acknowledged, and documented in recent industry practice (Antofagasta PLC, 2023; Capstone Copper Corp, 2025). However, for the purposes of this study, no incremental CAPEX is applied to the Plant → Pit scenarios. This assumption isolates the scheduling and operational effects of each integration strategy and quantifies associated impact on the overall value if the Plant → Pit approach is adopted. The author acknowledges that any real-world CAPEX requirement would reduce this benefit and may shift the economic advantage toward the Pit → Plant strategy which could be examined in future studies.

Assumption 3 – Mining capacity fixed at 100 Mt/a

In practice, variations in rock mass characteristics across ore domains including differences in rock strength (unconfined compressive strength – UCS), fracture density (rock quality designation – RQD), and Blastability Index (BI) can affect drill and blast productivity and load-and-haul cycle times, introducing variability in the annual mining rate. However, such domain-related productivity fluctuations are considered negligible in this study, and it is assumed that rock mass driven variations to have no impact on the overall annual mining rate. Mining capacity is therefore held fixed at 100 Mt/a across all scenarios.

Assumption 4 – Ore blending and stockpiling

In practice, ore blending is a widely adopted operational strategy for managing feed variability to the process plant, particularly in deposits exhibiting significant variation in grade, hardness, or mineralogy across ore domains. Blending through the development and management of multiple ore stockpiles assists with mill performance improvement and metallurgical recovery. However, establishing and operating dedicated blending stockpiles incurs additional CAPEX associated with stockpile infrastructure, reclaim systems, and materials handling equipment, as well as incremental OPEX from additional re-handling, equipment utilisation, and inventory management. It is acknowledged that incorporating blending strategies in practice would partially offset the feed variability effects modelled in the Plant → Pit scenarios, and would introduce additional capital and operating cost considerations beyond the scope of this analysis. Therefore, for simplicity and isolating the effects of strategic integration approaches on value, ore blending and stockpile management are not considered in this study and are recommended as a subject for future work.

Simulation scenario results

The simulation results for all 60 scenarios, comprising 30 Plant → Pit and 30 Pit → Plant runs, are presented in the Appendix Tables A1 and A2 respectively, reporting NPV and LOM outputs for each combination of milling capacity, copper recovery, and copper price.

The Plant → Pit simulation results show a consistent and quantifiable relationship between downstream operational performance and LOM value, confirming that process plant

underperformance arising from unmanaged feed variability imposes a measurable NPV penalty without a corresponding reduction in mining capacity. The simulation results for all 30 Plant → Pit scenarios are presented in Appendix Table A1.

On the other hand, the Pit → Plant simulation results show a consistent and positive relationship between upstream scheduling decisions and LOM value, where improvements in milling throughput and copper recovery through ore sequencing and feed grade control generate measurable NPV gains. The full results for all 30 Pit → Plant scenarios are presented in Appendix Table A2.

DISCUSSION

What-if adapting Process Plant → Pit?

The Plant → Pit simulation results demonstrate a consistent and quantifiable relationship between downstream operational performance and LOM value, confirming that process plant underperformance arising from unmanaged feed variability imposes a measurable NPV penalty without a corresponding reduction in mining capacity.

The simulation results for all 30 Plant → Pit scenarios is given in Figure 5, including NPV, and life-of-mine outputs for each combination of milling capacity, copper recovery, and copper price, are presented in Appendix Table A1.

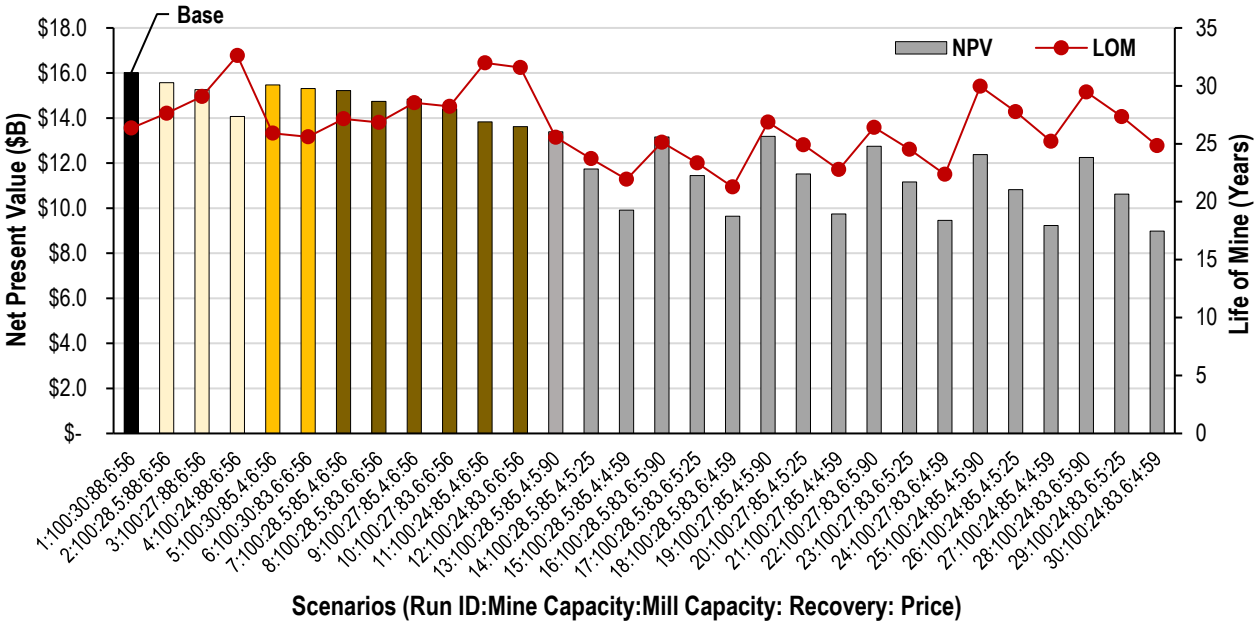


FIG 5 – NPV and LOM variations for 30 scenarios if adapting Process Plant → Pit (colours represent simulation four groups; base case shown in black).

Under base case conditions (*Scenario1: 30 Mt/a milling, 88 per cent recovery, and $6.56/lb copper price*) the Plant → Pit strategy yields an NPV of $16.02B over a 26.35-year mine life. All subsequent scenarios within this strategy represent departures from this reference, and the NPV differences quantify the impact of feed variability on LOM value.

Effect of throughput reduction (Group 1 – Scenarios 2 to 4)

Isolating the effect of milling throughput reduction at fixed recovery and copper price reveals a consistent and progressive NPV erosion. Reducing mill throughput by 5 per cent, 10 per cent, and 20 per cent yields NPV of $15.57B, $15.26B, and $14.07B respectively, representing value losses of $0.45B, $0.76B, and $1.95B against the base case. In this case, LOM extends from 26.35 years to 27.63, 29.08, and 32.61 years which is a direct consequence of processing fewer tonnes annually, or turning the process plant into a bottleneck. It is worthwhile to be noted that the 'throughput loss' penalises value through two compounding mechanisms simultaneously:

1. Reduced annual revenue generation.
2. Extended time-discounting of future cash flows over a longer LOM.

The implication is significant in an orebody where feed hardness variability is sufficient to destabilise the downstream performance and cause throughput reductions of 10 per cent to 20 per cent. That is, in this case, the Plant → Pit strategy exposed the operation to value losses of $0.76B to $1.95B relative to the base case, without any change in the resource or the copper price. This is the cost of mining only for 'value', or the typical 'grade-driven' mining approach without engineering processing flexibility to match.

Effect of recovery reduction (Group 2 – Scenarios 5 to 6)

Considering milling capacity fixed at 30 Mt/a and reducing copper flotation recovery by 3 per cent and 5 per cent down to 85.4 per cent and 83.6 per cent produces NPV of $15.47B and $15.31B respectively, with negligible impact on the LOM. Based on the results, the recovery-only NPV penalties of $0.55B and $0.71B are meaningful but smaller in magnitude than the equivalent throughput reductions. That is, throughput loss accelerates more steeply at higher reduction levels, making it the more consequential risk driver as feed hardness intensifies.

Combined effect of throughput and recovery reduction (Group 3 – Scenarios 7 to 12)

Combining throughput and recovery reductions simultaneously exposes the compounding nature of these two value drivers. At 28.5 Mt/a milling with 85.4 per cent and 83.6 per cent recovery (Scenarios 7 and 8), NPV significantly falls to $15.22B and $14.74B. At the most severe combined case within this group, 24 Mt/a milling and 83.6 per cent recovery (Scenario 12), NPV falls to $13.62B, a remarkable reduction of $2.40B from the base case, with the LOM extending to 31.58 years.

This group is the most operationally realistic within the fixed-price scenarios, as it directly represents the impact of mining eastward through a geologically 'undesirable orebody' without sufficient processing flexibility. Variable feed hardness reduces throughput while variable mineralogy simultaneously degrades recovery, and the combined effect is a significant value loss. The progression from Scenario 7 through Scenario 12 implies that as the extent of feed variability increases, economic sustainability becomes a challenge, reinforcing the central argument of this paper that upstream mining decisions propagate downstream as '*compounding operational penalties*' that cannot be recovered through short-term adjustments once the mining sequence is committed.

A further observation from this group concerns LOM extension under combined stress. Scenario 12 extends LOM by 5.23 years relative to the base case, which commonly is not beneficial for it is a symptom of value destruction, representing the operational limits of a process plant that is unable to process ore at the rate the pit can deliver it.

Combined throughput, recovery and copper price (Group 4 – Scenarios 13 to 30)

The fourth group introduces copper price stress of 10 per cent, 20 per cent, and 30 per cent below the base price across all throughput and recovery combinations from Group 3, representing the most demanding test of the Plant → Pit strategy, which would be simultaneous operational underperformance and sustained market downside.

At a 10 per cent price reduction ($5.90/lb), NPV ranges from $12.25B to $13.39B, representing an additional contraction of $0.68B to $1.82B below the Group 3 outcomes at equivalent operational parameters. At a 20 per cent price reduction ($5.25/lb), NPV contracts further to a range of $10.62B to $11.74B. At the most severe stress level, a 30 per cent price reduction at $4.59/lb, the NPV ranges from $8.98B to $9.91B, representing a reduction of 38 per cent to 44 per cent from the base case across all operational combinations.

Three observations strengthen the Strategic Mine to Mill argument from this group:

- First, the LOM shortens as copper price falls across all scenarios at $4.59/lb, to a range of 21.27 to 25.21 years compared to 24.84 to 32.61 years at base price for equivalent operational

parameters. This reflects fewer blocks being considered economic to process when lower prices render marginal ore blocks uneconomic within GEOVIA Whittle®, reducing the envelope for availability of processible feed.

- Second, the NPV spread attributable to operational variability narrows as price falls. At $6.56/lb the range between best and worst operational case in Group 3 is $2.40B; at $4.59/lb across equivalent combinations in Group 4, this spread narrows to $0.93B. This convergence reveals that under severe price stress, commodity market risk dominates operational risk as the primary value driver. This is a key consideration with direct implications for CAPEX allocation decisions under the Plant → Pit strategy, where significant processing flexibility investment may be difficult to justify when price downside can diminish operational gains.
- Third, and most critically, even under the most severe combined stress, 20 per cent throughput reduction, 5 per cent recovery loss, and 30 per cent price decline (Scenario 30) the operation retains a positive NPV of $8.98B. Again, this confirms that the deposit possesses sufficient economic resilience to remain viable across tested scenarios. However, the gap between this worst-case outcome and the base case NPV of $16.02B is $7.04B. It's a value difference that represents the maximum economic cost of adopting the Plant → Pit strategy on this deposit without engineering adequate processing flexibility to manage feed variability. This $7.04B range basically defines the upper bound of value at risk, and it is against this figure that any capital investment in processing flexibility must ultimately be evaluated.

What-if adapting Pit → Process Plant?

Figure 6 presents the NPV and LOM outcomes for all 30 Pit → Plant simulation scenarios. Under base case conditions (Scenario 31) of 30 Mt/a milling, 88 per cent recovery, and $6.56/lb, the Pit → Plant strategy yields an NPV of $12.97B over a LOM of 26.27 year. This establishes the reference for the northward mining direction, basically representing conventional Mine to Mill practice. Notably, the Pit → Plant base case NPV of $12.97B is $3.05B lower than the Plant → Pit base case of $16.02B under identical operational parameters, confirming that northward mining through less heterogenous/variable but lower-value ore domains sacrifices NPV in exchange for feed consistency, which is the fundamental trade-off between the two strategies.

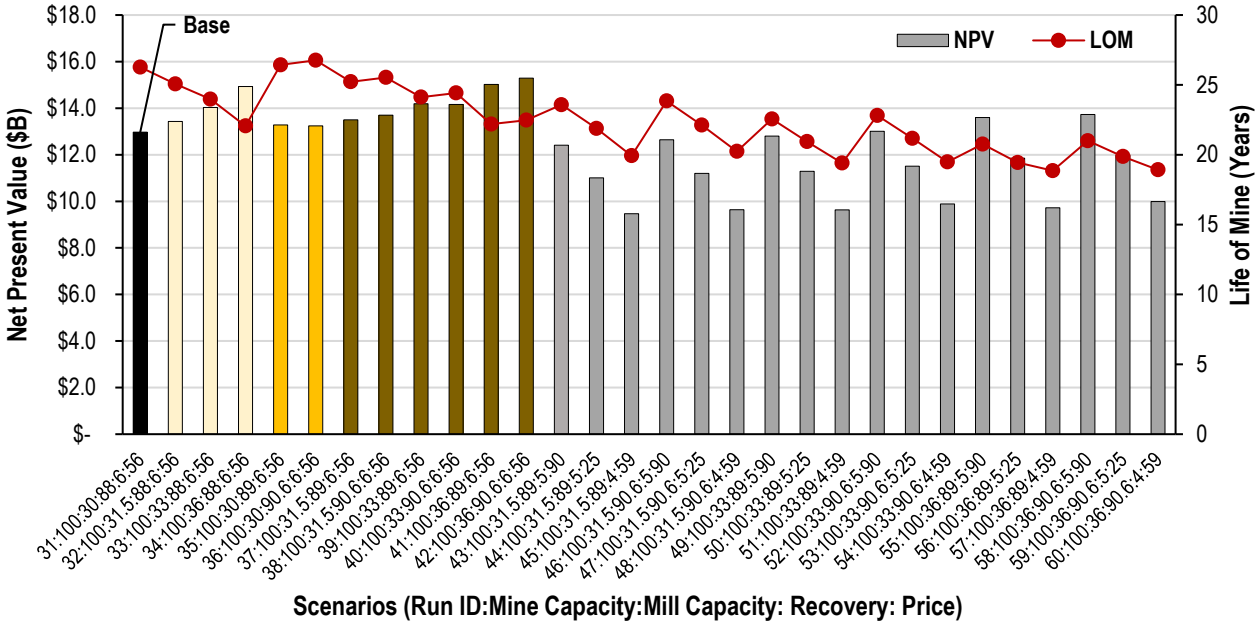


FIG 6 – NPV and LOM variations for 30 scenarios if adapting Pit → Process Plant (colours represent four simulation groups; base case shown in black).

Effect of throughput improvement (Group 1 – Scenarios 32 to 34)

Increasing mill throughput by 5 per cent, 10 per cent, and 20 per cent at fixed recovery and price yields NPV of $13.43B, $14.03B, and $14.93B respectively, with mine life shortening to 25.07, 23.98, and 22.06 years. The NPV gains of $0.46B, $1.06B, and $1.96B relative to the base case are progressive and increasingly significant, with the 20 per cent throughput improvement gaining about $2B in NPV through scheduling optimisation alone at zero incremental CAPEX. Mine life shortening is a positive outcome here as processing more tonnes annually against a fixed ore mining accelerates cash flow generation and reduces time-discounting effect, directly improving NPV.

The inverse relationship between throughput and mine life confirms that the Pit → Plant strategy, by sequencing softer ore domains earlier in the schedule, can meaningfully limit mine life and improve value without capital investment. A 20 per cent throughput improvement brings the Pit → Plant NPV to $14.93B, closing the gap with the Plant → Pit base case of $16.02B to $1.09B, achievable by smarter planning.

Effect of recovery improvement (Group 2 – Scenarios 35 to 36)

Improving copper recovery to 89.0 per cent and 90.6 per cent at fixed throughput and price yields NPV of $13.28B and $13.24B respectively, which delivers gains of $0.31B and $0.27B above the base case. Unexpectedly, the higher recovery of 90.6 per cent produces a marginally lower NPV than 89.0 per cent ($13.24B versus $13.28B), as the added ore reserve extended the LOM from 26.43 to 26.76 years. This counterintuitive result is worthy of attention as at fixed throughput, a higher recovery rate increases the economic value of marginal ore blocks, potentially accounting for lower-grade material into the optimal schedule and extending the LOM as dilutes the NPV benefit of the recovery improvement itself. This suggests that recovery improvement in isolation, without an accompanying throughput gain, most likely delivers limited incremental value under the Pit → Plant strategy.

Combined throughput and recovery improvement (Group 3 – Scenarios 37 to 42)

Combining throughput and recovery improvements produces the strongest results within the fixed price scenarios. At 31.5 Mt/a and 90.6 per cent recovery (Scenario 38), NPV reaches $13.70B. At 33 Mt/a and 89.0 per cent recovery (Scenario 39), NPV rises to $14.19B. The peak within this group is Scenario 42 where 36 Mt/a milling and 90.6 per cent recovery got $15.29B over 22.47 years, a gain of $2.32B above the Pit → Plant base case and only $0.73B below the Plant → Pit base case of $16.02B. This is the most important result in the fixed-price analysis as it shows that combining soft-domain sequencing for throughput uplift with feed grade control for recovery improvement can recover $2.32B in NPV at near-zero incremental capital, closing approximately 76 per cent of the value gap between the two strategies' base cases. The LOM reduction to 22.47 years confirms that both value mechanisms that accelerated cash flow and improved recovery per tonne are adding value simultaneously, producing a compounding benefit which is opposite to the compounding penalty observed in Group 3 of the Plant → Pit analysis.

Combined throughput, recovery and price (Group 4 – Scenarios 43 to 60)

Introducing copper price stress of 10 per cent, 20 per cent, and 30 per cent across throughput and recovery combinations tests the resilience of the Pit → Plant value gains under market downside. At $5.90/lb, NPV ranges from $12.41B to $13.73B depending on operational parameters, with the best outcome at 36 Mt/a and 90.6 per cent recovery (Scenario 58, $13.73B) exceeding the Pit → Plant base case NPV of $12.97B. These outcomes suggest that the throughput and recovery gains of the Pit → Plant strategy are sufficient to defend and exceed base case value even under a 10 per cent price reduction. At $5.25/lb, NPV ranges from $11.01B to $11.97B, and at $4.59/lb from $9.47B to $9.99B across valid scenarios.

Strategic comparison between Plant → Pit versus Pit → Plant

Figures 7 and 8 presents a direct comparison of NPV and LOM outcomes across all 60 simulated scenarios for both integration strategies, illustrating how NPV varies in response to applied changes in milling capacity, copper recovery, and copper price under each approach.

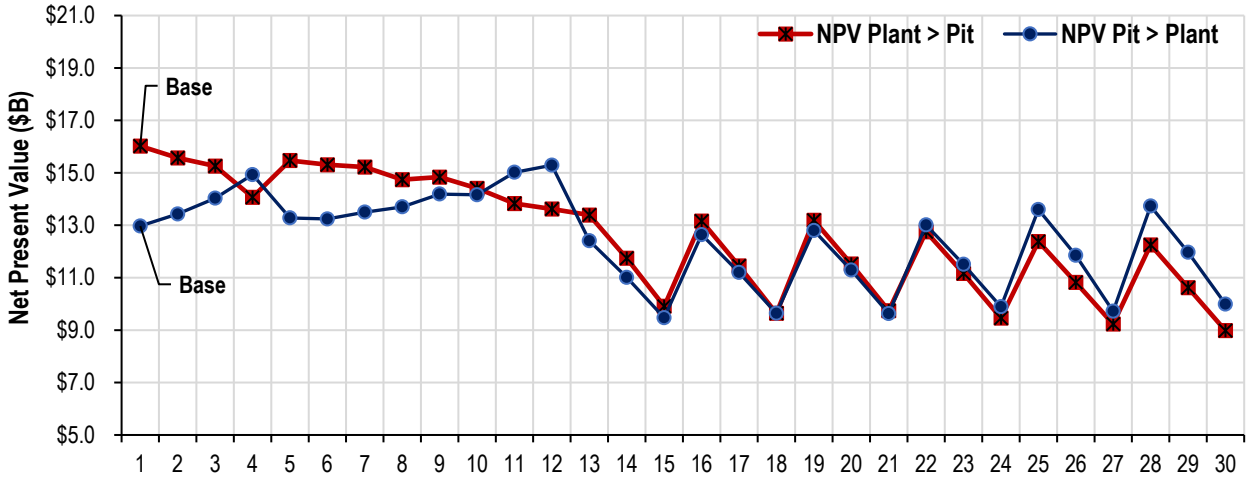


FIG 7 – Comparing NPVs for Plant → Pit (East) and Pit → Plant (North) scenarios.

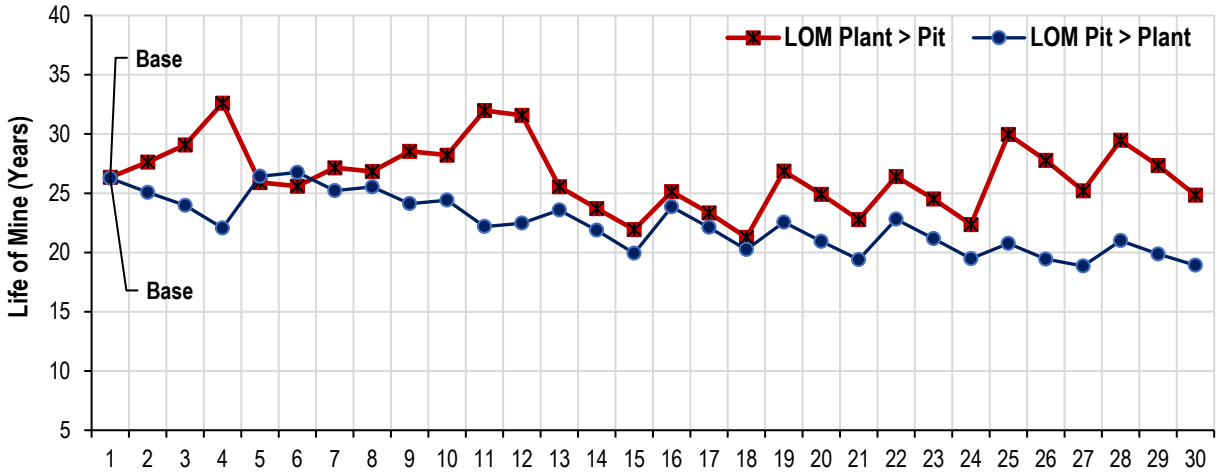


FIG 8 – Comparing LOMs for Plant → Pit (East) and Pit → Plant (North) scenarios.

The two integration strategies are closely compared based on the outcomes across seven key evaluation areas: a) base case NPV and mining direction, b) upside potential, c) downside risk, d) sensitivity to milling throughput, e) sensitivity to copper recovery, f) resilience under copper price stress, g) and life-of-mine implications.

The key objective of this comparison is to provide a focused basis for assessing the relative merits and limitations of each strategy under varying operational and market conditions. It is acknowledged, however, that further investigation incorporating additional deposit types, operational contexts, and geo-metallurgical complexity will be necessary to consolidate and strengthen these arguments.

1. Base case and mining direction

   Plant → Pit base case (Scenario 1) yields $16.02B NPV over 26.35 years which is $3.05B higher than the Pit → Plant base case (Scenario 31) at $12.97B over 26.27 years under identical operational parameters, confirming that eastward mining through higher-value ore domains delivers a materially superior starting position.

   This $3.05B advantage, however, is conditional on the plant performing at 30 Mt/a capacity. Any feed variability (hardness, PSD, mineralogy) could diminish it significantly.

2. Upside potential

   Without engineering flexibility into the process plant design Plant → Pit has no operational upside beyond its base case as all 29 scenarios represent NPV deterioration from $16.02B.

Pit → Plant best case (Scenario 42: 36 Mt/a, 90.6 per cent recovery, $6.56/lb) reaches $15.29B, assists with closing the gap with the Plant → Pit base case to just $0.73B through scheduling optimisation alone without deploying CAPEX. This represents a $2.32B NPV gain above the Pit → Plant base case. Based on the results, Pit → Plant offers genuine and achievable upside over Plant → Pit.

3. Downside risk

   Plant → Pit worst case (Scenario 30: 24 Mt/a, 83.6 per cent recovery, $4.59/lb) yields $8.98B, which is a $7.04B reduction from its base case, representing 44 per cent value loss under maximum combined stress.

   Pit → Plant worst case (Scenario 60: 36 Mt/a, 90.6 per cent recovery, $4.59/lb) yields $9.99B, which is a $2.98B reduction from its base case, representing 23 per cent value loss under equivalent price stress.

   Accordingly, Pit → Plant worst case exceeds Plant → Pit worst case by $1.01B, confirming superior downside resilience. Considering total NPV range across all scenarios, results show the difference between the best and worst cases in Plant → Pit is $7.04B; while for Pit → Plant the difference is $5.82B. With that, the results indicate that the Plant → Pit strategy exhibits greater NPV variance across the simulated scenario space.

4. Sensitivity to throughput

   Plant → Pit presents a 20 per cent milling throughput reduction results in $1.95B in NPV loss (Scenarios 1 to 4) while extending the LOM from 6.26 years to 32.61 years, compounding the discount penalty over a longer period.

   Pit → Plant shows a 20 per cent throughput increase gains $1.96B in NPV (Scenarios 31 to 34) while expectedly shortening the LOM by 4.21 years to 22.06 years, improving cash flows.

   Overall, the scenario outcomes show that milling throughput is a significant value lever in both strategies, generating equivalent magnitude effects in opposite directions.

5. Sensitivity to recovery

   Plant → Pit recovery losses of 3 per cent and 5 per cent reduce NPV by $0.55B and $0.71B respectively, which is meaningful but secondary to throughput.

   Pit → Plant recovery gains of 1 per cent and 3 per cent add only $0.31B and $0.27B respectively, which is relatively negligible, and counterintuitively diminishing at higher recovery levels due to marginal ore inclusion extending the LOM and increasing discounting. Recovery is a secondary value driver in both strategies.

6. Resilience to price volatility

   At $4.59/lb, Plant → Pit NPV ranges from $8.98B to $9.91B and for Pit → Plant ranges from $9.47B to $9.99B in a way the ranges overlap and converge, with Pit → Plant outperforming Plant → Pit across the majority of equivalent operational combinations.

   Under severe price stress, the Plant → Pit NPV advantage over Pit → Plant effectively disappears, while Pit → Plant maintains superior value through shorter LOM and improved cash flow generation.

7. Life-of-Mine (LOM)

   Plant → Pit LOM extends under operational stress, reaching 32.61 years under maximum throughput reduction, which increases pit utilisation and long-term commodity price exposure. It is worth noting that all Plant → Pit scenarios tested in this paper are based on the assumption that no engineering flexibility was incorporated into the process plant design to accommodate the feed variability inherent to this integration strategy.

   Pit → Plant LOM consistently shortens as performance improves, reaching 22.06 years at maximum throughput (Scenario 34), which potentially offers a better capital efficiency. The Pit

→ Plant strategy always allows for processing more ore and provides better operational flexibility over the LOM.

CONCLUSIONS

Mining is inherently probabilistic, and value losses arising from disconnected upstream and downstream planning cannot be recovered through operational adjustments alone. This paper argued that strategic integration choices made at the SMP stage, not short-term operational initiatives, are the primary drivers of long-term value realisation. Strategic Mine to Mill was proposed as a risk-informed life-of-mine value framework, evaluated through a 60-scenario simulation on an open pit copper deposit using GEOVIA Whittle®. Two integration strategies Plant → Pit and Pit → Plant were discussed and assessed across variations in milling throughput, copper recovery, and copper price. Based on the limited results given in this study, key conclusions are as follows:

- Plant → Pit yields a superior base case NPV of $16.02B, however this advantage is entirely depending on the process plant performance, offering no operational upside, and exposes the operation to a maximum value at risk of $7.04B under combined stress, caused the highest NPV variance of the two strategies.
- Pit → Plant, starting $3.05B below Plant → Pit, showed $2.32B of recoverable upside through scheduling optimisation without deploying CAPEX, closing 76 per cent of the value gap, with a worst-case NPV $1.01B superior to Plant → Pit and lower overall variance.
- Milling throughput is the dominant value lever in both strategies. Recovery is secondary and exhibits diminishing returns in isolation. Overall, Pit → Plant inherently enables greater ‘improvement opportunities’ in the downstream.
- Under severe copper price stress, which is an intrinsic characteristic of commodity markets, the Plant → Pit NPV advantage disappears, with both strategies converging, while Pit → Plant maintains superior value through shorter LOM and accelerated cash flows.

Strategic Mine to Mill assists in further exploring the missing links between pit-level decisions and downstream performance. It reframes Mine to Mill as a strategic, risk-informed value system rather than a short-term, downstream optimisation practice. Considering the interconnected nature of mining operations, developing more value-driven strategies through incorporating more levels of sophistication into evaluation practices would be critical to long-term and sustainable value realisation. Future investigation incorporating broader deposit types and explicit geo-metallurgical variability will further strengthen these findings.

ACKNOWLEDGEMENTS

The author thanks Dassault Systèmes for sponsoring the presentation of this work at the Strategic Mine Planning Conference 2026. Sincere appreciation is also extended to the conference organisers and reviewers for their efforts in bringing the mine planning community together and providing a valuable platform for the exchange of knowledge.

REFERENCES

Antofagasta PLC, 2023. Centinela Second Concentrator Project Approved for Development, Press Release, 20 December 2023. Available from: <https://www.antofagasta.co.uk/investors/news/2023/centinela-second-concentrator-project-approved-for-development>

Burger, B, McCaffery, K, Jankovic, A, Valery, W and McGaffin, I, 2006. Batu Hijau model for throughput forecast, mining and milling optimisation and expansion studies, SME 2006 Conference, St Louis, USA.

Bye, A, 2006. The strategic and tactical value of a 3D geotechnical model for mining optimization, Anglo Platinum, Sandsloot open pit, *Journal of the Southern African Institute of Mining and Metallurgy*, 106(2):97–104.

Bye, A, 2011. Case studies demonstrating value from geometallurgy initiatives, in Proceedings of the First International Geometallurgy Conference 2011, GeoMet 2011 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Capstone Copper Corp, 2025. Capstone Copper Announces Sanctioning of Mantoverde Optimized Project, Press Release. Available from: <https://capstonecopper.com/news/capstone-copper-announces-sanctioning-of-mantoverde-optimized-project/>

Carrasco, C, Keeney, L, Napier-Munn, T J and Bode, P, 2017. Unlocking additional value by optimising comminution strategies to process Grade Engineering® streams, *Minerals Engineering*, 103–104:2–10.

Dimitrakopoulos, R, 2018. Stochastic mine planning – Methods, examples and value in an uncertain world, *Advances in Applied Strategic Mine Planning*, 101–115.

Faramarzi, F, Jokovic, V, Morrison, R and Kanchibotla, S S, 2018. Quantifying variability of ore breakage by impact–Implications for SAG mill performance, *Minerals Engineering*, 127:81–89.

Faramarzi, F, Kanchibotla, S S and Morrison, R, 2019. Simulating the impact of ore competence variability on process performance – Case study of a large copper mine SAG Conference, Vancouver, Canada.

Faramarzi, F, Napier-Munn, T, Morrison, R and Kanchibotla, S S, 2020. The extended drop weight testing approach–What it reveals, *Minerals Engineering*, 157:106550.

Faramarzi, F, Amini, E, Bolden, L, Beaton, N and Koh, E, 2022. Scenario-based evaluation of potential value chain gains using Integrated Extraction Simulator – Mt Keith Nickel West case study, in *Proceedings of the IMPC Asia-Pacific 2022*, pp 1150–1163 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Faramarzi, F and Smith, R, 2023a. How does dynamic behaviour of comminution circuits impact strategic planning of mines?, 8th International Conference on Geology and Mine Planning, GEOMIN Mine Planning 2023, Santiago, Chile.

Faramarzi, F and Smith, R, 2023b. Forecasting value-chain performance for strategic and operational evaluations, SAG Conference 2023, Vancouver, Canada.

Godoy, M, 2018. A risk analysis-based framework for strategic mine planning and design—Method and application, in *Advances in Applied Strategic Mine Planning,* pp 75–90 (Springer).

Grundstrom, C, Kanchibotla, S, Jankovich, A, and Thornton, D, 2001. Blast fragmentation for maximising the sag mill throughput at Porgera Gold Mine, in *Proceedings of the 27th Annual Conference on Explosives and Blasting Technique*, pp 383–399 (The International Society of Explosives Engineers: Cleveland).

Hart, S, Valery, W, Clements, B, Reed, M, Song, M and Dunne, R, 2001. Optimisation of the Cadia Hill SAG mill circuit, in *Proceedings of the International Conference on Autogenous and semi-Autogenous Grinding Technology, SAG 2001*, pp 11–30 (University of British Columbia: Vancouver).

Kanchibotla, S S, Valery, W and Morrell, S, 1999. Modelling fines in blast fragmentation and its impact on crushing and grinding, Proceedings of Explo '99 – A Conference on Rock Breaking (The Australasian Institute of Mining and Metallurgy: Melbourne; and Western Australian School of Mines).

Kanchibotla, S S, Vizcarra, T G, Musunuri, S A R, Tello, S, Hayes, A and Moylan, T, 2015. Mine to mill optimisation at paddington gold operations, in *Proceedings of the International Conference on Autogenous and Semi-Autogenous Grinding Technology*, *SAG 2015*, pp 20–24.

Kanchibotla, S, Morrell, S, Valery, W and O'Loughlin, P, 1998. Exploring the effect of blast design on SAG mill throughput at KCGM, Mine to Mill Conference.

Kojovic, T, Michaux, S and McKenzie, C, 1995. Impact of blast fragmentation on crushing and screening operations in quarrying, Proceedings of Explo '95 Conference (The Australasian Institute of Mining and Metallurgy: Melbourne).

Lam, M, Jankovic, A, Valery, W and Kanchibotla, S, 2001. Maximising SAG mill throughput at Porgera Gold Mine by optimising blast fragmentation, in *Proceedings of the International Conference on Autogenous and semi-Autogenous Grinding Technology, SAG 2001*, pp 271–288 (University of British Columbia: Vancouver).

McKee, D J, 2013. *Understanding mine to mill*, Cooperative Research Centre for Optimising Resource Extraction (CRC ORE). Available from: <https://www.911metallurgist.com/C/What-is-Mine-to-Mill.pdf>

Morales, N, Seguel, S, Cáceres, A, Jélvez, E and Alarcón, M, 2019. Incorporation of geometallurgical attributes and geological uncertainty into long-term open pit mine planning, *Minerals*, 9(2):108.

Poblete, C J, Gonzlez, M A, Romero, J A, Fuentes, D L and Abdrashitova, O, 2016. Use of Robust Design Methodology for Production Scale Definition in Open Pit Mining, in *Proceedings Ninth AusIMM Open Pit Operators' Conference 2016*, pp 284–291 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Rimélé, A, Dimitrakopoulos, R and Gamache, M, 2020. A dynamic stochastic programming approach for open pit mine planning with geological and commodity price uncertainty, *Resources Policy*, 65:101570.

Smith, R, Faramarzi, F and Poblete, C, 2022. Strategic and tactical mine planning considering value chain performance for maximised profitability, in *Proceedings of the IMPC Asia-Pacific 2022*, pp 1204–1225 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Strohmayr, S and Valery, W, 2001. SAG mill circuit optimisation at Ernest Henry Mining, in *Proceedings of the International Conference on Autogenous and semi-Autogenous Grinding Technology, SAG 2001*, pp 11–42 (University of British Columbia: Vancouver).

Valery, W, Duffy, K and Jankovic, A, 2019. Mine to Mill optimization, *SME Mineral Processing and Extractive Metallurgy Handbook*, pp 335–343.

Young, M, Pease, J, Johnson, N and Munro, P, 1997. Developments in milling practice at the lead/zinc concentrator of Mount Isa Mines Limited from 1990, in *Proceedings of the AusIMM Sixth Mill Operators Conference*, 20 p (The Australasian Institute of Mining and Metallurgy: Melbourne).

APPENDIX

TABLE A1

Simulation results for Plant → Pit integrated strategy (mining direction: east).

No.	Mining capacity (Mt/a)	Milling capacity (Mt/a)	Copper recovery (%)	Copper price ($/lb)	NPV ($B)	Life-of-mine (years)
1	100	30.00	88.00	6.56	16.02	26.35
2	100	28.50	88.00	6.56	15.57	27.63
3	100	27.00	88.00	6.56	15.26	29.08
4	100	24.00	88.00	6.56	14.07	32.61
5	100	30.00	85.40	6.56	15.47	25.91
6	100	30.00	83.60	6.56	15.31	25.60
7	100	28.50	85.40	6.56	15.22	27.15
8	100	28.50	83.60	6.56	14.74	26.84
9	100	27.00	85.40	6.56	14.84	28.54
10	100	27.00	83.60	6.56	14.41	28.22
11	100	24.00	85.40	6.56	13.83	31.98
12	100	24.00	83.60	6.56	13.62	31.58
13	100	28.50	85.40	5.90	13.39	25.56
14	100	28.50	85.40	5.25	11.74	23.71
15	100	28.50	85.40	4.59	9.91	21.93
16	100	28.50	83.60	5.90	13.16	25.13
17	100	28.50	83.60	5.25	11.45	23.34
18	100	28.50	83.60	4.59	9.64	21.27
19	100	27.00	85.40	5.90	13.19	26.86
20	100	27.00	85.40	5.25	11.52	24.91
21	100	27.00	85.40	4.59	9.74	22.77
22	100	27.00	83.60	5.90	12.75	26.41
23	100	27.00	83.60	5.25	11.16	24.52
24	100	27.00	83.60	4.59	9.46	22.36
25	100	24.00	85.40	5.90	12.38	29.96
26	100	24.00	85.40	5.25	10.82	27.77
27	100	24.00	85.40	4.59	9.23	25.21
28	100	24.00	83.60	5.90	12.25	29.47
29	100	24.00	83.60	5.25	10.62	27.34
30	100	24.00	83.60	4.59	8.98	24.84

TABLE A2

Simulation results for Pit → Plant integrated strategy (mining direction: north).

No.	Mining capacity (Mt/a)	Milling capacity (Mt/a)	Copper recovery (%)	Copper price ($/lb)	NPV ($M)	Life-of-mine (years)
31	100	30.00	88.00	6.56	12.97	26.27
32	100	31.50	88.00	6.56	13.43	25.07
33	100	33.00	88.00	6.56	14.03	23.98
34	100	36.00	88.00	6.56	14.93	22.06
35	100	30.00	89.00	6.56	13.28	26.43
36	100	30.00	90.60	6.56	13.24	26.76
37	100	31.50	89.00	6.56	13.50	25.22
38	100	31.50	90.60	6.56	13.70	25.53
39	100	33.00	89.00	6.56	14.19	24.12
40	100	33.00	90.60	6.56	14.16	24.42
41	100	36.00	89.00	6.56	15.02	22.19
42	100	36.00	90.60	6.56	15.29	22.47
43	100	31.50	89.00	5.90	12.41	23.58
44	100	31.50	89.00	5.25	11.01	21.88
45	100	31.50	89.00	4.59	9.47	19.93
46	100	31.50	90.60	5.90	12.64	23.85
47	100	31.50	90.60	5.25	11.20	22.12
48	100	31.50	90.60	4.59	9.64	20.24
49	100	33.00	89.00	5.90	12.80	22.56
50	100	33.00	89.00	5.25	11.29	20.94
51	100	33.00	89.00	4.59	9.63	19.40
52	100	33.00	90.60	5.90	13.01	22.81
53	100	33.00	90.60	5.25	11.51	21.17
54	100	33.00	90.60	4.59	9.89	19.48
55	100	36.00	89.00	5.90	13.60	20.76
56	100	36.00	89.00	5.25	11.85	19.44
57	100	36.00	89.00	4.59	9.72	18.86
58	100	36.00	90.60	5.90	13.73	21.00
59	100	36.00	90.60	5.25	11.97	19.87
60	100	36.00	90.60	4.59	9.99	18.92

Assessing the impact of nonlinear combinations from the geometallurgical models and geological uncertainty within the strategic mine planning

G Gariba[1], D Alegre[2], M Faria[3], A Mayrink[4], L C Andrade[5], R Peroni[6] and J L V Mariz[7]

1. Consultant, KPI Mining, Belo Horizonte MG, Brazil. Email: gabriel.gariba@kpimining.com
2. Consultant, KPI Mining, Belo Horizonte MG, Brazil. Email: douglas.alegre@kpimining.com
3. Director of Technical Services, KPI Mining, Saint-Laurent QC, Canada. Email: matheus.faria@kpimining.com
4. Consultant, KPI Mining, Belo Horizonte MG, Brazil. Email: alisson.mayrink@kpimining.com
5. Consultant, KPI Mining, Saint-Laurent QC, Canada. Email: laura.andrade@kpimining.com
6. Professor, Mining Engineering Department, Universidade Federal do Rio Grande do Sul, Porto Alegre RS, Brazil. Email: peroni@ufrgs.br
7. Professor, Mining Engineering Department, Federal University of Pernambuco, Recife PE, Brazil. Email: jorge_valenca@hotmail.com

ABSTRACT

Strategic mine planning production forecasts and economic outcomes depend on spatial distribution of grades and contacts between different ore types, complex recovery functions, cut-off policies governing flows to multiple destinations, product specifications, and other operational constraints. Conventional mine planning techniques often misrepresent ore tonnages, metal production, recovery performance, and financial outcomes, leading to poor reconciliation between forecasts and actuals. This limitation arises because these approaches typically rely on single, smoothed estimated resource models, which are unable to adequately integrate geological uncertainty and variability, or to capture the non-linear relationships inherent in recovery estimations. This study evaluates the impact of geological uncertainty and recovery non-linearity on mine planning performance in a gold deposit characterised by sulfide and oxide mineralisation. A set of equiprobable joint geostatistical simulations of multiple grades and material types was generated to capture spatial variability and correlations. Using these simulations, a three-step methodology was implemented. First, these simulations were used to perform a risk analysis of a deterministic mine plan. Second, a stochastic mine plan was developed using a two-stage simultaneous optimisation framework. Finally, both deterministic and stochastic models were reconciled against grade control data using F1 factors, enabling the quantification of deviations in ore tonnage, metal content, recovery and production estimates. The risk analysis of the deterministic plan indicates a production risk with a P_{50} = -5.9 per cent in the life-of-mine, while recoveries are underestimated due to uncertainty in lithological contacts, resulting in significant production risks. This overestimation in production generates a cash flow risk of, in average, -13.6 per cent when compared with the deterministic value. In contrast, stochastic optimisation demonstrates the potential to increase expected net present value (NPV) by up to 10.7 per cent, primarily through improved sequencing decisions, an alternative optimised cut-off grade policy, that account for uncertainty and non-linear recovery behaviour. Comparison against grade control data further highlights differences between deterministic and simulated models. While average grades are relatively close for all the models, the application of recovery functions, economic parameters, and cut-off policies significantly amplifies discrepancies. The deterministic model underestimates recoveries, presenting an F1 factor of 87.0 per cent, while simulations show a better representation, ranging from 105.9 per cent to 107.1 per cent. Recovery factors vary between from 101.4 per cent to 101.8 per cent for simulations and 104.2 per cent for the deterministic model. These results demonstrate that the improvements achieved through stochastic optimisation are directly linked to a more realistic representation of geological variability and uncertainty. Incorporating these factors into strategic mine planning enhances forecast reliability, supports better decision-making, and ultimately maximises project value.

INTRODUCTION

Long-term mine planning aims to maximise project value and metal production through decisions based on economic forecasts, geological interpretations, and process performance models.

However, these decisions are commonly supported by limited sampling information and models that are subject to significant uncertainty and variability during actual operation. Plant recovery plays a critical role, since it is directly proportional to product masses, and, in some cases when it is badly estimated, can lead to misspecification of the scale of projects, generating serious value losses (Dunham and Vann, 2007).

Geometallurgical recovery estimation remains a major challenge in mining because recoveries are controlled by multiple geological and operational factors that are not always adequately represented in conventional resource models. Although geometallurgical testing can improve recovery prediction, its implementation is often limited by high testing costs and by challenges associated with sample support and scale effects (Deutsch, 2013; Lamberg, 2011; Batelochi, 2018). As recovery estimates are frequently derived from estimated grades and material classifications, uncertainty in geological variables propagates directly into uncertainty in recovery, recovered metal, and project cash flow. Furthermore, the non-linear behaviour of recovery functions amplifies this effect, meaning that a recovery calculated based on an average estimated grade can differ to the average recovery calculated based on the possible grade values, and relatively small variations in estimated grades may generate disproportionately large differences in predicted metallurgical performance. Current studies have been developed on the usage of different methodologies for representing uncertainty its impact on recoveries in mine planning (Jiang and Dimitrakopoulos, 2025; Coward and Dowd, 2015).

Several studies have been developed to improve recovery forecasting through geometallurgical modelling and advanced predictive approaches (Leal *et al,* 2016; Mendonça, Vieira and Costa, 2015; Costa *et al,* 2025). However, these methodologies are commonly calibrated using operational plant data and are rarely evaluated considering the spatial uncertainty and geological variability inherent to long-term mine planning models. As a result, their predictive performance may not fully represent the uncertainty conditions under which strategic planning decisions are made.

Stochastic mine planning emerges as an approach to explicitly incorporate uncertainty and non-linearity into the planning process. Two-stage simultaneous stochastic optimisation of mining complexes integrate anticipative and adaptive decision mechanisms to address uncertain geological and economic inputs. These methodologies rely on multiple equiprobable geological scenarios generated through geostatistical simulation techniques that reproduce the statistical distributions, spatial continuity, and geological relationships observed in the sampled data. By preserving variability and reducing the smoothing effects inherent to conventional estimation methods, stochastic approaches allow mine planning decisions to better represent operational uncertainty while maximising expected economic performance under production constraints (Goodfellow and Dimitrakopoulos, 2016, 2017). Works on stochastic mine planning have been developed to incorporate diverse sources of uncertainty and mine planning horizons, ranging long-term planning accounting for geological uncertainty (Quigley, Dimitrakopoulos and Grammatikopoulos, 2018) to short-term planning accounting for carbon emission factor uncertainty (Balboa-Espinoza and Dimitrakopoulos, 2026).

This study evaluates the application of a stochastic planning framework to improve confidence and robustness in long-term mine planning by incorporating geometallurgical variables. Geostatistical simulations were generated for domain contacts and grades and subsequently used to perform a risk analysis of a deterministic mine plan, quantifying the impact of geological uncertainty, and recovery non-linearity on production and economic performance. A stochastic optimisation approach was then applied to generate an alternative production schedule for the first ten years of the life-of-mine and assess potential improvements in plan performance. Finally, the deterministic and simulated models were reconciled against grade-control data through F1 factors calculated independently for the first year of operation, allowing the evaluation of how closely each approach reproduces actual operational outcomes. The results aim to quantify the contribution of geometallurgy, represented by the recovery curves, geostatistical simulations, and stochastic mine planning toward generating more reliable production forecasts and more resilient strategic mine plans.

METHODOLOGY

Reconciliation with short-term data

Model validation was performed through reconciliation against short-term operational data using F1 factors, a metric commonly applied to evaluate the accuracy of *in situ* reserve estimation and production forecasting (Parker, 2012; Menin, 2017). The F1 factor is calculated as:

$$F1 = \frac{Short\ range\ model\ depletions}{Long\ range\ model\ depletions} \quad (1)$$

In this study, grade-control data were adopted as the reference representation of actual production outcomes, while both the geostatistical simulations and the deterministic estimated model were treated as forecast scenarios. F1 factors were calculated independently for plant feed gold grade, recovered gold metal, total ore tonnage, and metallurgical recovery. An ideal F1 value is equal to 1.0, indicating perfect reconciliation match between forecasted and actual results. Deviations from this target may originate from multiple sources, including data uncertainty, sampling quality, change-of-support effects, limitations in geological estimation methodologies, and simplifications introduced in the mine planning process. Lower values indicate overestimation of the corresponding quantity, while higher values indicate underestimation.

The factors were calculated for each block model independently, in an area which encompasses around one year worth of grade control data, collected during mining operation. Each block for each of the models was classified as ore or waste based on a combination of lithology and a marginal block value assumptions. For reporting and comparative purposes, reconciliation results obtained from the simulation ensemble were summarised using P_{10}, P_{50}, and P_{90} statistics, allowing the evaluation of the variability and uncertainty associated with each forecasting approach.

The objective of this reconciliation analysis was to assess whether the results generated by calculating recoveries through simulations actually provide a better representation of the plant performance than the deterministic values, and if the values reported by a stochastic plan, due to the usage of simulations, are confirmed by the mine operation.

Stochastic plan

The stochastic production schedule was generated using a two-stage stochastic integer programming framework. This optimisation step uses multiple equiprobable geological simulations to generate a production schedule that is resilient to geological variability, maximising expected net present value (NPV) while simultaneously respecting operational constraints across all scenarios. The optimisation framework incorporates both strategic and adaptive decision mechanisms. Strategic decisions, such as extraction sequencing, are shared among all geological realisations, while operational responses are allowed to adapt according to the characteristics of each scenario. The optimisation is solved using an algorithm that combines simulated annealing techniques with advanced search and allocation procedures to balance value maximisation and compliance with production targets within a unified objective function.

The objective function adopted in the optimisation can be expressed as:

$$Max\left(\frac{1}{|\mathbb{S}|}\sum_{s\epsilon\mathbb{S}}\sum_{t\epsilon\mathbb{T}}\sum_{h\epsilon\mathbb{H}}\sum_{i\epsilon\mathcal{I}} p_{h,i,t} * v_{h,i,t,s} - \frac{1}{|\mathbb{S}|}\sum_{s\epsilon\mathbb{S}}\sum_{t\epsilon\mathbb{T}}\sum_{h\epsilon\mathbb{H}}\sum_{i\epsilon\mathcal{I}} c^{+}_{h,i,t} * d^{+}_{h,i,t,s} + c^{-}_{h,i,t} * d^{-}_{h,i,t,s}\right) \quad (2)$$

The first term represents the expected value optimisation component, where $s\epsilon\mathbb{S}$ represents each geological scenario in the set of available scenarios, $t\epsilon\mathbb{T}$ represents each period in the set of production periods, $h\epsilon\mathbb{H}$ are the hereditary attributes (variables of interest in locations of interest), $i\epsilon\mathcal{I}$ denotes each location in the set of locations in the mining complex(mines, stockpiles, plants and dumps), $p_{h,i,t}$ refer to each attribute's unitary economic contribution and $v_{h,i,t,s}$ refers to the attribute value. The second term incorporates target compliance through penalty functions. Variables $c^{+}_{h,i,t}$ and $c^{-}_{h,i,t}$represent deviations from upper and lower production targets, respectively, while $d^{+}_{h,i,t,s}$ and $d^{-}_{h,i,t,s}$ correspond to the penalties associated with exceeding or falling short these limits. Because penalties are evaluated independently for each geological realisation, the optimisation process seeks solutions that remain feasible and economically robust across the complete uncertainty space.

In practical terms, this formulation promotes the generation of risk-resilient mine plans by reducing sensitivity to geological uncertainty, minimising target violations across multiple possible outcomes, and prioritising extraction sequences that maximise expected value while maintaining operational stability. A more detailed mathematical formulation can be found in the referenced papers from Goodfellow and Dimitrakopoulos (2017).

Application in an open pit mine

Mining context

The study was conducted in an open pit gold mining complex characterised by sulfide and oxide mineralisation. Mining is performed through a conventional open pit operation, with sulfide material being classified into direct feed, medium-grade stockpiles, and low-grade stockpiles according to the long-term planning strategy. Oxide material is processed directly and is not stockpiled due to its relatively limited tonnage compared to the sulfide ore inventory. The long-term production plan adopted as the baseline for this study corresponds to a 10-year life-of-mine schedule developed using a conventional deterministic mine planning methodology. This plan served as the reference scenario for subsequent risk analysis, stochastic optimisation, and reconciliation assessments performed throughout the study.

Value chain modelling

A mathematical representation of the mineral value chain was developed using KPI-COSMO software. This model was applied consistently in both the risk analysis and stochastic optimisation stages and included the integrated representation of mining, stockpiling, and processing activities. The optimisation framework assumes that each block is extracted in a single mining period; however, material destinations may vary across stochastic scenarios according to the specific characteristics of each simulation. Material routing decisions were controlled through a cut-off grade policy that dynamically allocates mined material to the processing plant, sulfide stockpiles, or waste dump based on block grades, recoveries and economic criteria. This structure allows geological uncertainty and recovery variability to propagate through the production system and impact the resulting operational and economic performance. The configured value chain model used for the risk analysis and stochastic optimisation can be seen in Figure 1.

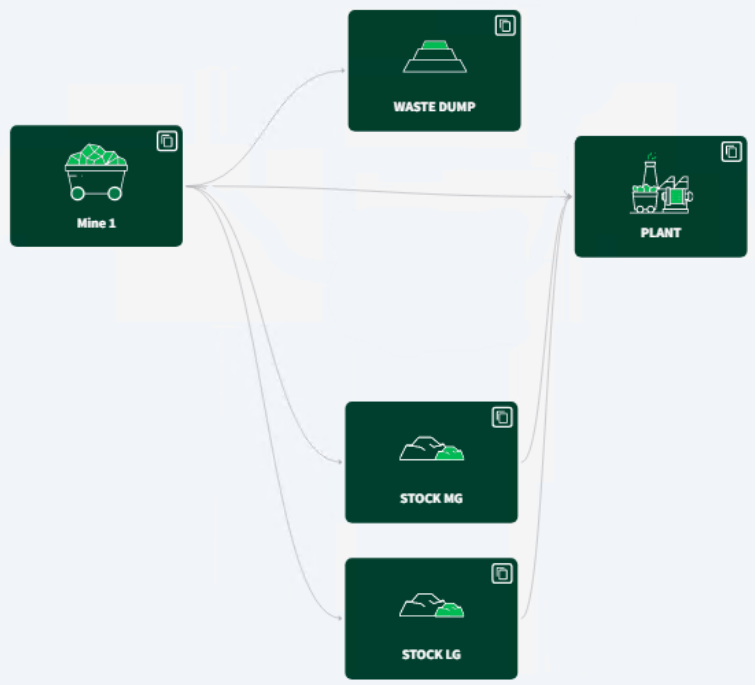


FIG 1 – Value chain model built for risk analysis and stochastic optimisation.

Geometallurgical recovery is expressed through a regression calculated independently for each block and each geological realisation using a logarithmic recovery function parameterised as a function of gold grades. This approach allows analysing geological variability to directly influence recovery estimates and consequently affect metal production and cash flow projections. Figure 2 shows the recovery curve used in this study, based on the Au grade. Mining costs were modelled as a function of depth through a constant incremental factor applied to deeper blocks. No phase constraints were incorporated into the optimisation process to avoid introducing additional biases into the comparison between deterministic and stochastic approaches, as the stochastic optimisation framework adopted in this study does not explicitly incorporate pushback or phase designs.

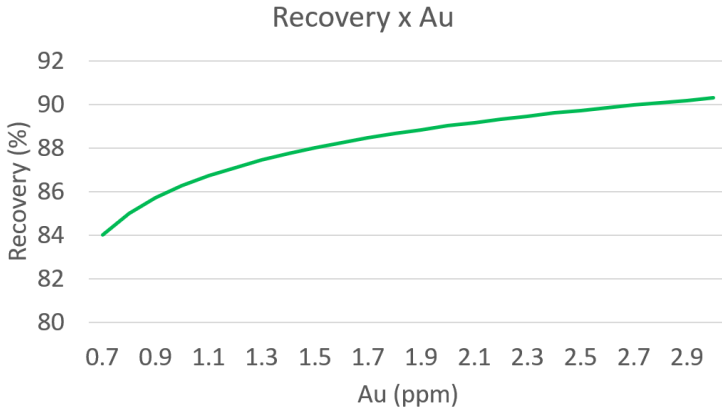


FIG 2 – Recovery x Au Grade curve used for calculating block recoveries.

The deterministic production schedule was evaluated using the value chain model and interrogated against the geological simulations to generate probabilistic forecasts of production and economic performance. The resulting distributions were summarised using P_{10}, P_{50}, and P_{90} statistics, providing a consistent basis for comparison with the stochastically optimised production schedule. The main performance indicators evaluated were plant feed gold grade, metallurgical recovery, recovered gold production, and economic performance. These metrics were aggregated over the complete life-of-mine to quantify the overall impact of geological uncertainty and recovery non-linearity on mine planning outcomes.

RESULTS AND DISCUSSIONS

Reconciliation with short-term data

Reconciliation against short-term operational data showed that the grade-control model reported a total ore tonnage 13.0 per cent lower than that estimated by the deterministic long-term model, resulting in an F1 factor of 87.0 per cent for the evaluated period. In contrast, the stochastic models produced substantially lower deviations, with a P_{50} F1 factor of 106.6 per cent, corresponding to an error of +6.6 per cent and an absolute deviation considerably smaller than that observed in the deterministic approach. The P_{10} and P_{90} of the stochastic models showed values of 105.9 per cent and 107.1 per cent, also with errors much lower than the deterministic value.

For gold grade reconciliation, the deterministic model achieved an F1 factor of 101.3 per cent for the blocks classified as ore, while the simulated models produced values ranging from 100.0 per cent to 102.7 per cent. Although the ranges reported by the simulations indicate some geological uncertainty, the simulated interval encompassed the deterministic estimate and included the 100 per cent value, which showed the capacity of the simulations to better represent the grade distributions for the analysed area when compared to the deterministic model.

Similar behaviour was observed for contained gold, where the deterministic model produced an F1 factor of 87.9 per cent, with a deviation of -12.1 per cent when compared to the grade control model. This can be attributed to the ore tonnage error, which had an impact on the total metal content estimation in the area. The simulated models, however, showed F1 factors varying between a P_{10} = 107.5 per cent and a P_{90} = 110.9 per cent. While still closer to the 100 per cent benchmark than the deterministic model, reconciliation error for contained gold was greater than the errors obtained

independently for ore tonnage and gold grade and can be explained by error propagation associated with the multiplication of tonnage and grade.

Finally, for the recovery, the deterministic model showed an F1 factor of 104.2 per cent, while the simulations ranged from a P_{10} = 101.4 per cent to a P_{90} = 101.8 per cent, with a P_{50} = 101.6 per cent. It was interesting to see that, even though the P_{50} of the gold grade was very similar to the deterministic model, the deterministic recovery was 2.5 per cent lower than the median given by the simulated models. This effect can be explained by the non-linearity of the recovery curves, and the impact that using average inputs can have when calculating geometallurgical variables, more than doubling the estimation error (from an error of +1.6 per cent when calculated on the P_{50} to +4.2 per cent obtained when using the deterministic value), solely due to the propagation of uncertainty through the recovery calculation.

Overall, the reconciliation results suggest that deterministic models provided overestimated values for ore tonnage and metal content, while underestimating grades and recoveries. The ranges given by the simulations (P_{10} to P_{90}) were closer to the benchmark (100 per cent) for all the measurements except for the grades, for which the deterministic values ware close to the probabilistic P_{50}. It is worth noting that the methodology showed that using simulations for approximating recoveries proved to be a better approach than estimating through a deterministic model, reducing the error propagation from using average inputs. The risk profiles present low variability for recovery values when analysing the area as a whole due to the analysis being performed in an area with higher grades, in which the recovery curve presents lower slope, although local variability shows pockets with expressive differences between models. Table 1 and Figure 3 summarise the F1 factors calculated for ore tonnage, grade, metal content and recovery.

TABLE 1

F1 factors for the deterministic model and simulations.

		Deterministic model	Stochastic model	
Material	**Quantity**	**Average**	**P_{50}**	**P_{10}–P_{90}**
Ore	Ore tonnage	87.0%	106.6%	105.9–107.1%
	Au grade	101.3%	101.5%	100.5–102.7%
	Au content	87.9%	109.3%	107.5–110.9%
	Recovery	104.2%	101.6%	101.4–101.8%

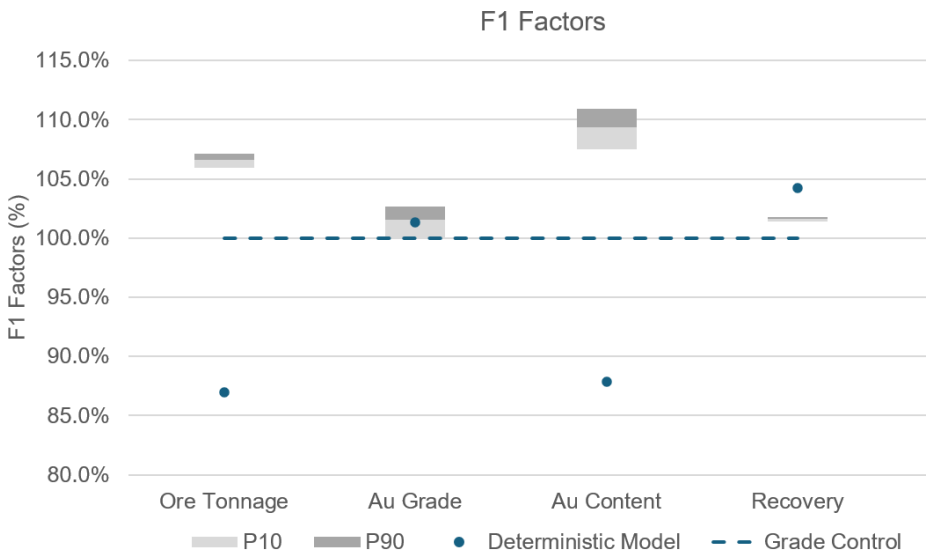


FIG 3 – F1 factors for the simulations and deterministic model.

Visual inspection of selected benches supported these observations and showed that, although deterministic and median simulated recoveries presented similar values across most areas of the deposit, localised regions exhibited substantial differences between the two approaches. In these areas, the simulations indicated that average recoveries could be significantly higher or lower than those predicted by the conventional deterministic estimation (Figure 4).

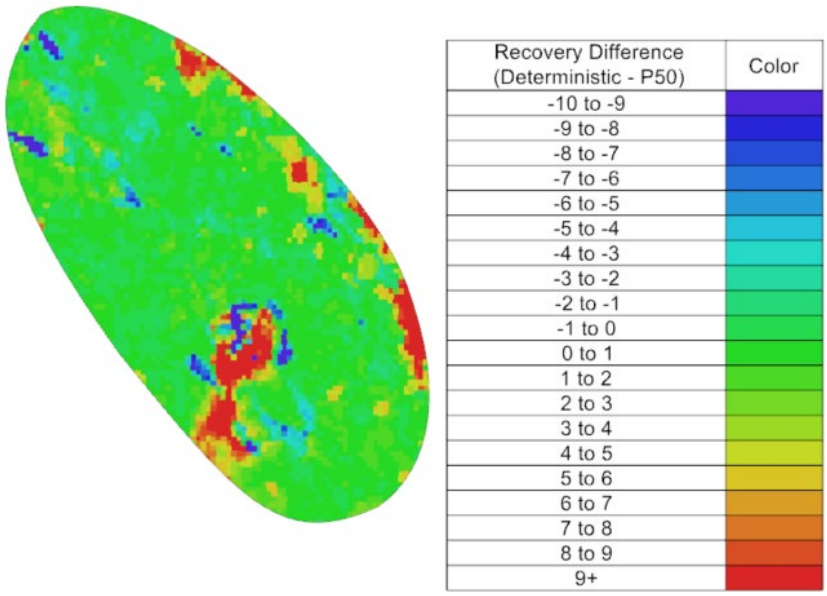


FIG 4 – Bench map of the absolute difference between deterministic recovery and the average recovery given by the simulations.

These localised discrepancies become particularly relevant when evaluated over localised domains, such as the area considered in the reconciliation study, where local variability has a greater influence on aggregated production results. Consequently, the inability of deterministic estimation to preserve this variability may explain the differences observed in recovery and production forecasts between the reconciliation approaches.

Stochastic optimisation of the mine schedule

By evaluating the deterministic mine plan across the individual geological simulations, it was possible to quantify a consistent downside probability in gold production when considering geological uncertainty. Compared with the deterministic baseline, probabilistic outcomes showed a production risk with a P_{50} = -5.9 per cent, with production deviations ranging from P_{10} = -7.4 per cent to P_{90} = -4.6 per cent. This reduction propagated directly into the economic performance of the operation, resulting in a cash flow risk with a P_{50} of -13.6 per cent (P_{10} = -15.1 per cent; P_{90} = -12.6 per cent).

These results highlight how moderate deviations in metal production may translate into substantially larger impacts on project economics. The amplification of economic risk can be attributed to the non-linear relationship between production, recoveries, operating costs, and revenue generation, causing uncertainty in geological inputs to propagate disproportionately into cash flow outcomes.

Application of the stochastic optimisation framework generated a more robust production schedule and resulted in improved operational and economic performance. Compared with the median results obtained from the risk analysis, the stochastic plan achieved a production increase with a P_{50} of +10.7 per cent, varying between P_{10} = +9.8 per cent and P_{90} = +11.8 per cent. Similarly, cash flow improvements ranged from P_{10} = +9.6 per cent to P_{90} = +12.0 per cent, with a P_{50} increase of +10.7 per cent.

Unlike the deterministic mine plan vulnerability highlighted in the risk analysis, where relatively small production losses generated amplified economic impacts, the gains obtained through stochastic optimisation showed a nearly proportional relationship between production and cash flow improvements. This behaviour is attributed to the high revenue contribution of additional gold production relative to operating costs, resulting in incremental production translating directly into

higher project value. Figure 5 shows a summary of the impact on the cash flow and gold in concentrate perceived in the risk analysis, and the gains generated during the stochastic optimisation step.

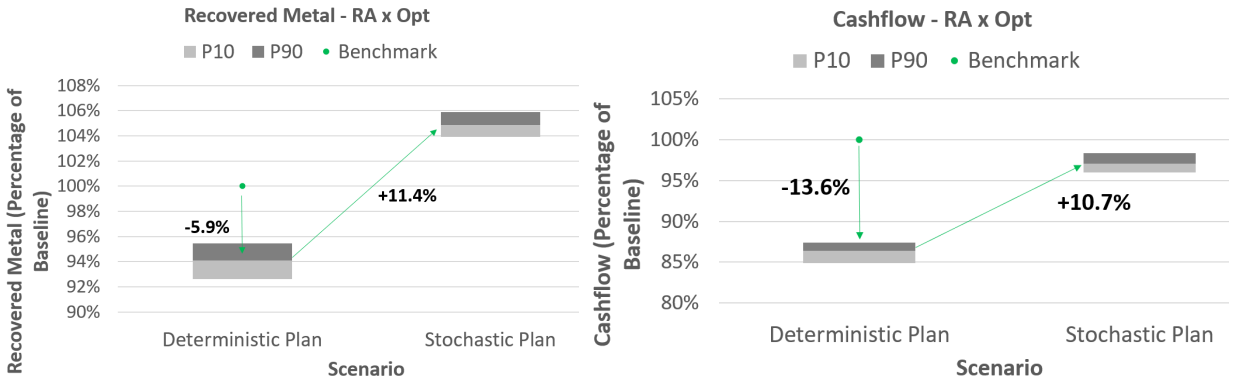


FIG 5 – Comparison of recovered gold production and cash flow between the deterministic baseline (green marker), simulated forecasts generated with the deterministic sequence (left), and stochastic optimisation results (right). Box plot limits represent P_{10} and P_{90} values, while the transition between the darker and lighter regions indicates the P_{50}.

In general, these results show that the stochastic optimisation was able to compensate the production risk in the deterministic plan by providing an expected range for gold production even higher than the original deterministic forecast. The resulting cash flow ranges for the stochastically optimised plan, however, continued lower than the original deterministic forecasts, mainly due to the elevated risk (P_{50} = -13.6 per cent), but the methodology was able to considerably mitigate the risk when compared to the forecasts generating by considering uncertainty in the deterministic plan through simulations.

CONCLUSIONS

The reconciliation study using grade-control data showed that the stochastic simulations presented value ranges (P_{10} to P_{90}) closer to the benchmark for all the measures taken, except for the gold grade in ore material, where the deterministic value approximated the simulated P_{50}. Even though the deterministic grade approximated the median of the simulations, substantial differences emerged when metallurgical recoveries were introduced through non-linear parametric recovery functions. The deterministic model showed an error in recovery estimation of +4.2 per cent, while simulation-based forecasts reduced this error to a range between +1.4 per cent and +1.8 per cent. These findings indicate that deterministic estimates amplified the error due to the non-linearity of the recovery curves, while simulations were able to better represent plant performance.

Evaluation of the risk of the deterministic mine plan through simulations, a significant exposure to geological uncertainty was identified, resulting in a consistent downside risk for both gold production and economic performance. The probabilistic forecasts generated on the deterministic schedule presented production deviations ranging from P_{10} = -7.4 per cent to P_{90} = -4.6 per cent and cash flow deviations ranging from P_{10} = -15.1 per cent to P_{90} = -12.6 per cent. The application of stochastic optimisation techniques reduced this exposure through the generation of a more risk-resilient extraction sequence, increasing expected production (P_{10} = +9.8 per cent; P_{90} = +11.8 per cent) and improving cash flow (P_{10} = +9.6 per cent; P_{90} = +12.0 per cent) relative to the average outcomes obtained in the deterministic risk analysis.

The comparison between deterministic and stochastic approaches further demonstrated that the smoothing effect inherent to conventional estimation methods becomes increasingly critical as uncertainty propagates through non-linear value chain calculations. While the deterministic model reproduced grades and tonnages with acceptable accuracy, the combination of smoothing and recovery non-linearity generated amplified biases in production and economic forecasts.

Overall, this study demonstrates that integrating geostatistical simulations into strategic mine planning allows uncertainty to be explicitly incorporated into production scheduling and economic

evaluation. The results indicate that stochastic optimisation not only generates mine plans with improved economic performance and lower exposure to geological risk, but also produces forecasts that are more consistent with operational reconciliation outcomes. These findings support the adoption of stochastic mine planning frameworks in operations where geological variability and non-linear geometallurgical responses significantly influence long-term value generation and decision-making.

ACKNOWLEDGEMENTS

The authors would like to acknowledge KPI Mining for providing the computational tools, methodologies, and technical environment that supported the development of this study, particularly for enabling the implementation of geostatistical simulation and stochastic mine planning workflows applied throughout the analysis.

The authors also express their gratitude to the Graduate program in Mining, Metallurgy and Materials of the Universidade Federal do Rio Grande do Sul in Brazil (PPGE3M/UFRGS) for the technical support, academic guidance, and access to data that made this research possible. The authors also acknowledge the CAPES/PROEX program that provides partial financial support for the researchers and experiments needed to develop this research. Their contribution was fundamental to the development, validation, and interpretation of the results presented in this work.

REFERENCES

Balboa-Espinoza, V and Dimitrakopoulos, R, 2026. Joint fleet management and stochastic short-term production Planning in mining complexes with carbon footprint reduction, *International Journal of Mining, Reclamation and Environment*, pp 1–22. https://doi.org/10.1080/17480930.2026.2628757

Batelochi, M, 2018. Simulação geoestatística e regressão multivariada aplicadas na modelagem geometalúrgica de um μdepósito de zinco (Geostatistical simulation and multivariate regression applied to the geometallurgical modelling of a zinc deposit), MSc dissertation, Universidade Federal do Rio Grande do Sul, Porto Alegre. Available from: <http://hdl.handle.net/10183/193390>

Costa, J F C L, Niquini, F G F, Schneider, C L, Alcântara, R M, Capponi, L N and Rodrigues, R S, 2025. Geometallurgical cluster creation in a niobium deposit using dual-space clustering and hierarchical indicator kriging with trends, *Minerals*, 15(7):755. https://doi.org/10.3390/min15070755

Coward, S and Dowd, P, 2015. Geometallurgical models for the quantification of uncertainty in mining project valuechains, in *Proceedings of 37th International Symposium on the Application of Computers and Operations Research in the Mineral Industry*, pp 360–369.

Deutsch, C V, 2013. Geostatistical modelling of geometallurgical variables—Problems and solutions, in *Proceedings of the Second AusIMM International Geometallurgy Conference, (GeoMet) 2013*, pp 7–16 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Dunham, S and Vann, J, 2007. Geometallurgy, geostatistics and project value—does your block model tell you what you need to know, in *Proceedings of the Project Evaluation Conference*, pp 189–196 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Goodfellow, R and Dimitrakopoulos, R, 2016. Global optimization of open pit mining complexes with uncertainty, *Applied Soft Computing*, 40:292–304. https://doi.org/10.1016/j.asoc.2015.11.038

Goodfellow, R and Dimitrakopoulos, R, 2017. Simultaneous Stochastic Optimization of Mining Complexes and Mineral Value Chains, *Mathematical Geosciences*, 49(3)341–360. https://doi.org/10.1007/s11004-017-9680-3

Jiang, Y and Dimitrakopoulos, R, 2025. An application of simultaneous stochastic optimisation on an open-pit copper mining complex with supply, recovery and market uncertainties, in *International Journal of Mining, Reclamation and Environment,* 39(1):74–92. https://doi.org/10.1080/17480930.2024.2381904

Lamberg, P, 2011. Particles-the bridge between geology and metallurgy, in Konferens i mineralteknik (Conference in Minerals Engineering), Luleå University of Technology.

Leal, R S, Peroni, R L, Costa, J F C L, Pereira, S G, Martins, R M and Capponi, L N, 2016. Geostatistics applied to geometallurgical modeling, in Proceedings of the 24th World Mining Congress, Brazil (IBRAM: Brasília).

Mendonça, A, Vieira, M and Costa, J F C L, 2015. Métodos geoestatísticos aplicados à modelagem geometalúrgica (Geostatistical methods applied to geometallurgical modelling), in Anais do XXVI Encontro Nacional de Tratamento de Minérios e Metalurgia Extrativa, Poços de Caldas, Brazil, pp 813–820.

Menin, R Q, 2017. Definição de espaçamento amostral para classificação de recursos minerais baseada em dados de reconciliação da produção (Defining sample spacing for mineral resource classification based on production reconciliation data), MSc dissertation, Universidade Federal do Rio Grande do Sul, Porto Alegre.

Parker, H M, 2012. Reconciliation principles for the mining industry, *Mining Technology*, 121(3):160–176. https://doi.org/10.1179/1743286312Y.0000000007

Quigley, M, Dimitrakopoulos, R and Grammatikopoulos, T, 2018. Risk-resilient mine production schedules with favourable product quality for rare earth element projects, *Mining Technology*, 127(1):41–55. https://doi.org/10.1080/14749009.2017.1323172

Debottlenecking underground stoping operations using an integrated Mine to Mill-to-backfill digital twin and optimisation model

A Kumar[1], R Goodfellow[2] and K Dempster[3]

1. Manager, Mine Optimization, Newmont, Montreal QC H8N2Z4, Canada.
   Email: ashish.kumar@newmont.com
2. Director, Mine Optimization, Newmont, Denver CO 80237, USA.
   Email: ryan.goodfellow@newmont.com
3. Lead Engineer, Long Term Planning, Newmont, Melbourne Vic 3000, Australia.
   Email: kali.dempster@newmont.com

ABSTRACT

Underground stoping operations are governed by tightly coupled constraints across mining, material handling, processing, tailings management and backfilling. Bottlenecks can emerge from interactions between mobile equipment capacity, ventilation availability, working-face temperature, haulage and hoisting limits, processing throughput, tailings handling and paste backfill readiness. These interactions evolve over time and are difficult to quantify using conventional long-term planning tools, which can result in schedules that appear feasible but underperform during execution. This paper presents an integrated, Newmont-developed Mine to Mill-to-backfill digital twin and strategic optimisation workflow for an operating underground longhole open stoping mine. The digital twin is calibrated using historical operating data and validated by replaying a historical plan, achieving simulated P50 delivery within 10 per cent of actual outcomes. Bottleneck insights from the digital twin are translated into explicit constraints in a strategic optimisation model that generates rapid life-of-mine schedules at quarterly scale. In the case study, the optimised schedule delivered an improvement up to 10 per cent in discounted cash flow relative to the base schedule. The quarterly strategic plan is then passed through a constraint programming based resource allocation model that preserves the annual strategic targets while adjusting the timing and sequence of tasks within each year to reflect resource availability and operational constraints. The resulting plan is then decomposed into shift based tasks and re tested in the digital twin. This final simulation step confirms whether the improvement in discounted cash flow remains achievable under realistic variability, resource constraints and system interactions.

INTRODUCTION

Underground stoping plans can appear robust in aggregate but underperform during execution when the interactions between mining, material movement, processing and backfill are not fully represented. In longhole open stoping operations, individual activities are often planned as discrete steps, while realised system performance depends on the dynamic interaction between development, stoping, haulage, hoisting, processing, tailings generation, paste preparation and backfill delivery. A constraint in one part of the value chain can delay dependent activities elsewhere, creating emergent bottlenecks that are difficult to identify using deterministic planning methods alone. Figure 1 outlines the system boundary represented in the digital twin, linking underground activity with material handling, processing, tailings handling and paste backfill return. The problem is both a simulation and optimisation challenge. The digital twin represents variability, queueing and dynamic system behaviour, while the strategic optimisation model generates schedules that maximise value while respecting practical execution constraints (Cacciuttolo *et al*, 2025). This paper presents a closed-loop workflow that links these capabilities to improve the executability of long-term underground plans.

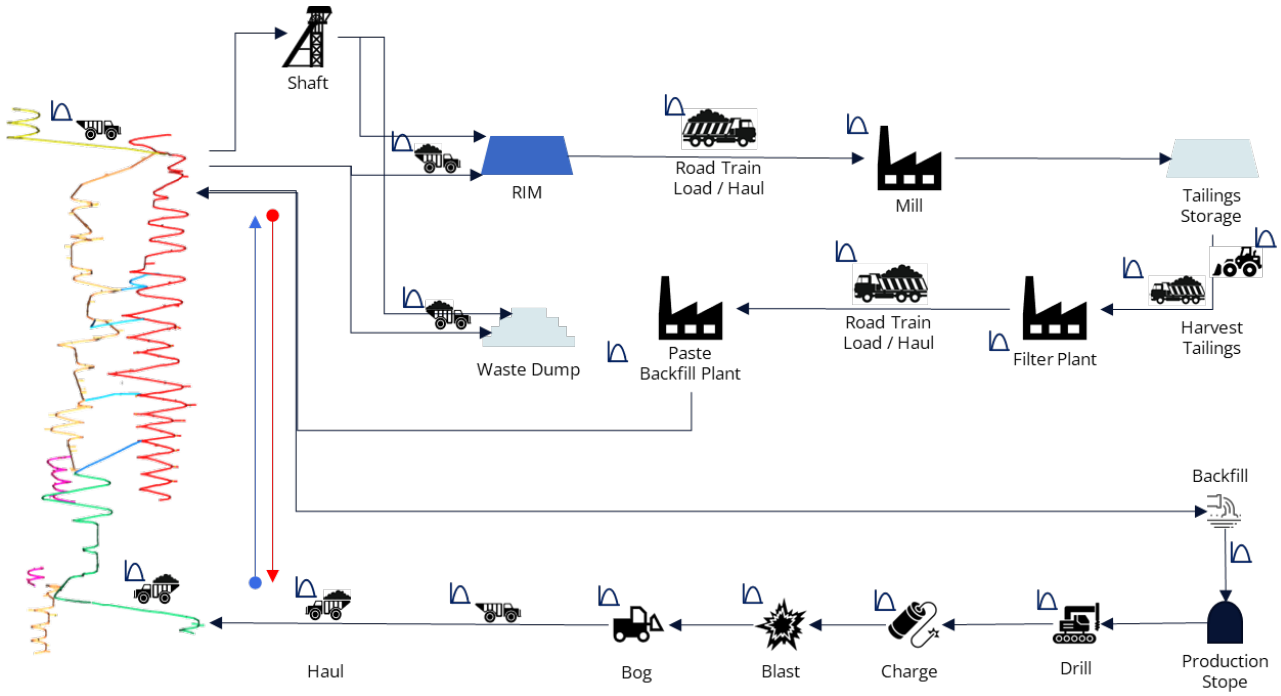


FIG 1 – Underground activity value chain represented in the digital twin, from stoping through to processing, tailings handling and paste backfill return.

INTEGRATED DIGITAL TWIN AND OPTIMISATION WORKFLOW

Workflow

The workflow starts with a long-term schedule, resource assumptions and historical operating data. These inputs are used to build and calibrate an integrated Mine to Mill-to-backfill digital twin that represents key activities, dependencies and constraints across the value chain. Figure 2 summarises the closed-loop workflow. The digital twin is first used to identify where the plan is constrained and which limits should be represented explicitly in optimisation. These insights are converted into planning constraints and passed into the strategic optimisation model.

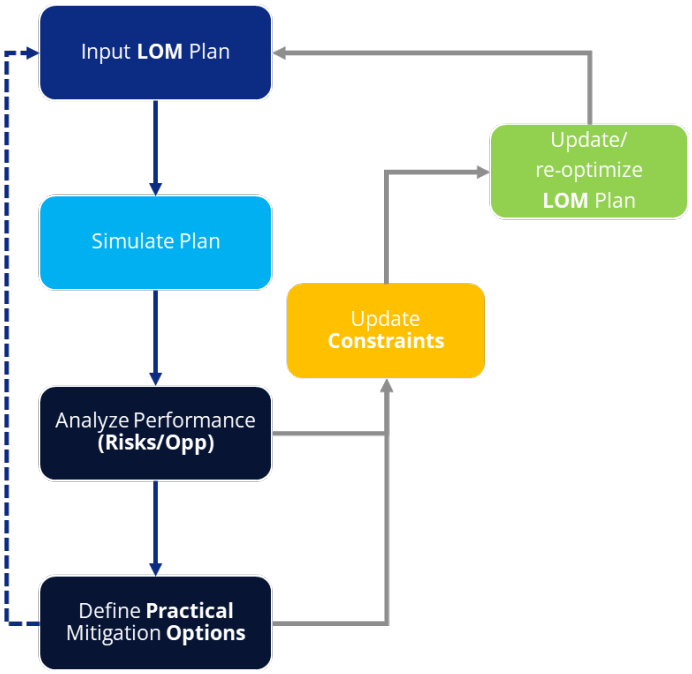


FIG 2 – Closed-loop workflow linking the integrated digital twin and optimisation model for debottlenecking underground stoping operations.

The strategic optimisation model uses a Bienstock and Zuckerberg based solve strategy to generate rapid life-of-mine schedules at quarterly scale (Bienstock and Zuckerberg, 2010; Muñoz *et al,* 2018). This quarterly strategic plan is then passed through a resource allocation optimisation model formulated as a constrained scheduling problem. The model adjusts the timing and sequence of activities within each year of the strategic plan and assigns constrained resources while enforcing weekly, monthly and annual quantity, quality, ventilation, heat and resource utilisation limits, consistent with constraint programming approaches for resource constrained scheduling and allocation (Rossi, van Beek and Walsh, 2006). The resource allocated plan is then decomposed into shift based tasks and re tested in the digital twin, where variability, queueing, resource interactions and operational constraints are reintroduced. This creates a closed loop process in which simulation, strategic optimisation and resource allocation optimisation inform each other rather than operating as separate planning steps.

Digital twin and calibration

The digital twin was developed for an operating underground longhole open stoping mine and captures the full cycle from underground activity through to surface processing and backfill return. It includes underground development and stoping tasks, secondary haulage, shaft hoisting or truck haulage to surface, processing, tailings generation, solar drying, paste production and paste backfill delivery underground. The model explicitly represents equipment rates, task sequencing, ventilation demand, working-face temperature impacts, and the dependency between processing performance, tailings availability and backfill readiness.

Calibration uses historical operating data, including equipment rates, cycle times, productivity and observed system behaviour. The objective is not only to match individual activity averages, but to reproduce how the full system behaves when activities, resources and constraints interact. The model is validated by replaying a historical plan through the digital twin and comparing simulated P50 delivery against actual outcomes. In the case presented, calibration within 10 per cent of actual performance provides confidence that the model captures the main drivers of delivery across the Mine to Mill-to-backfill value chain. Once calibrated, the digital twin is used to identify bottlenecks driven by haulage congestion, hoisting capacity, paste availability, local ventilation restrictions, temperature limits, processing constraints or task sequencing effects. These insights identify which constraints should be represented in the optimisation model.

Optimisation model

The strategic optimisation model converts bottleneck insights from the digital twin into explicit planning constraints and determines when each mining unit is extracted and where it is sent across the planning horizon. The model operates at quarterly scale for life-of-mine planning and seeks to maximise total discounted value while respecting the practical limits that govern underground execution.

$$max \sum_{i \in I} \sum_{t \in T} \sum_{d \in D} v_{i,t,d} \left(x_{i,t,d} - x_{i,t-1,d} \right)$$

where $x_{i,t,d}$, indicates whether mining unit i has been scheduled by time t to destination d, and $v_{i,t,d}$ represents the discounted value of extracting mining unit i in period t to destination d. The model is constrained by mining precedence and extraction sequence, single extraction and destination assignment, equipment and fleet capacity, ventilation and working-face temperature limits, processing and filter plant throughput, tailings and paste availability, and haulage effort constraints such as tonne-kilometres where relevant. The optimisation generated a schedule with an improvement of up to 10 per cent in discounted cash flow compared with the base schedule. After resource allocation, the plan is re-tested in the digital twin to confirm that this improvement remains achievable under more realistic operating conditions, including resource constraints, variability, queueing and system interactions, without creating new bottlenecks elsewhere in the value chain.

CONCLUSIONS

This paper presents a practical workflow for linking an integrated Mine to Mill-to-backfill digital twin with strategic optimisation and resource allocation optimisation to improve the executability of long-

term underground stoping plans. The approach uses calibrated simulation to identify system bottlenecks, converts those bottlenecks into explicit planning constraints, generates rapid life-of-mine strategic schedules at quarterly scale, applies resource allocation optimisation to produce an executable resource-constrained schedule, and then decomposes the resource-allocated plan into shift-based tasks for digital twin risk assessment. The case study demonstrates that the digital twin can reproduce historical delivery within 10 per cent of actual outcomes and that the optimised schedule can improve discounted cash flow up to 10 per cent relative to the base case. More importantly, the closed-loop workflow provides a structured method for confirming whether this value improvement remains achievable once resource limits, variability, queueing and full value-chain interactions are considered. This provides a more realistic basis for long-term planning, capital prioritisation and risk-informed decision-making across the Mine to Mill-to-backfill system. Future work could integrate supply uncertainty, including grades and geological domains, to test the robustness of optimised schedules under variable feed conditions.

ACKNOWLEDGEMENTS

The authors acknowledge the support of Newmont and the operational and technical teams who contributed data, domain knowledge and review input to the development of the integrated planning workflow. Company approval to publish this work is gratefully acknowledged.

DISCLAIMER

The technical and financial information presented in this paper are site-specific to Newmont operations and have been developed solely for internal operational management, academic discussion, and industry knowledge sharing. Nothing in this paper is to be considered as forward-looking observations and forecasting. These statements and models are based on current data, monitoring configurations, and geological assumptions, actual geomechanical responses, subsidence trends, and operational outcomes may differ materially due to localised geological anomalies and changes in systems or operational details. Newmont Corporation and the authors make no warranties, express or implied, regarding the applicability, suitability or outcomes of these methodologies when applied to other mining operations or distinct geological environments. Nothing contained herein constitutes investment, financial, or legal advice, nor should it be relied upon to evaluate the financial performance or market securities of Newmont Corporation.

REFERENCES

Bienstock, D and Zuckerberg, M, 2010. Solving LP relaxations of large scale precedence constrained problems, in *Proceedings of the 10th Conference on Integer Programming and Combinatorial Optimization*, pp 1–17 (Springer: Berlin).

Cacciuttolo, C, Atencio, E, Komarizadehasl, S and Lozano Galant, J A, 2025. Development of an advanced multi layer digital twin conceptual framework for underground mining, *Sensors*, 25(21):6650.

Muñoz, G, Espinoza, D, Goycoolea, M, Moreno, E, Queyranne, M and Rivera, O, 2018. A study of the Bienstock Zuckerberg algorithm: applications in mining and resource constrained project scheduling, *Computational Optimization and Applications*, 69:501–534.

Rossi, F, van Beek, P and Walsh, T, 2006. *Handbook of Constraint Programming* (Elsevier: Amsterdam).

Geometallurgy in mine planning

Quantifying geometallurgical uncertainty for mine planning using machine learning and conditional simulation

C Artica[1], M Noppé[2] and N Fox[3]

1. FAusIMM(CP), WH Bryan Mining Geology Research Centre, The University of Queensland, Brisbane Qld 4000. Email: cecilia.artica@outlook.com
2. JKTech Pty Ltd, Brisbane Qld 4000.
3. WH Bryan Mining Geology Research Centre, The University of Queensland, Brisbane Qld 4000.

ABSTRACT

Block models are essential tools in mining, underpinning resource estimation, mine design, and production forecasting. While conventional models focus on revenue-generating metals such as copper and gold, there is growing recognition of the importance of other variables that can affect profitability, such as deleterious elements, mineralogical content, and rock hardness. This study addresses the need for integrating these geometallurgical variables into three-dimensional block models, with a focus on grindability, represented by Bond Work Index (BWi), using data from Rio Tinto's Resolution Copper Project.

The optimised model achieved a coefficient of determination (R^2) of 0.716 and root mean squared error (RMSE) of 0.689 on all data, and maintained an R^2 of 0.484 and RMSE of 0.929 on the independent test set, demonstrating useful predictive capability while highlighting the expected reduction in performance under validation.

To address spatial variability and uncertainty, the study integrated conditional simulations of both lithological and continuous variables. Lithology was simulated using the Hierarchical Truncated Pluri-Gaussian method, while continuous variables were transformed using the Projection Pursuit Multivariate Transform and simulated via the Turning Bands method. The simulations preserved both the spatial and multivariate structure of the input data, enabling robust prediction of BWi values and variability across the deposit.

This research demonstrates the potential for combining machine learning and geostatistical simulation to develop predictive, spatially informed geometallurgical models. Future work will focus on building a full BWi block model, with applications in mine planning, processing optimisation, and uncertainty quantification.

INTRODUCTION

Geometallurgical modelling has become a critical tool in mining operations, linking geological variability to metallurgical performance and supporting decisions around mine planning, blending, and processing (Wright *et al,* 2013; Esper *et al,* 2013; Butler *et al,* 2013; Ehrig, 2013; Baumgartner, Gomez and Escobar, 2016; King and Macdonald, 2016; Leichliter, Jahoda and Montoya, 2013; Koch, 2019; Deutsch, 2015; Scott *et al,* 2023).

Examples include BHP's Iron Ore operations in Western Australia, where models help optimise blasting and blending to maximise crusher throughput (Wright *et al,* 2013), and Morro de Ouro in Brazil, where improved ore-type definition enhanced reconciliation and planning (Esper *et al,* 2013). At DeGrussa, geometallurgical models guide blend strategies to improve recovery and product quality (Butler *et al,* 2013), while at Olympic Dam, they inform comminution circuit design by linking mineralogical properties with metallurgical parameters (Ehrig, 2013). Similar approaches have improved flotation recoveries at Cerro Corona (Baumgartner, Gomez and Escobar, 2016), supported early-stage throughput and acid consumption estimates at Productora (King and Macdonald, 2016), and optimised metallurgical test work selection at La Colosa (Leichliter, Jahoda and Montoya, 2013).

Recent studies have integrated machine learning techniques to predict metallurgical properties such as ore hardness and blast furnace performance from geological features (Koch, 2019; Scott *et al,* 2023). However, these models are often deterministic and lack explicit spatial uncertainty representation. While geostatistical simulation techniques, including hierarchical truncated pluri-Gaussian (HTPG) and multivariate simulations, have been applied to capture spatial variability

(Deutsch, 2015), their integration with machine learning to predict processing-related variables remains limited.

Addressing this gap, this study proposes an integrated framework combining machine learning and conditional simulation to predict Bond Work Index (BWi), using Rio Tinto's Resolution Copper deposit as a case study. This approach preserves spatial and multivariate relationships and explicitly quantifies uncertainty, providing a more robust foundation for operational decision-making and mine planning.

The novelty of this work lies in integrating machine learning with conditional simulation of both categorical and continuous geological variables to generate spatially explicit predictions of Bond Work Index while preserving multivariate relationships and uncertainty.

METHODOLOGY

The modelling framework presented in this study was developed to predict BWi within a spatially explicit geometallurgical context. It combines machine learning (ML) with geostatistical simulation to generate a three-dimensional block model that reflects both spatial variability and uncertainty in key geological and metallurgical variables. The methodology is summarised here in further detail.

Data checks and preprocessing

The first stage involved rigorous data validation and exploratory analysis to assess the quality and representativeness of the available data sets. This included the evaluation of:

- Spatial distribution and clustering of sample locations.
- Differences in support and data resolution across variables.
- The extent and patterns of missing data.
- Identification and treatment of statistical outliers.
- Detection of non-linear relationships between geological inputs and metallurgical responses.

This stage was critical for understanding potential sources of bias and ensuring that subsequent modelling steps were based on robust and interpretable data.

Machine learning algorithm development

The second stage focused on developing a supervised ML model capable of predicting BWi from geological variables. Following the approach of Scott *et al* (2023) the key steps included:

- Feature selection, using Random Forest and Gradient Boosting algorithms to reduce dimensionality and retain only the most informative inputs.
- Data transformation, comparing normal score and robust scaler approaches to mitigate skewness and outlier influence.
- Algorithm selection and evaluation, benchmarking Linear Regression, Random Forest, and Gradient Boosting models.
- Hyperparameter tuning, using grid search to optimise model performance based on cross-validation metrics, including coefficient of determination (R^2), root mean squared error (RMSE) and residuals.

Categorical lithology data were also encoded and tested for inclusion, with one-hot encoding used to assess their predictive contribution. The final model was selected based on its generalisation performance on a withheld test set.

Simulation of categorical variables

To incorporate spatial variability and uncertainty in lithology, HTPG simulation was employed following the approach of Figueiredo Silva (2018). This entailed:

- Trend removal and variogram analysis of indicator variables representing lithological categories.
- Latent Gaussian transformation, ensuring compatibility with pluri-Gaussian frameworks.
- Definition of hierarchical truncation rules, based on geological contact relationships and depositional sequences.
- Simulation of multiple realisations, each honouring geological continuity, spatial statistics, and contact structures.
- Validation of simulation outputs to ensure geological realism and internal consistency.

HTPG simulation provided a means to generate equally probable lithological scenarios that reflect observed spatial patterns and categorical relationships.

Simulation of continuous variables

Continuous variables identified as important predictors (Al, Be, Cs, Mg, Mo, Zn) were simulated using a multivariate approach. The process included:

- Multiple imputation to handle missing data and ensure consistency across input grids.
- Multivariate transformation using the Projection Pursuit Multivariate Transform (PPMT) (Barnett, Manchuk and Deutsch, 2014) to preserve joint distributions.
- Variogram modelling for each transformed variable to capture spatial continuity.
- Turning Bands simulation to generate multiple realisations of each variable, conditioned to existing data.
- Back-transformation to recover the simulated variables in their original space.

The simulations were validated to ensure fidelity to both the statistical properties and spatial structures of the original data set.

RESULTS

Random Forest and Gradient Boosting algorithms were used in the feature selection process by ranking based on the number of splits attributed to each variable. The resulting importance rankings (Figure 1) consistently identified six key features: Mo, Cs, Mg, Be, Zn, and Al as the dominant contributors to BWi prediction. Both algorithms exhibited a clear step change in importance at the fifth to sixth feature, indicating that these top variables captured most of the predictive power. Features ranked below this threshold demonstrated a gradual decline in importance, suggesting diminishing returns from their inclusion. This analysis supports the decision to focus on the top six features for training the final predictive model, improving model efficiency while maintaining high predictive accuracy.

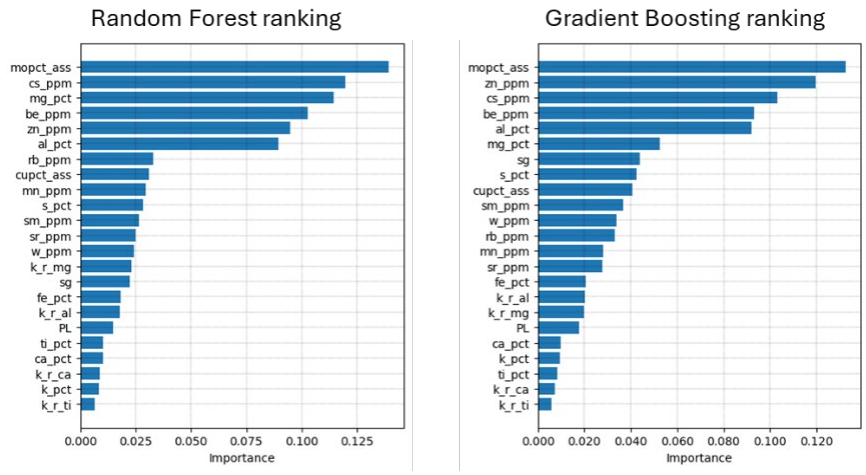


FIG 1 – Stage 2 feature ranking based on Random Forest and Gradient Boosting.

After comparing Linear Regression, Random Forest, and Gradient Boosting approaches, Gradient Boosting with a robust scaler transform was selected as the final model due to its superior ability to capture non-linear relationships and overall predictive accuracy. The model was built on the selected subset of six key numeric features (Al, Be, Cs, Mg, Mo, Zn), and further enhanced by including lithology as a categorical feature using one-hot encoding. When evaluated on all data, the final tuned model achieved an R^2 of 0.716, an RMSE of 0.689, and a Spearman correlation coefficient (ρs) of 0.817, showing a good fit with minimal bias. On the independent test set, it maintained an R^2 of 0.484, an RMSE of 0.929, and a ρ_s of 0.655, reflecting good generalisation. The inclusion of lithology was particularly effective in reducing overfitting and improving model robustness. These results underscore the importance of integrated feature selection, data transformation, and algorithm optimisation in developing a reliable framework for predicting BWi in complex geological settings.

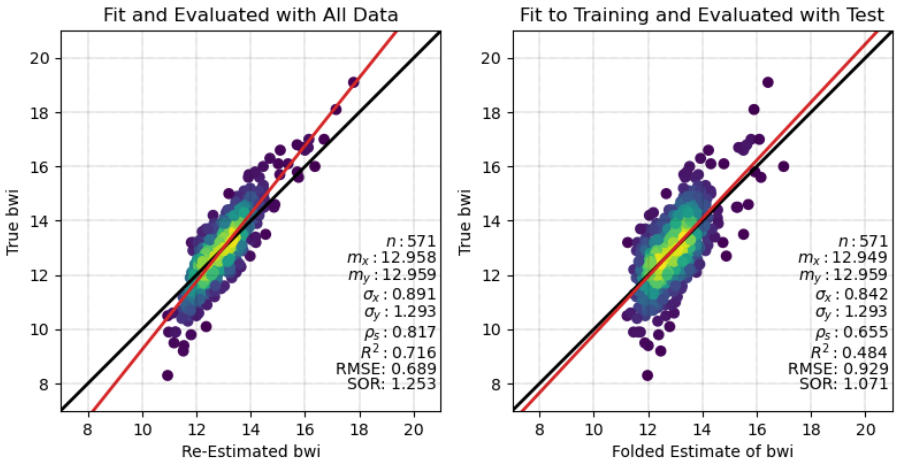


FIG 2 – Cross validation scatterplots – best scenario with lithology features added.

Figure 3 shows two realisations of simulated lithological categories generated using the HTPG simulation method. The cross-sections illustrate how the truncation rules and the numerically derived Gaussian variograms reproduce the expected hierarchical relationships and contacts defined in the conceptual model. The main lithological domains, including the diabase units (Diab1 to Diab3), quartzite subgroups (Qzite1 to Qzite3A), and Mesc, are clearly represented with coherent spatial continuity. The comparisons between realisations highlight the stochastic nature of the simulation while maintaining global proportions and major structural trends. Some local differences, particularly in the geometry of smaller units such as Qzite3A and Diab3, point to areas where limited data or edge effects affect geological realism.

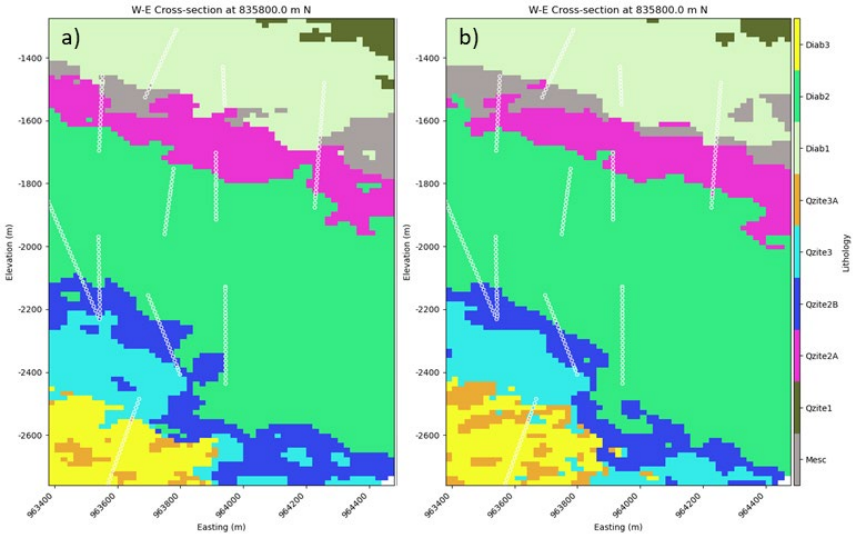


FIG 3 – West–east cross-sections at 835800mN of simulated geological categories: (a) Realisation 1; (b) Realisation 2.

Figure 4 illustrates the outcomes of the conditional simulation of continuous variables for Diabase 2, presented as plan views of the aluminium percentage (Al_pct). The upper left panel shows the original data distribution while the adjacent and lower panels display three simulated realisations, each representing an equally probable spatial distribution of Al_pct after imputation, multivariate transformation, trend removal, and conditional simulation. The realisations exhibit coherent spatial structures and honour both the global trends and local variability defined by the variogram models and multivariate relationships. The consistent reproduction of high and low value zones across realisations demonstrates the robustness of the simulation process, while the differences among realisations reflect the inherent uncertainty captured by the stochastic approach. These simulated continuous variable models provide critical inputs for the final BWi predictions, ensuring that both spatial variability and uncertainty are rigorously integrated into the geometallurgical framework.

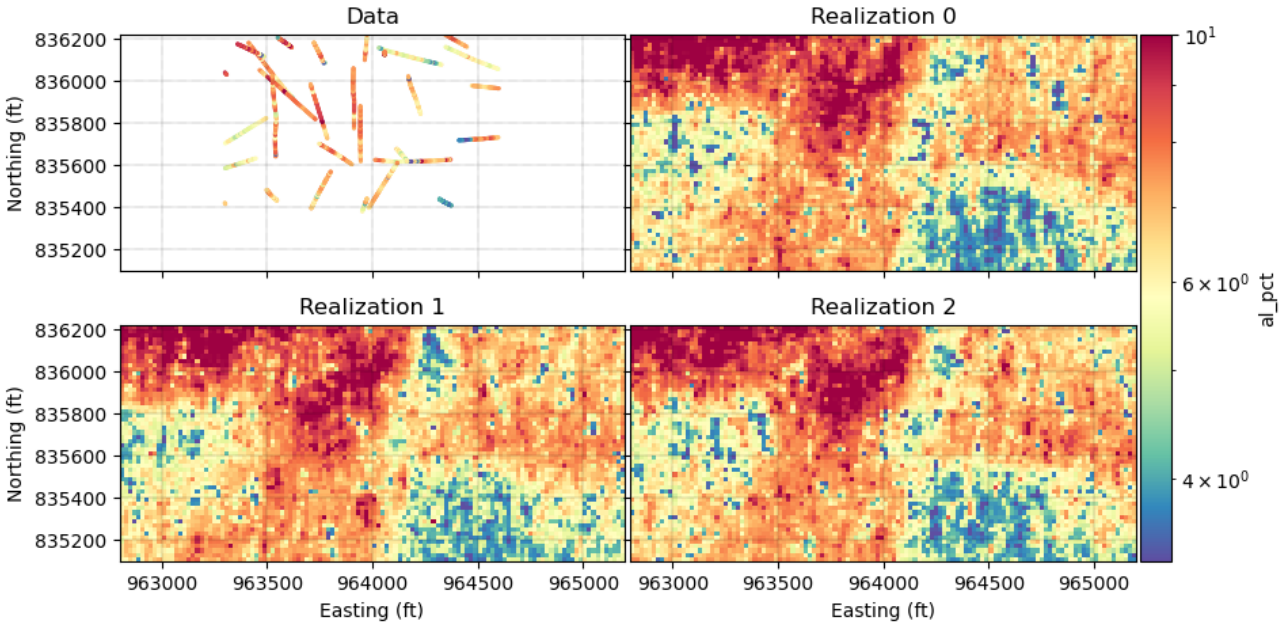


FIG 4 – Plan view showing the Al input data, and three example realisations for Diabase2.

DISCUSSION

Where a machine learning model is used to predict the BWi from simulated geological variables, the accuracy of the machine learning model relies heavily on the correct spatial representation of the intricate, non-linear dependencies among these variables (Dunham and Vann, 2007; Deutsch *et al*, 2016). Ensuring that these relationships were retained and accurately restored in the final three-dimensional product was a key focus of this research so that the machine learning model will be able to make accurate predictions from the simulated block model.

In future work, realisations from categorical and continuous simulations will be combined and input into the trained machine learning model to predict BWi across the deposit where:

- Each categorical simulation realisation (eg lithology) will be paired with a corresponding continuous simulation realisations (eg Al, Mo etc).
- For each pairing, the trained ML algorithm will be used to predict BWi at the point scale.
- This process will be repeated across all realisation pairs to generate an ensemble of equally probable BWi block models.
- The outputs can then be summarised at the Selective Mining Unit (SMU) scale to facilitate integration with mine planning systems.

This final step is expected to produce a robust geometallurgical block model that captures both predictive relationships and spatial uncertainty, enhancing its value for operational decision-making and scenario evaluation.

CONCLUSIONS

This research has successfully integrated machine learning and geostatistical conditional simulation aimed at enhancing the prediction of the BWi within a geometallurgical framework. As the industry transitions to exploiting lower-grade and more mineralogically complex orebodies, the demand for advanced geometallurgical models capable of capturing spatial variability, multivariate relationships, and uncertainty has become paramount. By combining simulations of both categorical and continuous variables with modern machine learning techniques, this study advances the field of geometallurgy, offering a comprehensive approach to optimise resource extraction and beneficiation. The novel framework developed here underscores the importance of integrating diverse data types and predictive modelling to improve decision-making and operational efficiency.

ACKNOWLEDGEMENTS

The authors thank Rio Tinto for permission to use data from the Resolution Copper Project and acknowledge the support of The University of Queensland.

REFERENCES

Barnett, R M, Manchuk, J G and Deutsch, C V, 2014. Projection pursuit multivariate transform, *Mathematical Geosciences*, 46:337–359.

Baumgartner, R, Gomez, P and Escobar, G, 2016. Comprehensive mineralogical characterisation at the Cerro Corona Cu-Au porphyry mine - the fundamental key for geometallurgical applications, in *Proceedings of the Third AusIMM International Geometallurgy Conference, Geomet 2016*, pp 221–230 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Butler, C, Dale, R, Robinson, S and Turner, A, 2013. DeGrussa Copper Mine - feasibility to first production: a case study in developing a working geometallurgy model, in *Proceedings of the Second AusIMM International Geometallurgy Conference, Geomet 2013,* pp 59–66 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Deutsch, J L, 2015. Multivariate Spatial Modelling of Metallurgical Rock Properties, PhD thesis (unpublished), University of Alberta.

Deutsch, J, Palmer, K, Deutsch, C, Szymanski, J and Etsell, T, 2016. Spatial modelling of geometallurgical properties: techniques and a case study, *Natural Resources Research*, 25:161–181.

Dunham, S and Vann, J, 2007. Geometallurgy, geostatistics and project value - does your block model tell you what you need to know?, in *Proceedings of the Project Evaluation Conference*, pp 1–8 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Ehrig, K, 2013. Geometallurgy - what you really need to know from exploration through to production?, in *Proceedings of the Metallurgical Plant Design and Operating Strategies, MetPlant 2013*, pp 28–33 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Esper, E, Rugolo, R, Moller, J and Pains, A, 2013. Morro do Ouro geological model with a metallurgical view, in *Proceedings of the Second AusIMM International Geometallurgy Conference, Geomet 2013,* pp 67–74 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Figueiredo Silva, D S, 2018. Enhanced Geological Modelling of Multiple Categorical Variables, PhD thesis, University of Alberta.

King, G and Macdonald, J, 2016. The business case for early-stage implementation of geometallurgy - an example from the Productora Cu-Mo-Au deposit, Chile, in *Proceedings of the Third AusIMM International Geometallurgy Conference, Geomet 2016,* pp 125–133 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Koch, P, 2019. Computational Methods and Strategies for Geometallurgy, PhD thesis, Luleå University of Technology.

Leichliter, S, Jahoda, R and Montoya, P, 2013. Using geometallurgical models to aid in metallurgical testwork for pre-feasibility projects, La Colosa, Colombia, in *Proceedings of the SME Annual Meeting*, pp 1–4 (Society for Mining, Metallurgy and Exploration: Denver).

Scott, T R, Gous, C, Muller, J and Deutsch, J L, 2023. Applying machine learning to predict the impact of changing iron ore properties on blast furnace performance, in Proceedings of the SAIMM Geometallurgy Conference 2023 (The Southern African Institute of Mining and Metallurgy: Johannesburg).

Wright, B, Carroll, T, Flanagan, M, Sua, T and Van Dijken, B, 2013. BHP Billiton West Australian iron ore geometallurgy - past, present and future, in *Proceedings of the Second AusIMM International Geometallurgy Conference, Geomet 2013,* pp 105–110 (The Australasian Institute of Mining and Metallurgy: Melbourne).

A geometallurgical framework using operational mining data for plant performance prediction and stochastic mine production scheduling

C Both[1]

1. Optimization Specialist, Newmont, Montreal QC H4M1L5, Canada.
   Email: christian.both@newmont.com

INTRODUCTION

Geometallurgy aims to characterise and model the relationships between spatially distributed rock properties within a mineral deposit and the resulting metallurgical behaviour throughout the mining value chain. These relationships are critical, as key processing performance indicators—such as plant throughput, metal recovery, and reagent consumption—are important drivers of revenue and operating costs in mining operations. Consequently, improving the prediction of these responses and their integration into mine planning and production scheduling is of high economic importance.

However, accurately predicting plant performance remains challenging. Processing responses are typically complex, non-linear functions of multiple rock characteristics and, importantly, of blended material streams rather than individual ore units. This blending introduces additional complexity, as several geometallurgical variables—such as rock hardness—are inherently non-additive. As a result, conventional approaches often produce biased predictions when applied at production scale, which limit their usefulness in decision-making and mine scheduling optimisation.

Traditional geometallurgical workflows typically rely on laboratory testing conducted on a limited number of drill core samples to assess comminution, recovery, and other processing responses. These tests, including Bond work index and SAG power index, are costly, time-intensive, and spatially sparse. While proxy models have been developed to estimate these properties from more readily available geological attributes, such models are typically calibrated on very small data sets and defined at drill core support. Consequently, they do not necessarily scale correctly to larger mining units or blended material streams.

Furthermore, conventional approaches are difficult to update as new information becomes available during operations. This creates a disconnect between geometallurgical modelling and the evolving mining system. However, modern mining operations continuously generate large volumes of operational data, providing an opportunity to develop adaptive, data-driven geometallurgical frameworks that better reflect actual processing behaviour.

GEOMETALLURGICAL PREDICTION MODEL USING OPERATIONAL MINING DATA

This work presents a novel geometallurgical framework that constructs predictive models of processing plant performance directly from operational mining data. The approach integrates multiple production data sets to establish empirical relationships between material extracted from the mineral reserve and observed plant performance (Figure 1).

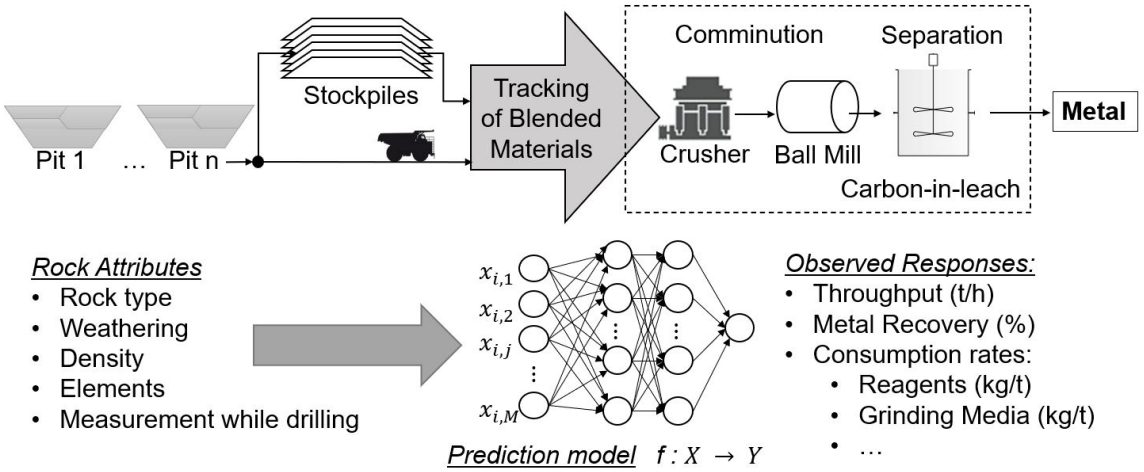


FIG 1 – Overview for constructing geometallurgical prediction models using production data.

A key advantage of this framework is the use of processing plant data as direct measurements of geometallurgical response variables at the appropriate support scale, labelled as observed responses in Figure 1. These data are nowadays abundantly available in operating mining complexes and reflect actual system performance under operating conditions.

Three additional data sources are integrated. First, measurement-while-drilling (MWD) data from blasthole drilling provide high-resolution information related to rock strength and hardness through penetration rates. This data set serves as a proxy for comminution-relevant properties and offers significantly greater spatial coverage compared to traditional metallurgical sampling. Second, geological models provide information on material types, densities, and metal grades, enabling characterisation of orebody variability.

Third, detailed material tracking is performed using truck cycle data from fleet management systems, allowing reconstruction of material flows from the mine to the processing plant at fine temporal resolution. This tracking accounts for multiple sources, stockpiling, and delays to estimate time-resolved blended feed characteristics. These are aligned with recorded plant operating data, allowing consistent pairing of predictor and response variables.

A central innovation of the framework is the treatment of non-additive geometallurgical variables. Instead of averaging hardness-related attributes, penetration rates are transformed into distributions or proportions of different hardness classes within each mining unit. These proportions better capture the disproportionate influence of harder material fractions on comminution performance and allow more accurate representation of blended behaviour.

Once blended rock characteristics and plant responses are aligned over consistent time intervals (eg daily or weekly), predictive models are constructed. Both multiple linear regression models and neural networks are employed to model relationships between geometallurgical attributes and processing responses such as throughput, reagent consumption, and grinding media usage.

CASE STUDIES AND INTEGRATION INTO STOCHASTIC MINE PRODUCTION SCHEDULING

Two case studies from an operating gold mining complex demonstrate the proposed framework. The complex consists of multiple open pit mines feeding a centralised processing facility including crushing, high-pressure grinding rolls (HPGR), ball mill, and carbon-in-leach circuits.

In the first case study, a ball mill throughput model is developed using blended hardness-related attributes derived from MWD data, combined with geological variables. Spatial uncertainty is accounted for via geostatistical simulation, and hardness proportions are computed within mining units to address non-additivity. After material tracking and temporal aggregation, regression models are calibrated using several months of plant data.

The resulting throughput model is integrated into a simultaneous stochastic mine production scheduling framework. Predicted throughput, based on blended material characteristics, replaces

conventional fixed tonnage constraints. Results show a significant reduction in the risk of not meeting production targets compared to conventional scheduling.

The second case study focuses on predicting reagent and consumable usage (grinding media, caustic soda, and acid). In addition to hardness-related attributes, blended grades and other geological variables are incorporated. Models are developed using weekly aggregated data and achieve strong predictive performance with correlation coefficients exceeding 0.7.

These models are also integrated into the stochastic scheduling framework, enabling processing costs to be modelled as functions of blended material characteristics rather than fixed unit costs. The resulting schedules yield different material selection strategies and improved economic outcomes. In the presented case, expected profit increases by approximately 3 per cent over a one-year planning horizon, while also improving reliability in meeting production targets.

CONCLUSION

This work demonstrates that operational mining data can be effectively leveraged to construct high-resolution, adaptive geometallurgical prediction models for key processing performance indicators. By integrating MWD data, geological information, and material tracking with plant operating data, the proposed framework provides a practical and scalable alternative to traditional laboratory-based geometallurgical programs.

Finally, the integration of these data-driven prediction models into a simultaneous stochastic mine production scheduling framework enables improved decision-making regarding material extraction, blending, and processing. The results show increased expected profitability, reduced production risk, and improved alignment with plant performance targets.

Disclaimer: The views and conclusions expressed in this paper are those of the author(s) and do not necessarily reflect the views, positions, or official policies of Newmont Corporation or its affiliates. Any forward looking statements are based on current assumptions and are subject to risks and uncertainties that may cause actual results to differ in a material way.

Mine scheduling, algorithms and pit optimisation

Multi-objective genetic algorithm for strategic open pit mine planning – generating pareto-optimal schedules for value-feasibility trade-offs

J Andre[1], E Fosdike[2] and M Roberts[3]

1. Regional Customer Success Lead – Mine Engineering, Maptek, Blainville QC J7C 5Z4, Canada. Email: jeremy.andre@maptek.com
2. Software Engineer, Maptek, Adelaide SA 5065, Australia. Email: edward.fosdike@maptek.com.au
3. Global Strategy Manager, Maptek, Perth WA 6163, Australia. Email: mark.roberts@maptek.com.au

INTRODUCTION

Strategic open pit mine planning requires balancing multiple competing objectives, including project value, equipment utilisation, environmental and process-related objectives. While optimisation techniques traditionally focus on maximising a single economic objective, such as net present value (NPV), planning decisions are driven by trade-offs between value creation and operational feasibility. Strategic scheduling problems are non-polynomial time hard (Caccetta and Hill, 2003), and solutions quality is often limited by computational complexity in mathematical programming approaches. Population-based metaheuristics address this challenge by trading optimality guarantees for computational efficiency, using stochastic search mechanisms, and provide best-effort solutions when an optimum is computationally impractical (Deb and Jain, 2014). However, single-objective metaheuristics can lead to restricted convergence and impractical results. They optimise a single objective while prioritising constraint satisfaction, limiting solutions' convergence towards higher-quality solutions. Also, single-objective optimisation requires to embed subjective trade-offs before the optimisation, rather than exposing competing objectives. This case study a multi-objective genetic algorithm to generate Pareto-optimal life-of-mine schedules under a common strategic open pit model. The Pareto-front identifies non-dominated Pareto-optimal solutions (Deb and Myburgh, 2016), where improving one objective deteriorates at least one other, enabling planners to select value-feasibility trade-offs aligned with site priorities.

MULTI-OBJECTIVE GENETIC ALGORITHM FOR STRATEGIC PLANNING

Based on a block model, each candidate schedule, or individual, represents a complete life-of-mine mining sequence, including period allocation, material routing, stockpiling strategy and cut-off grades defined for each pit-stage-bench-lithology combination. Derived from a single-objective framework (Myburgh and Deb, 2017), this method optimises selected economic objectives, such as NPV, cut-off grade variations and wasted metal, alongside operational objectives for the total mining capacity, blending and minimum accumulation. Minimum accumulation represents a minimum quantity, such as tonnes or ounces, that must be scheduled. For any schedule $s \in S$, where S is the set of schedules satisfying all hard constraints, the objective vector F is:

$$\min_{s \to S} F(s) = [\frac{1}{f_{NPV}(s)}, f_{wmet}(s), f_{COG}(s), f_{blend}(s), f_{acc}(s), f_{min}(s)]$$

Where:

- $f_{NPV}(s)$ is the schedule NPV
- $f_{wmet}(s)$ is a cumulative wasted metal measure
- $f_{COG}(s)$ is a cumulative cut-off-grade variation between scheduling periods
- $f_{blend}(s)$ is a cumulative grade blending violation measure
- $f_{acc}(s)$ is a cumulative violation measure about total tonnage or volume moved
- $f_{min}(s)$ is a cumulative minimum-accumulation violation measure

The multi-objective genetic algorithm optimises cut-off grades, the mining order for each bench and stage, as well as the stockpile strategy and capital allocation. Each operational objective is assessed

against user-defined upper and lower bounds for each period. Compliance means meeting the target in all schedule periods, while any deviation is treated as a violation to be minimised. In this framework, thousands of candidate schedules are populated over many generations and each evaluated for fitness. Candidate schedules are ranked using all selected objectives, with non-dominated solutions preferentially retained to preserve solution quality and diversity across generations (Deb and Jain, 2014). Over successive generations, crossover and mutation operators generate new feasible schedule alternatives, and the population evolves toward increasingly favourable value-feasibility trade-offs. Metaheuristics are designed to explore large and complex search spaces for potential solutions, which make them well-suited for multi-objective strategic planning problems.

CASE STUDY AND SUPPORT FOR DECISION-MAKING

A Pareto-optimal schedule is one in which no objective can be improved without degrading another. Unlike single-objective optimisation, multi-objective optimisation produces a set of trade-off solutions rather than a single optimum. This approach efficiently generates diverse Pareto-optimal schedules for large, complex planning problems (Figure 1).

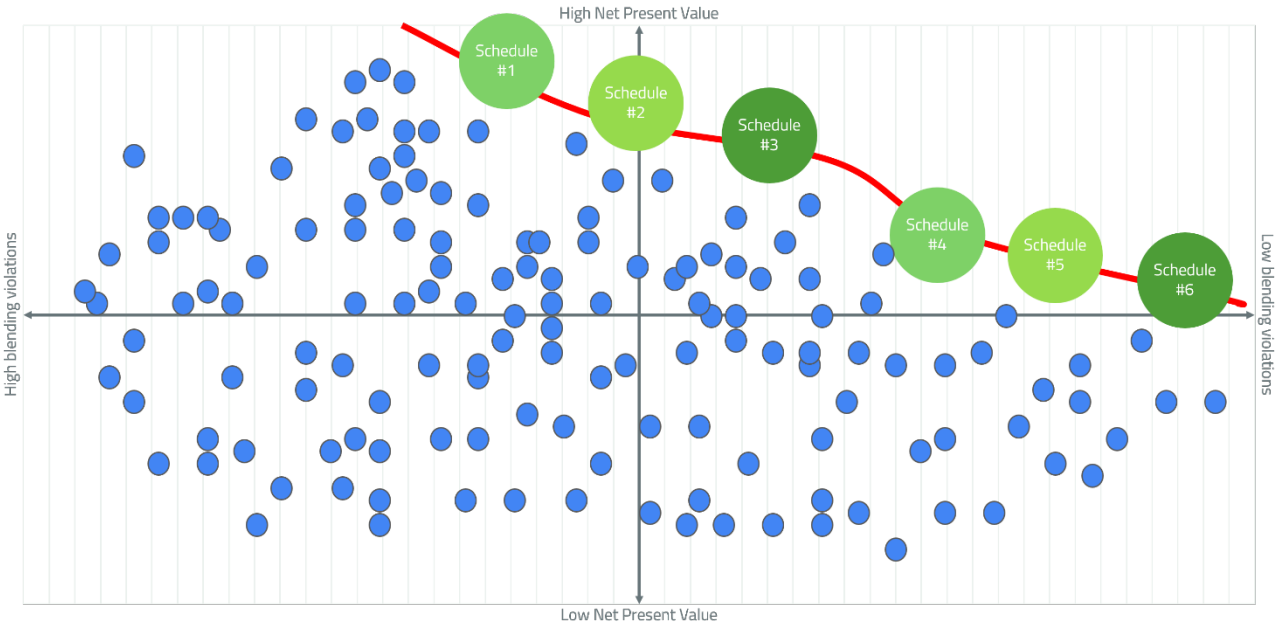


FIG 1 – Pareto-front representation (red line) of the trade-off between NPV and blending objectives.

The case study is based on a gold mine with two pits, eight phases and three lithologies, resulting in 702 stage-bench-lithology combinations. Results were generated with a population of 200 individuals, up to 100 generations and a stopping criterion of ten generations without improvement (Table 1). Across ten runs, the average runtime was 29.5 mins, compared with 13 mins for the equivalent single-objective model.

TABLE 1

Comparative schedule outcomes illustrating trade-offs between economic performance and operational objectives.

Solution	Period number	NPV (in US$ million)	Minimum accumulation violations period number	Earliest period violating minimum accumulation	Blending violation period number	Cumulative cut-off grade variation
A	8	973.9	3	6	3	0.03
B	9	967.0	3	7	3	4.80
C	9	959.6	3	7	3	0.81
D	9	955.3	3	7	5	3.36
E	8	936.3	2	7	6	12.88
F	9	887.6	4	2	7	11.95

In Table 1, the Pareto-optimal schedules trade off NPV, blending, minimum accumulation compliance and cumulative cut-off grade variation. Schedule A delivers the highest NPV, but this outcome is accompanied by earlier minimum accumulation target violations and limited blending compliance. While economically attractive, this schedule may require greater operational flexibility to manage mining and processing performance. At the opposite end of the Pareto frontier, schedules E and F achieve the highest number of blend-compliant periods and improved minimum accumulation performance, although at the cost of a substantial reduction in NPV and significantly greater cut-off grade variability. Intermediate solutions illustrate the range of available compromises. Relative to schedule A, schedule C delays the first minimum accumulation violation and improves operational performance while sacrificing USD14.3 million in NPV. Schedule D further improves blending compliance, achieving five blend-compliant periods while maintaining an NPV above USD955 million. These results demonstrate no schedule simultaneously optimises all objectives and highlights the value of Pareto-front analysis. Rather than focusing on identifying a single optimal solution, planners can select a value-feasibility trade-off aligned with site-specific priorities.

CONCLUSIONS

Strategic planning inherently involves competing objectives and single-objective optimisation techniques tend to obscure considered trade-offs. Multi-objective genetic algorithms provide a practical mechanism for simultaneously evaluating economic and operational objectives and generating diverse Pareto-optimal schedules. The case study demonstrates substantial improvements in blending and minimum accumulation compliance can be achieved with relatively modest reductions in NPV, highlighting the value of Pareto-based optimisation to support more informed strategic planning decisions. The multi-objective optimisation framework implemented in Maptek Evolution Strategy shifts strategic scheduling from a prescriptive search for a single optimal solution towards a trade-off-based decision support process. By generating diverse Pareto-optimal schedules, the approach enables mine planners to align planning outcomes with site-specific operational and business objectives.

REFERENCES

Caccetta, L and Hill, S P, 2003. An application of branch and cut to open pit mine scheduling, *Journal of Global Optimization*, 27:349–365.

Deb, K and Jain, H, 2014. An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, Part I: Solving problems with box constraints, *IEEE Transactions on Evolutionary Computation*, 18(4):577–601.

Deb, K and Myburgh, C, 2016. Breaking the billion-variable barrier in real-world optimization using a customized evolutionary algorithm, in *Proceedings of the Genetic and Evolutionary Computation Conference 2016,* pp 653–660 (Association for Computing Machinery: New York).

Myburgh, C and Deb, K, 2017. Use of derived heuristics in improved performance of evolutionary optimization: an application to gold processing plant, in *Proceedings of the 2017 IEEE Congress on Evolutionary Computation*, pp 830–837 (IEEE: Piscataway).

Pit optimisation for large models in the presence of geologic and economic uncertainty

M Deutsch[1], J Deutsch[2] and C Deutsch[3]

1. Director of Software Development, GeologicAI, Denver CO 80126, USA. Email: mdeutsch@geologicai.com
2. President Resource Modeling Solutions, GeologicAI, Calgary AB T2A 6R1, Canada. Email: jdeutsch@geologicai.com
3. Advisor, GeologicAI, Calgary AB T2A 6R1, Canada. Email: cdeutsch@geologicai.com

ABSTRACT

Making mine planning decisions in the presence of geologic and economic uncertainty requires minimising downside risk and maximising expected performance. This multi-objective optimisation problem has historically proven to be computationally challenging for realistic models with tens of millions of blocks and hundreds of realisations. This paper introduces a methodology for generating solutions that is applicable to large models without sacrificing quality or requiring too much time. Lagrangian relaxation is applied to sample the efficient frontier and construct a set of Pareto optimal solutions, enabling decision-makers to balance risk and expected return. This framework is applicable to large pit optimisation problems that include tonnage and stripping ratio constraints.

INTRODUCTION

The planning of open pit mining operations must contend with geologic and economic uncertainty. The small fraction of the deposit that is directly sampled by drilling is not sufficient to provide an exact model of the subsurface and the chaotic nature of global markets and economies mean that the parameters and inputs to the mine planning process are all uncertain. Geostatistical methods for handling uncertainty typically involve generating multiple realisations of the subsurface that must be considered together as a whole for downstream planning activities. However, for practical reasons many mine planning algorithms are designed to operate on a single model at a time with a single set of global parameters.

There have been several approaches developed to consider uncertainty within mine planning. Some work to include uncertainty considerations directly into the objective function or the constraints (Dimitrakopoulos and Ramazan, 2008). Others try to apply conventional mine planning algorithms to a representative subset of geostatistical/economic realisations. In general, the inclusion of uncertainty complicates the problems substantially and they struggle to contend with large realistic models which limits their real-world applicability.

An ideal solution approach must scale to an adequate number of geostatistical/economic realisations, over 100, and must work well within a reasonable modelling resolution. On the order of 10 million active blocks (not counting air or extra blocks around the periphery of the deposit) must be practical. It must consider all realisations together to generate a small set of designs that actively balance performing well across all realisations and avoiding downside risks.

METHODOLOGY

The conventional ultimate pit optimisation problem identifies a subset of material which together maximises the overall extracted value subject to precedence constraints (Lerchs and Grossmann, 1965). To mine any deeper would require the removal of so much waste as to have a net negative impact on the total value. This problem is simplistic but is still of use in some parts of the mine planning process. A simple formulation of this problem with uncertainty is:

$$\begin{array}{lll} maximize & \sum_{l\in L}\left(\sum_{b\in B} v_{b,l} X_b\right) & \\ s.t & X_b - X_{\hat{b}} \le 0 & \forall b \in B, \hat{b} \in \hat{B}_b \\ & 0 \le X_b \le 1 & \forall b \in B \end{array} \quad (1)$$

Where:

$b \in B$	is the set of all blocks
$l \in L$	is the set of all geostatistical / economic realizations using different underlying geologic models and different economic parameters
$\hat{b} \in \hat{B}_b$	is the set of antecedent blocks that must be mined if block b is to be mined
$v_{b,l}$	is the economic block value of block b for realization l
X_b	is the proportion of block b which is mined in the ultimate pit

The objective in Equation 1 is then to maximise the performance of a single design across all realisations on average. Because X_b is independent of l solving a single ultimate pit problem using the block average (the e-type model) will maximise expected performance across all realisations.

The next step is to add a consideration for minimising risk across realisations. It may feel natural to minimise a symmetric measure of risk, such as the variance or standard deviation. However, this is not reflective of mining goals. Minimising the standard deviation will penalise a plan that overachieves the mean on some realisations. Instead, an asymmetric measure of variance such as a lower partial moment is preferred (Dimitrakopoulos, Martinez and Ramazan, 2007; Fishburn, 1977).

The semi-deviation is an appropriate lower partial moment which squares the deviations below some target to more heavily penalise large underperformances. In this context it rewards a plan that is robust across realisations and would not dramatically fail on any given set of inputs. The target T may be a specific number known to the operator or it may be the mean. The semi-deviation objective is given in Equation 2:

$$minimize \quad \sqrt{\frac{1}{L}\sum_{l \in L}\left(\max\{0, T - \sum_{b \in B} v_{b,l} X_b\}\right)^2} \tag{2}$$

Incorporating this second objective into the optimisation process is not particularly straightforward, especially if retaining the network-flow based reformulation is desired. Instead of trying to balance the two objectives directly it seems prudent to generate a set of Pareto optimal solutions. These solutions are not dominated by another solution in either maximising expected value or minimising downside risk and form an efficient frontier as demonstrated in Figure 1.

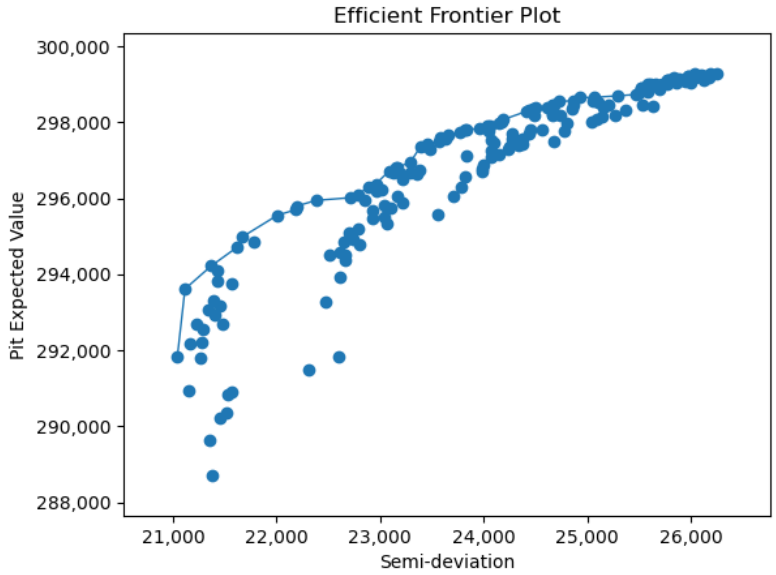


FIG 1 – Example efficient frontier showing several possible ultimate pits measured with both their expected value across all realisations and their semi-deviation.

Each point is an ultimate pit that has been evaluated against all realisations using the objective function in Equations 1 and 2 to obtain the pit expected value and semi-deviation values respectively.

The solutions along the top left of the cloud of solutions (connected by the line) are Pareto optimal and form the efficient frontier. The operator should likely choose one of these solutions considering the cost for a particular level of risk tolerance. To generate these solutions an iterative approach based on Lagrangian relaxation is applied. Similar approaches to slightly different problems have been applied in Dagdelen and Johnson (1986), Asad and Dimitrakopoulos (2013), Dimitrakopoulos, Farelly and Godoy (2002) and others.

To approximate the downside risk measure, one can add a constraint to ensure that all realisations achieve a combined value higher than some realisation specific target, relax those constraints into the objective, and solve the resulting network flow model repeatedly with appropriate Lagrange multipliers.

This approximation of the downside risk measure performs well in practice and is computationally efficient. Given that a suitably sophisticated implementation of pseudoflow can solve a 10 million block ultimate pit problem in a second, many thousands of candidate Lagrange multipliers can be evaluated in a reasonable amount of time (Hochbaum, 2008; Gallo, Grigoriadis and Tarjan, 1989). In practice this approach generates solutions with only minor reductions in value and significant reductions in downside risk exposure. The biggest challenge with this method is the same as with conventional nested pit analysis which uses a similar scheme. The gap problem may prevent us from identifying exactly the pit that we want as a small change in the Lagrange multipliers might lead to a big change in the output pit (Meagher, Dimitrakopoulos and Avis, 2014).

CONCLUSIONS

This extended abstract presents a computationally efficient approach for ultimate pit optimisation under geological and economic uncertainty. Lagrangian relaxation is used to generate a set of Pareto optimal pit designs that balance expected value against downside risk while retaining the efficiency of network-flow optimisation. The approach is applicable to large block models with many realisations and can incorporate practical constraints such as tonnage and stripping ratio limits. Although the gap problem may prevent the generation of some intermediate pit sizes, the method provides decision-makers with a practical range of alternatives for evaluating the trade-off between expected performance and risk.

REFERENCES

Asad, M W A and Dimitrakopoulos, R, 2013. Implementing a parametric maximum flow algorithm for optimal open pit mine design under uncertain supply and demand, *Journal of the Operational Research Society*, 64(2):185–197.

Dagdelen, K and Johnson, T B, 1986. Optimum open pit mine production scheduling by Lagrangian parameterization, in *Proceedings of the 19th International Symposium on Application of Computers and Operations Research in the Mineral Industry* (ed: R V Ramani), pp 127–142 (Society of Mining Engineers: Littleton).

Dimitrakopoulos, R and Ramazan, S, 2008. Stochastic integer programming for optimising long-term production schedules of open pit mines: methods, application and value of stochastic solutions, *Transactions of the Institutions of Mining and Metallurgy, Mining Technology*, 117(4):155–160.

Dimitrakopoulos, R, Farrelly, C T and Godoy, M, 2002. Moving forward from traditional optimization: grade uncertainty and risk effects in open-pit design, *Transactions of the Institution of Mining and Metallurgy*, *Mining Technology*, 111(1):A82–A88.

Dimitrakopoulos, R, Martinez, L and Ramazan, S, 2007. A maximum upside/minimum downside approach to the traditional optimization of open pit mine design, *Journal of Mining Science*, 43(1):73–82.

Fishburn, P C, 1977. Mean-risk analysis with risk associated with below-target returns, *American Economic Review*, 67(2):116–126.

Gallo, G, Grigoriadis, M D and Tarjan, R E, 1989. A fast parametric maximum flow algorithm and applications, *SIAM Journal on Computing*, 18(1):30–55.

Hochbaum, D S, 2008. The pseudoflow algorithm: a new algorithm for the maximum-flow problem, *Operations Research*, 56(4):992–1009.

Lerchs, H and Grossmann, I F, 1965. Optimum design of open-pit mines, *Transactions of the Canadian Institute of Mining and Metallurgy*, 68:17–24.

Meagher, C, Dimitrakopoulos, R and Avis, D, 2014. Optimized open pit mine design, pushbacks and the gap problem—a review, *Journal of Mining Science*, 50(3):508–526.

Strategic conveyor guided mine sequencing for cost stability and value optimisation in nickel open pit operations

A R Wilogo[1]

1. Mining Engineering Department, PT Roda Jaya Sakti, East Jakarta, Indonesia.
   Email: andikaramdhani11@gmail.com

ABSTRACT

Conveyor based material handling in open pit mining is commonly evaluated as an alternative haulage system after the broader mine development strategy has been established. However, fixed conveyor infrastructure introduces spatial and sequencing constraints that can influence where, when and how material is accessed. In heterogeneous nickel laterite deposits, this challenge is further complicated by the need to maintain selective mining and material routing across variable limonite, transition, saprolite and waste domains.

This paper proposes an Adaptive Hybrid Continuous Mining Framework (AHCMF) integrating mining suitability assessment, conveyor corridor definition, hybrid excavation and artificial intelligence (AI) assisted selective material classification into strategic mine sequencing. Mining domains are classified according to geological continuity, ore thickness, grade variability, material characteristics and operating constraints to identify areas suited to continuous, hybrid or conventional extraction. Conveyor corridors are then evaluated together with long-term extraction sequences rather than treated as downstream material handling infrastructure.

A conceptual AI assisted classification box is positioned within the conveyor material flow to support selective routing. Geological and grade control information is combined with potential real time inputs including elemental measurements, moisture, belt mass flow and machine vision to support routing into direct process feed, blending material, marginal material or waste. Geological personnel retain final control of classification decisions.

A conceptual large scale nickel laterite operation representative of conditions encountered in Central Sulawesi, Indonesia, with an assumed mining rate of 10 million bcm/a, is used to illustrate the framework. The assessment does not represent a validated equipment design or feasibility level economic study. Instead, it identifies interactions between geological suitability, infrastructure utilisation, sequence flexibility, material selectivity and long-term cost exposure.

The framework indicates that the strategic value of conveyor guided mining depends not on conveyor transport alone, but on the integration of mining domain selection, infrastructure commitment, extraction sequence and selective material routing.

INTRODUCTION

Nickel laterite operations require mine plans capable of responding to geological variability while maintaining consistent material supply to downstream processing. Lateritic profiles typically comprise limonite, transition and saprolite horizons developed above ultramafic bedrock. These horizons can vary laterally and vertically in thickness, grade, moisture, mineralogy and physical characteristics. Such variability creates a continuing requirement for selective extraction, grade control and destination management.

Conventional truck excavator mining is well suited to this operating environment because mobile equipment provides substantial tactical flexibility, consistent with established open pit mining and mine planning principles (Hartman and Mutmansky, 2002; Hustrulid, Kuchta and Martin, 2013). Loading units can be reassigned between mining areas, trucks can be redirected to alternative destinations and haul routes can be modified as pit geometry develops. This flexibility has considerable strategic value where geological interpretation, feed requirements or operational priorities change over time.

The same flexibility, however, creates continued dependence on mobile haulage. As an operation expands, increasing haul distance, changing road geometry, fleet availability, fuel consumption, tyre consumption and road maintenance can progressively affect cost and schedule performance.

Material transport therefore becomes more than an operational fleet selection issue and can influence long-term mine value.

Continuous and semi continuous material handling systems provide an alternative operating architecture. In-pit crushing and conveying systems have been studied as alternatives to conventional truck haulage because fixed infrastructure may reduce the requirement for extensive truck fleets under suitable mining conditions. However, these systems require planning approaches that explicitly consider infrastructure location, relocation and utilisation throughout the mine life. Nehring *et al* (2018) show that strategic mine planning approaches materially affect the evaluation of in-pit crush and convey (IPCC) systems, while Kamrani, Pourrahimian and Askari Nasab (2025) integrate road and conveyor networks directly into long-term scheduling.

Nickel laterite introduces an additional planning challenge because a fully continuous mining system may not be appropriate across the complete deposit. Irregular ore contacts, internal waste, variable material competency, high moisture and selective grade requirements may reduce the effectiveness of continuous excavation in some areas. The strategic question is therefore not whether conveyors should replace trucks across an operation, but where fixed infrastructure and continuous mining provide sufficient value and where mobile flexibility should remain.

A second challenge relates to selectivity after material enters a continuous transport stream. Conveyor systems can provide stable material movement but may reduce the discrete destination control characteristics associated with truck haulage. Maintaining reliable classification therefore becomes important where different laterite domains must be directed to separate processing, blending, stockpiling or waste destinations.

Sensor based sorting has demonstrated that real time material information can support waste rejection and diversion into alternative process routes, although effectiveness depends on ore characteristics, particle presentation, sensor technology and process configuration (Robben and Wotruba, 2019). This selective routing concept is also consistent with the broader role of material characterisation and separation in mineral processing (Wills and Finch, 2015).

This paper proposes an Adaptive Hybrid Continuous Mining Framework that combines three strategic elements: mining suitability classification, conveyor guided mine sequencing and artificial intelligence (AI) assisted selective material routing. The proposed AI system is not intended to replace geological interpretation or grade control. Its role is to provide an additional decision support layer capable of integrating geological information with real time material measurements within a continuous transport environment.

The purpose of the paper is to establish a conceptual planning architecture for determining where continuous mining may be applicable, how conveyor corridors interact with the long-term mine sequence and how selective material routing may be retained. The paper does not seek to demonstrate a completed engineering design or predetermined economic advantage.

Conceptual case study basis

A conceptual large scale nickel laterite operation in Central Sulawesi, Indonesia, is used to illustrate the proposed framework. The case study does not represent a specific operating mine, Mineral Resource, Ore Reserve, pre-feasibility study or feasibility study.

The operation is assumed to contain lateritic mineralisation comprising limonite, transition and saprolite horizons. Geological continuity and material characteristics are assumed to vary across the deposit, resulting in areas with different requirements for selectivity and different potential suitability for continuous excavation.

A nominal mining rate of 10 million bcm/a is adopted to provide scale for the strategic assessment. The operation is assumed to be sufficiently large and long lived for alternative material handling infrastructure to warrant strategic consideration.

The conceptual case study basis is summarised in Table 1.

TABLE 1

Conceptual case study basis.

Parameter	Conceptual basis
Location	Central Sulawesi, Indonesia
Deposit type	Nickel laterite
Nominal mining rate	10 million bcm/a
Primary material domains	Limonite, transition and saprolite
Reference mining system	Truck excavator mining
Alternative mining architecture	Hybrid continuous and conventional
Primary material transport	Conveyor guided
Selective routing concept	AI assisted classification with geological oversight
Assessment level	Conceptual strategic study

No site specific equipment quotation, detailed resource model, hydrological model, geotechnical design or feasibility level economic estimate is applied. Numerical assumptions are used only where necessary to establish the scale and logic of the framework.

STRATEGIC PLANNING PROBLEM

Conventional mobile haulage planning

In a conventional truck excavator mining operation, pit development and extraction sequences can be modified with relatively high spatial flexibility. Haulage routes change progressively as working areas advance, and trucks can be redirected between loading areas and destinations.

This flexibility is particularly valuable in heterogeneous ore deposits because the operation can react to changing geological information and short-term feed requirements.

The strategic disadvantage is that the material handling system remains strongly exposed to haulage distance and mobile fleet requirements. As pits expand or deepen, the number of truck hours required to move the same volume of material may increase unless additional infrastructure or alternative haulage strategies are introduced.

Conveyor guided planning

Conveyor infrastructure creates a different planning relationship. Alignment, gradient, transfer locations, drive stations, working space, drainage and relocation requirements must be considered before infrastructure is installed.

A conveyor corridor that is poorly aligned with future mining areas may experience low utilisation or require premature relocation. Conversely, a mining sequence developed without considering fixed infrastructure may create technically feasible production schedules that are strategically incompatible with conveyor development.

The resulting relationship can be summarised as:

- The mine sequence determines conveyor utilisation, while the conveyor corridor influences the practical mine sequence.
- This interdependence forms the principal strategic planning problem addressed by Adaptive Hybrid Continuous Mining Framework (AHCMF). Figure 1 compares the conventional mobile haulage planning logic with the conveyor guided planning logic adopted in AHCMF.

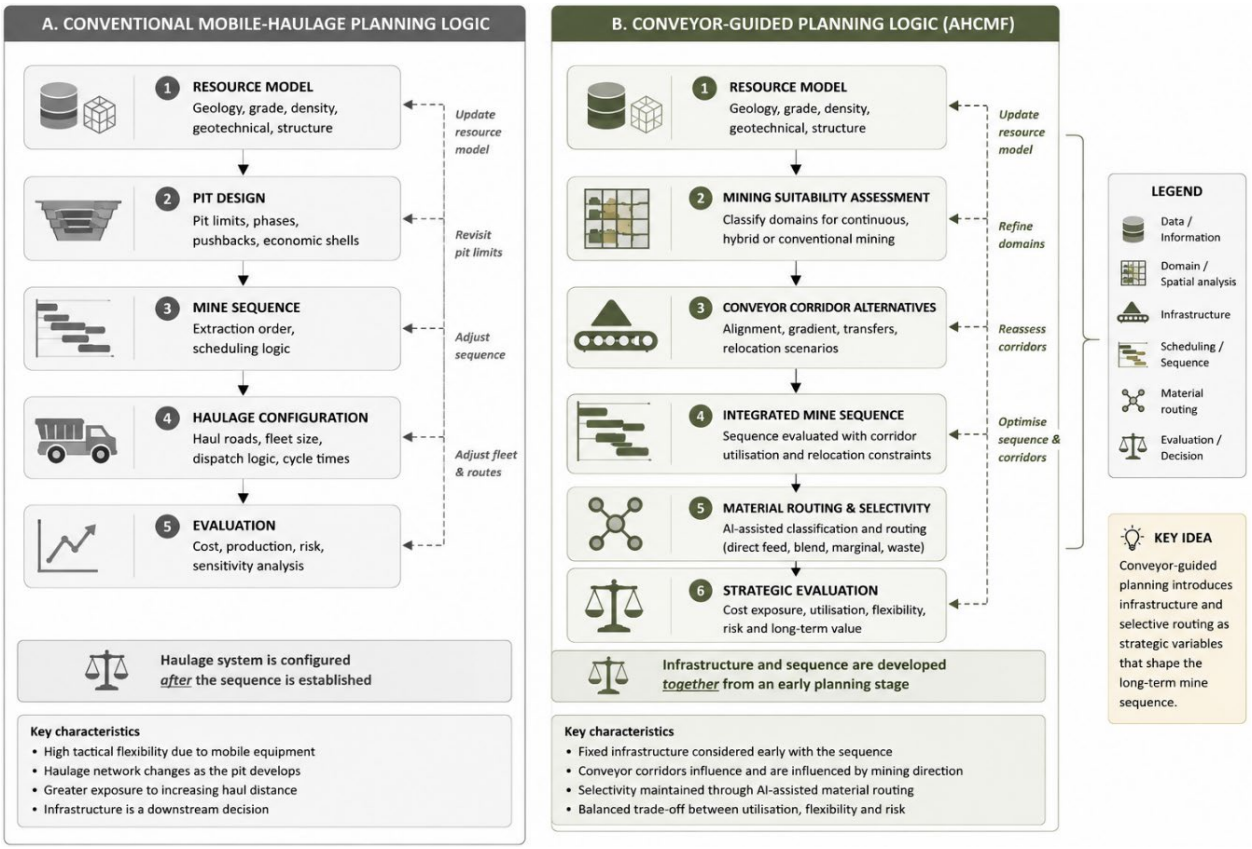


FIG 1 – Conventional and conveyor guided strategic planning logic.

The conventional planning pathway develops the resource model, pit design and mine sequence before finalising haulage configuration. AHCMF introduces mining suitability and conveyor corridor alternatives earlier in the process so that mining method, infrastructure and extraction sequence are assessed together.

MINING SUITABILITY ASSESSMENT

Continuous mining should not be applied uniformly across a heterogeneous nickel laterite deposit. AHCMF therefore begins with a mining suitability assessment that separates the deposit into areas where continuous, hybrid or conventional mining may be strategically appropriate.

The assessment includes geological, material, geotechnical, infrastructure and strategic factors.

Geological criteria include ore thickness, lateral continuity, internal waste, contact complexity and grade variability. Material criteria include diggability, moisture, boulder occurrence and competency. Infrastructure criteria include ground stability, available working width, drainage interaction and conveyor accessibility. Strategic criteria include expected mining duration, material volume, sequence position and potential infrastructure utilisation.

Three conceptual mining classes are defined.

Class A – continuous mining domain

Class A represents relatively continuous areas where geometry and material characteristics provide the strongest potential for sustained continuous excavation and conveyor utilisation.

These domains generally require lower levels of local selectivity and provide sufficient mining duration to support infrastructure utilisation.

Class B – hybrid mining domain

Class B contains material that may support continuous excavation but includes local variability requiring selective support from conventional equipment.

Potential conditions include irregular contacts, local waste lenses, isolated boulders, variable material characteristics or temporary access constraints.

Class B represents the primary operating environment for AHCMF because it combines the continuity advantages of conveyor based systems with the selectivity of mobile equipment.

Class C – conventional mining domain

Class C contains thin, highly irregular, discontinuous or strongly selective areas where the operational value of mobile mining equipment exceeds the potential benefits of continuous excavation. Hydraulic excavators and truck haulage remain the principal mining method in these areas. The conceptual characteristics of Classes A, B and C are summarised in Table 2.

TABLE 2

Conceptual mining suitability classification.

Characteristic	Class A	Class B	Class C
Geological continuity	High	Moderate	Low
Ore thickness consistency	Higher	Variable	Irregular
Selectivity requirement	Low–moderate	Moderate	High
Continuous excavation potential	High	Conditional	Low
Conveyor compatibility	High	Conditional	Limited
Conventional equipment role	Supporting	Significant	Primary

The classification is not intended to estimate productivity directly. Its purpose is to provide a strategic screening process that links orebody characteristics with mining system selection before long-term sequencing is finalised. Figure 2 illustrates how the assessment inputs are translated into mining suitability classes and corresponding mining-system selections.

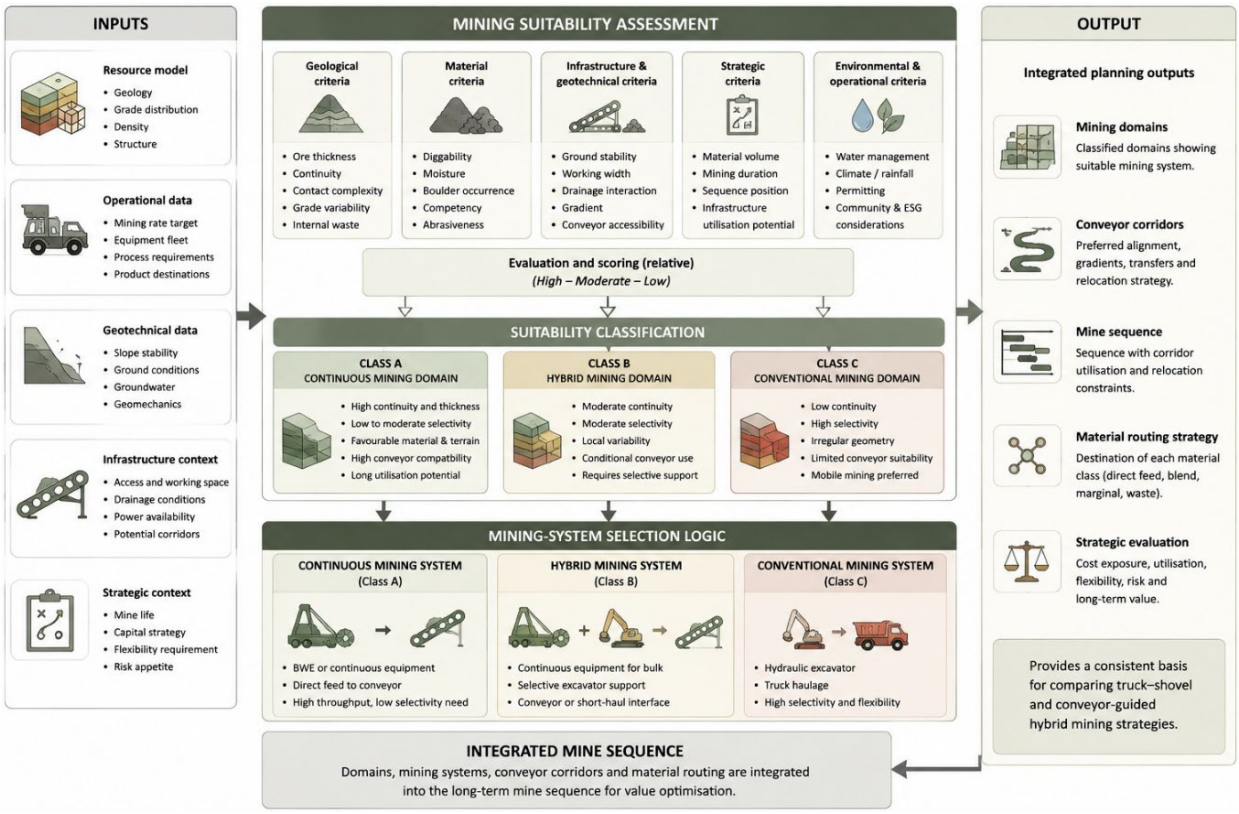


FIG 2 – Mining suitability and mining system selection logic.

The Class A–C boundaries are intentionally qualitative at this conceptual stage. Numerical thresholds for geological continuity, thickness variability, selectivity and material competence would require site-specific calibration using resource, geotechnical and operational data.

ADAPTIVE HYBRID CONTINUOUS MINING FRAMEWORK

The Adaptive Hybrid Continuous Mining Framework connects resource knowledge, mining suitability, extraction method, conveyor corridors, mine sequence and destination control within a single planning structure. Figure 3 presents the integrated structure of the proposed framework.

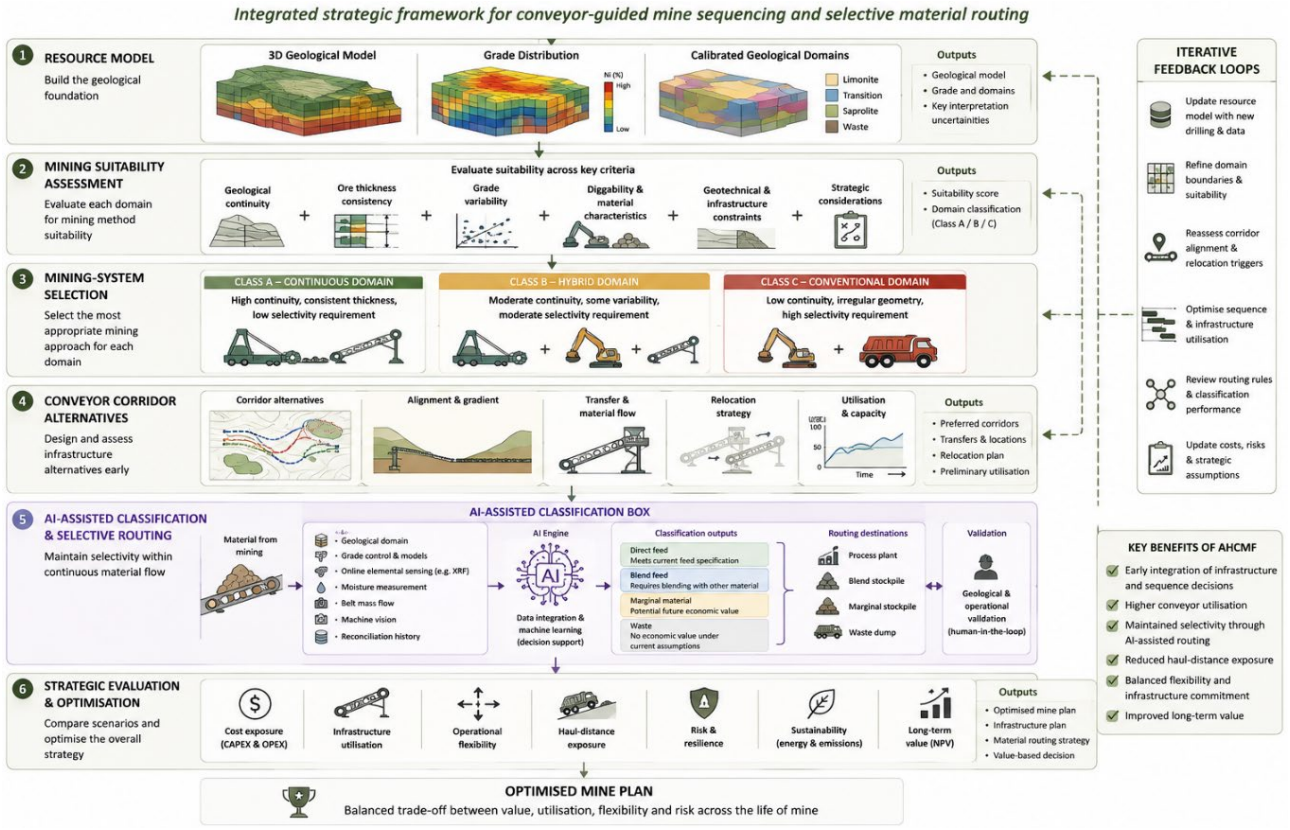


FIG 3 – Adaptive hybrid continuous mining framework.

Resource and mining domain definition

The resource model provides geological and grade information from which mining suitability domains are established.

Unlike conventional material classification based primarily on ore and waste definitions, AHCMF adds a mining system dimension. Each domain is evaluated according to whether continuous, hybrid or conventional mining is the more appropriate strategic response.

Strategic conveyor corridor definition

Potential conveyor corridors are identified according to mining direction, elevation development, material volume, location of suitable mining domains, transfer positions and expected duration of use.

The objective is not simply to establish the shortest transport distance. Corridor selection must consider the relationship between current infrastructure investment and future access to material.

A corridor with slightly greater initial length may provide greater strategic value if it remains usable through multiple mining stages and reduces relocation requirements.

Sequence infrastructure integration

Alternative mining sequences are evaluated together with corridor utilisation. A sequence that maximises early ore access may result in poor infrastructure utilisation or repeated conveyor relocation. A sequence developed solely around conveyor utilisation may delay valuable material or reduce operating flexibility.

AHCMF therefore treats the planning objective as a balance between resource value, infrastructure utilisation and sequence flexibility. Recent long-term scheduling research demonstrates the importance of integrating road and conveyor networks directly into production scheduling rather than treating them independently (Kamrani, Pourrahimian and Askari Nasab, 2025).

Figure 4 illustrates the intended relationship between mine development and infrastructure life. A principal conveyor backbone remains where practical while feeder systems and short haul interfaces move progressively with active mining areas.

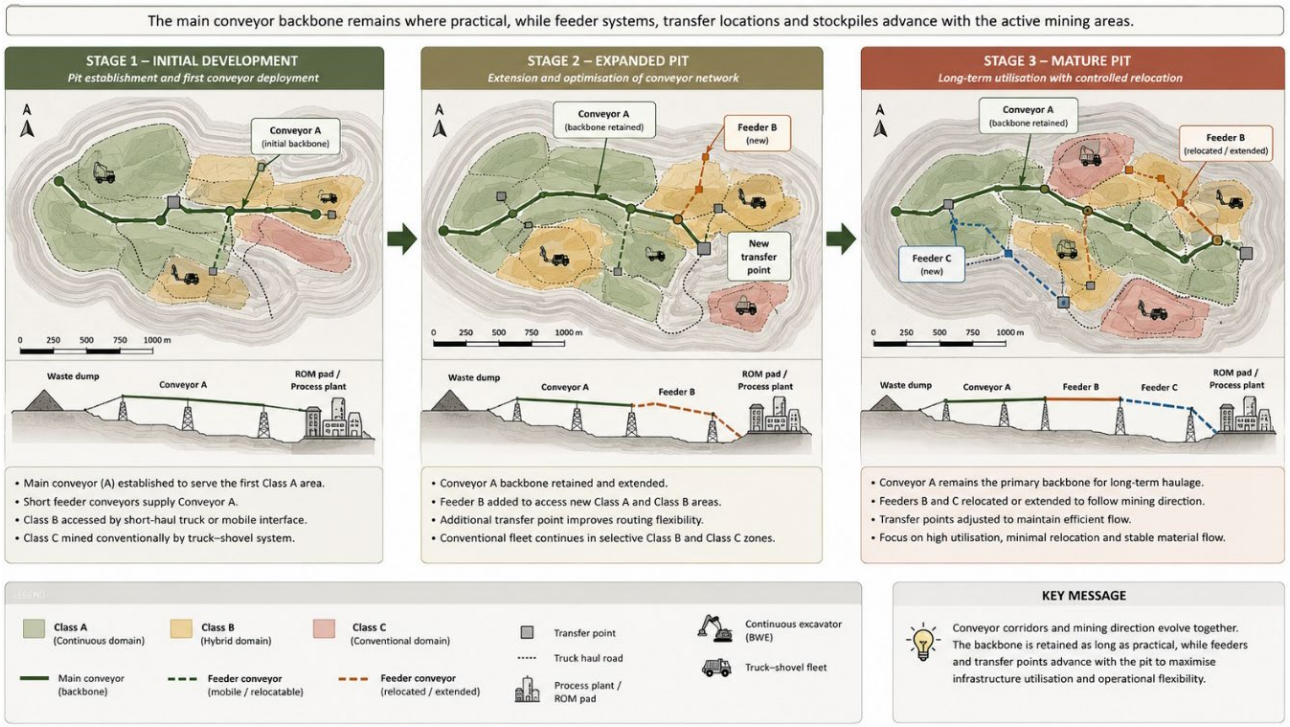


FIG 4 – Conceptual evolution of conveyor corridors with progressive mine development.

Material routing

Material extracted from each mining domain is routed according to its geological and economic characteristics.

Potential destinations include:

- direct process feed
- blending material
- marginal material stockpile
- waste storage.

Selective routing becomes particularly important where continuous mining produces a relatively uninterrupted material stream.

Strategic evaluation

Alternative mining configurations are assessed according to their combined effect on:

- infrastructure utilisation

- production continuity
- haulage exposure
- relocation requirements
- selective flexibility
- operational resilience
- capital commitment
- long-term value.

The framework therefore produces a strategic comparison rather than automatically selecting continuous mining as the preferred option.

HYBRID MINING ARCHITECTURE

AHCMF does not propose full replacement of conventional mobile mining. Continuous excavation equipment, represented conceptually by a bucket wheel excavator (BWE), is assigned to suitable Class A domains and selected Class B areas where material characteristics allow sustained excavation.

Conventional equipment remains necessary because laterite deposits may contain local conditions that reduce continuous mining efficiency. Hydraulic excavators, dozers and limited mobile haulage perform three principal functions. Figure 5 illustrates the conceptual hybrid mining architecture across the three mining suitability domains.



FIG 5 – Conceptual hybrid mining architecture across variable laterite domains.

Selective mining

Conventional equipment manages irregular ore contacts, internal waste and local areas requiring greater selectivity.

Development

Mobile equipment supports initial pit access, local infrastructure development and preparation ahead of continuous mining.

Contingency

Auxiliary equipment provides limited alternative production capacity during maintenance or disruption of the fixed system. This configuration prevents AHCMF from becoming dependent on a single extraction method and retains a degree of local operational flexibility.

Class A domains principally feed the conveyor network through continuous excavation. Class B domains combine continuous excavation and selective conventional equipment. Class C domains remain predominantly truck excavator mining based.

AI ASSISTED SELECTIVE MATERIAL ROUTING

Selective material routing is a central component of AHCMF because increasing continuity of material movement must not remove the ability to distinguish different material destinations.

AI assisted classification box concept

A conceptual AI assisted classification box is positioned along the conveyor system before final material destination.

Rather than relying on a single sensor, the framework proposes integration of several information sources.

Potential inputs include:

- resource and geological domains
- grade control data
- laboratory reconciliation
- online elemental sensing
- moisture measurements
- belt mass flow
- machine vision
- historical routing and reconciliation data.

The AI component is conceptualised as a data-fusion and decision-support layer rather than a specific trained algorithm. Its final architecture, classification logic and performance requirements would require site-specific development, calibration and validation.

Sensor based sorting systems have been applied to mineral processing for waste elimination and diversion between process routes, although the suitability of each sensing technology depends on the material and application (Robben and Wotruba, 2019).

AHCMF therefore does not assume that online X-ray fluorescence (XRF) provides independent and universally accurate nickel classification. The AI box is instead treated as a multi input decision support system.

Material classification outputs

The proposed classification architecture contains four principal routing categories.

Direct process feed

Material that satisfies the current destination requirements and can be routed directly toward processing or product preparation.

Blend material

Material that does not independently meet feed specifications but provides value when combined with complementary material.

Marginal material

Material that is not selected for immediate processing but may retain future value under different market, recovery or processing conditions.

Waste

Material without an identified processing or strategic stockpiling destination under the current planning assumptions. Figure 6 illustrates the proposed AI-assisted selective material routing architecture and its relationship with the four routing categories.

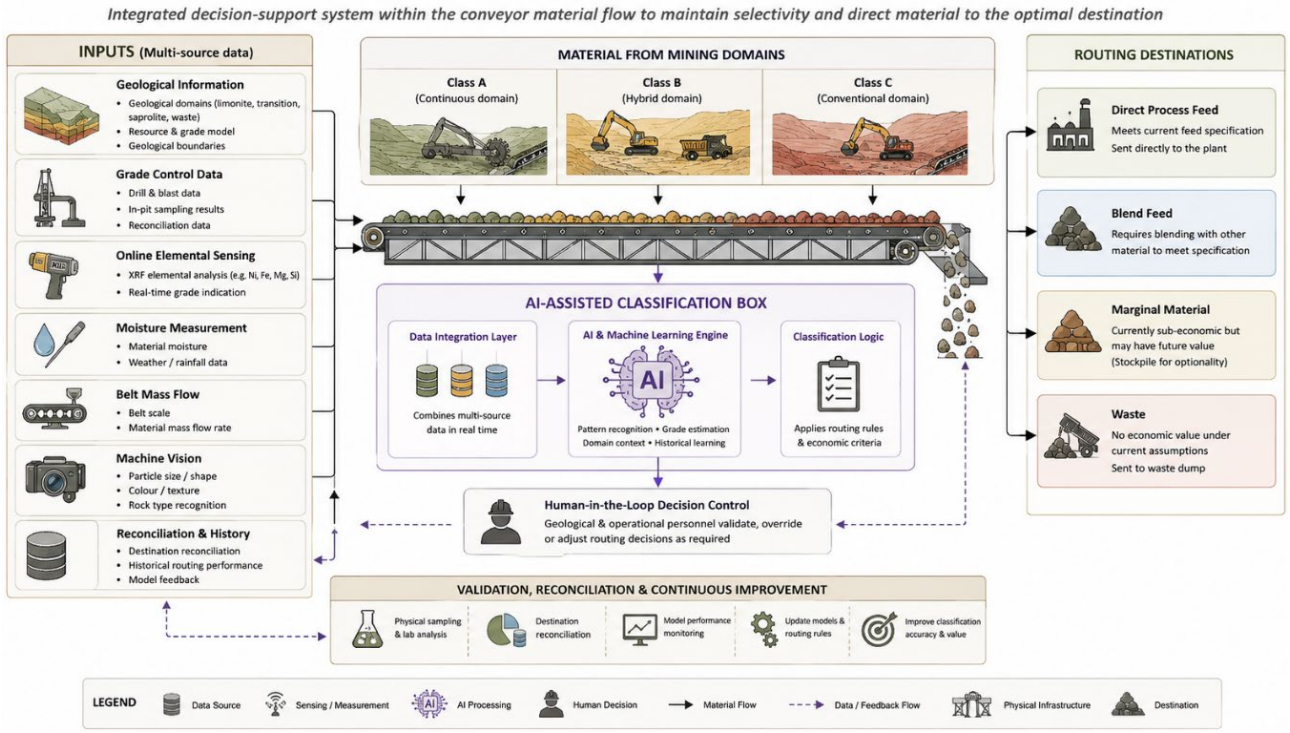


FIG 6 – AI assisted selective material routing architecture.

Human in the loop control

The proposed AI system provides classification recommendations rather than final autonomous decisions.

Geological and operational personnel retain responsibility for:

- validation of classification rules
- reconciliation
- changes to cut-off or feed criteria
- exception management
- destination governance.

Human oversight is particularly important because classification errors within a continuous material stream could propagate more rapidly than isolated routing errors within discrete truck movements.

The strategic role of AI is therefore to integrate multiple material information streams quickly enough to support selective routing while maintaining technical accountability.

WATER, GEOTECHNICAL AND RELIABILITY CONSTRAINTS

Water management

High rainfall is an important constraint in tropical laterite operations because rainfall influences pit accessibility, material moisture, ground conditions and infrastructure reliability.

Fixed conveyor corridors require greater consideration of drainage because infrastructure cannot be moved as readily as mobile haul routes. Surface water diversion, bench drainage, sumps and pumping locations therefore influence corridor suitability.

Water management is not developed as a standalone engineering design within this paper. It is incorporated as a strategic screening constraint during corridor and mining domain evaluation.

Geotechnical conditions

Continuous excavation and conveyor infrastructure require sufficiently stable operating areas.

Geotechnical constraints can affect corridor alignment, working width and excavation direction. Information concerning slope stability, floor condition and groundwater behaviour should therefore influence mining suitability before infrastructure commitment.

System reliability

Fixed infrastructure creates a different risk structure from truck fleets.

Truck haulage distributes capacity across multiple mobile units. Failure of one truck reduces capacity but generally does not interrupt the entire material flow.

A major conveyor or transfer point failure can affect a much larger proportion of production.

AHCMF therefore retains:

- auxiliary mobile equipment
- surge capacity
- buffer stockpiles
- planned maintenance integration.

The purpose of redundancy is not to reproduce full truck excavator mining capacity but to reduce exposure to complete loss of material movement during temporary fixed system disruption.

CONCEPTUAL STRATEGIC ECONOMIC ASSESSMENT

The economic assessment evaluates the difference in cost structure and exposure between conventional truck excavator mining and AHCMF. The assessment follows the strategic principle that mining alternatives should be compared in terms of their long-term economic implications rather than isolated operating costs. It does not represent a feasibility level capital estimate.

Truck excavator mining generally requires lower initial commitment to fixed material handling infrastructure, while substantial expenditure remains associated with mobile equipment, haul road development, fuel, tyres, maintenance and workforce.

AHCMF increases early infrastructure commitment through continuous mining equipment, conveyors, transfer systems, electrical infrastructure and material classification technology. The economic question is therefore not whether one system universally costs less.

The relevant strategic question is whether a higher fixed infrastructure commitment can be justified by reducing long-term exposure to variable mobile haulage requirements.

Nehring *et al* (2018) demonstrate that IPCC and truck excavator mining comparisons are sensitive to the strategic planning approach adopted, while Kamrani, Pourrahimian and Askari Nasab (2025) show that conveyor network configuration and crusher placement can alter long-term scheduling and economic performance. The resulting strategic cost-structure comparison is summarised in Table 3.

TABLE 3

Strategic cost structure comparison.

Characteristic	Truck excavator mining	AHCMF
Initial fixed infrastructure commitment	Lower	Higher
Mobile fleet dependence	Higher	Lower
Exposure to increasing truck haul distance	Higher	Potentially lower
Diesel dependency	Higher	Lower
Electricity dependency	Lower	Higher
Tactical relocation flexibility	High	Lower
Infrastructure utilisation requirement	Lower	High
Major system outage concentration	Lower	Higher
Selective mining capability	High	Maintained through hybrid fleet
Digital material routing requirement	Lower	Higher

AHCMF – Adaptive Hybrid Continuous Mining Framework.

The strategic economic hypothesis is that AHCMF becomes increasingly attractive where a long mine life, sustained material movement and increasing truck haul distances allow fixed infrastructure to maintain high utilisation. Conversely, high relocation frequency, short mine life or extensive selective mining can reduce the economic justification for fixed infrastructure.

No universal cost reduction percentage is proposed.

STRATEGIC SCENARIO ASSESSMENT

The applicability of AHCMF is evaluated using three conceptual operating environments.

Scenario A – low suitability

Scenario A represents operations characterised by short mine life, limited material volume, high geological variability, strong selectivity requirements and frequent changes in mining direction.

These conditions reduce conveyor utilisation and increase the value of mobile flexibility. Conventional truck excavator mining is therefore expected to retain a strategic advantage.

Scenario B – hybrid suitability

Scenario B represents deposits containing substantial continuous mining areas together with zones requiring selective conventional mining.

Conveyor infrastructure can serve sustained material flows while mobile equipment manages irregular mining domains.

This scenario represents the principal conceptual application of AHCMF.

Scenario C – high suitability

Scenario C contains long mine life, high material movement, relatively predictable mining direction, substantial continuous domains and increasing exposure to long truck haul distances.

Under these conditions, early evaluation of conveyor corridors becomes increasingly relevant. Figure 7 illustrates the directional suitability envelope across the principal strategic drivers, while Table 4 summarises the three conceptual operating scenarios.

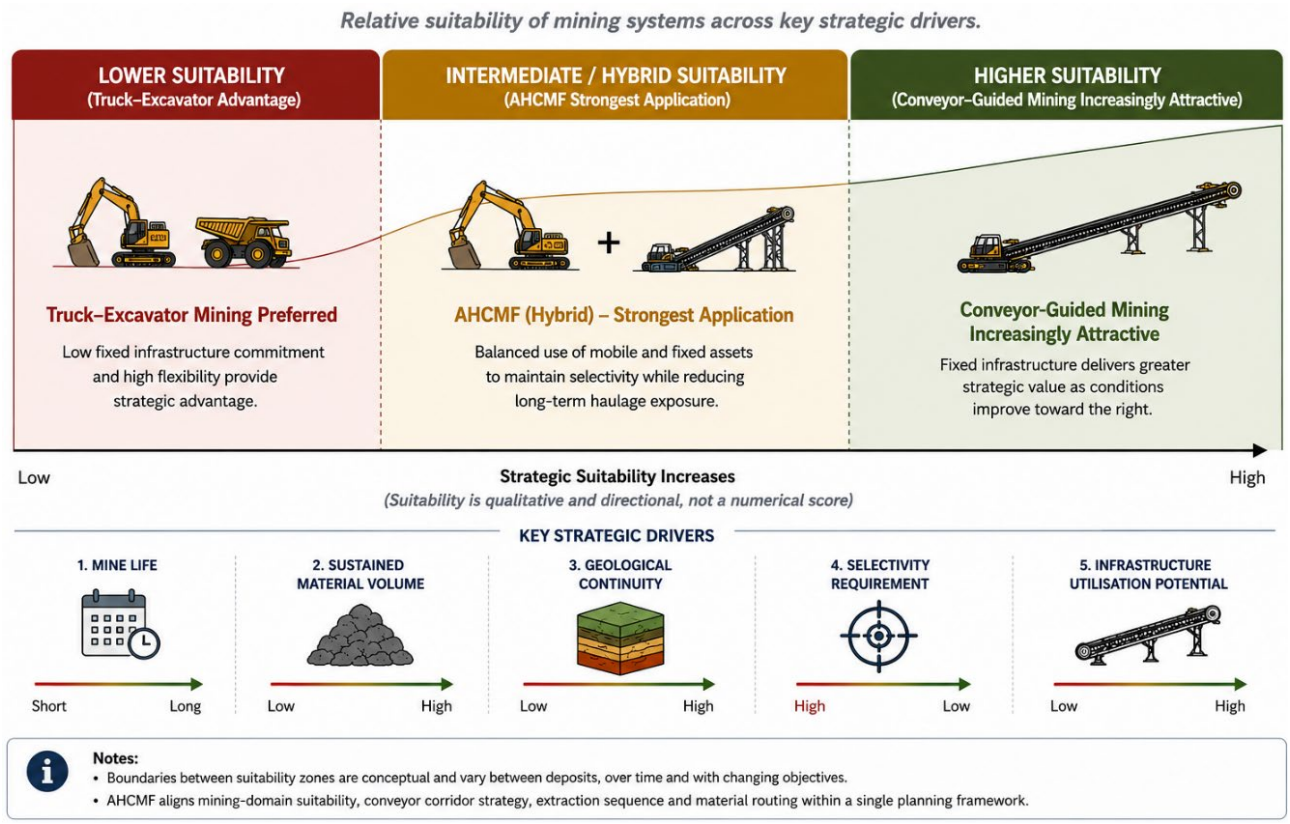


FIG 7 – Strategic suitability envelope for conveyor guided sequencing.

TABLE 4

Strategic suitability scenarios.

Strategic factor	Scenario A	Scenario B	Scenario C
Mine life	Short	Moderate	Long
Material volume	Low	Moderate–high	High
Geological continuity	Low	Moderate	High
Selectivity requirement	High	Moderate	Low–moderate
Conveyor utilisation	Low	Moderate	High
Infrastructure relocation	Frequent	Manageable	Limited
Truck haul distance exposure	Low	Increasing	High
AHCMF suitability	Low	Strong hybrid case	High

The resulting strategic relationship is directional rather than deterministic. Increasing mine life, material volume, haulage distance, geological continuity and corridor utilisation generally strengthen the case for conveyor guided sequencing. Increasing selectivity, sequence volatility and relocation frequency have the opposite effect.

STRATEGIC MARGINAL MATERIAL MANAGEMENT

Laterite mining commonly generates material close to current processing or economic thresholds. Immediate classification of all sub economic material as permanent waste can remove future optionality where commodity price, processing technology, metallurgical recovery or blending requirements change.

AHCMF therefore introduces strategic marginal material management. The terminology intentionally avoids treating marginal stockpiles as Ore Reserves or Mineral Resources. Material identified as

marginal can be directed to dedicated stockpiles where separation is justified by expected future option value and handling cost.

The AI assisted classification system can support this approach by maintaining consistent material categories and linking physical stockpiles with geological and reconciliation information. The value of the concept is not an assumed increase in recoverable tonnes. Its value is preservation of optionality.

Lane (1988) demonstrates the importance of economic cut-off decisions in defining ore value, and changes in economic conditions can alter the preferred treatment of marginal material. Strategic stockpiling therefore provides a mechanism for postponing irreversible disposal decisions where future value remains plausible.

STRATEGIC RISK ASSESSMENT

AHCMF creates both opportunities and new sources of risk. The principal risks are related to geological suitability, infrastructure commitment, reliability and material classification. Table 5 summarises the principal strategic risks and corresponding planning responses.

TABLE 5

Principal strategic risks and planning responses.

Risk	Strategic consequence	Planning response
Geological variability	Reduced continuous mining utilisation	Class A/B/C classification
Poor corridor selection	Under utilised or stranded infrastructure	Alternative corridor scenarios
Sequence change	Reduced conveyor utilisation	Schedule sensitivity testing
Conveyor outage	Material flow interruption	Mobile contingency and surge capacity
Classification error	Dilution or material misrouting	Human in the loop validation
Sensor failure	Reduced routing confidence	Redundant/manual verification
High water inflow	Corridor or operating disruption	Early water management integration
Geotechnical instability	Infrastructure and access constraints	Geotechnical screening
Short mine life	Insufficient infrastructure utilisation	Early strategic suitability assessment
Capital escalation	Reduced project attractiveness	Scenario based economic evaluation

NB: The risk assessment reinforces the central premise of the framework: conveyor selection should not be separated from the mine planning context in which the infrastructure must operate.

DISCUSSION

Conveyor rigidity as a strategic planning variable

Conveyor systems are frequently characterised by lower tactical flexibility than truck haulage. This remains a genuine disadvantage and should not be minimised. However, rigidity also creates a requirement to define mining direction, infrastructure utilisation and sequence stability earlier in the planning process.

The relevant comparison is therefore not simply flexibility versus inflexibility. It is a comparison between different forms of strategic value. Truck excavator mining systems preserve option value through mobility. Conveyor guided systems exchange part of that mobility for increased infrastructure

commitment and potentially lower exposure to long distance mobile haulage. AHCMF seeks a balance between these two characteristics rather than maximising either one.

Selectivity within continuous material movement

The second strategic issue is selectivity. Continuous excavation and conveying are most easily applied to relatively homogeneous bulk materials. Nickel laterite does not necessarily provide this condition. Different material domains may require separate destinations because of nickel grade, iron, magnesium, silica, moisture, mineralogy or processing requirements.

For this reason, the value of AHCMF depends not only on the ability to transport material continuously but also on the ability to preserve material intelligence throughout the transport chain. The AI assisted classification box responds to this requirement. Its contribution is not autonomous ore definition. It is the integration of multiple information sources into a faster material routing decision process.

Role of the AI classification box

The AI box is particularly important in Class A and Class B domains where material enters a higher throughput conveyor system. Classification information can potentially be updated continuously as material passes through the system.

However, real time sensing does not eliminate the need for resource modelling, grade control drilling, face observations, sampling or laboratory reconciliation. Sensor based sorting systems have demonstrated value in material diversion across several mineral applications, but sensor performance remains application specific (Robben and Wotruba, 2019).

Nickel laterite application therefore requires site specific calibration and validation. The AI box should consequently be viewed as a decision support architecture rather than an assumed fully autonomous technology.

Strategic contribution

The main contribution of AHCMF is the integration of four decisions that are frequently treated separately:

1. Where continuous mining is geologically suitable.
2. Where fixed material handling infrastructure should be located.
3. How extraction sequence should respond to infrastructure commitment.
4. How selectivity should be maintained after material enters continuous transport.

The framework therefore moves beyond a simple comparison of truck haulage and conveyor operating cost.

Its central proposition is that mining method, infrastructure and material information should be integrated before long-term sequencing decisions become fixed.

LIMITATIONS AND FUTURE WORK

The framework is conceptual and has not been validated against a specific operating mine.

No detailed three dimensional resource model, production schedule, BWE productivity simulation, conveyor mechanical design, hydrological model, geotechnical design or feasibility level economic estimate is presented.

The assumed 10 million bcm/a mining rate provides a representative planning scale only. The suitability of BWE excavation would require testing of material strength, abrasivity, moisture, boulder occurrence, trafficability and selectivity. Similarly, the AI assisted classification system requires site specific evaluation.

The study does not demonstrate that online XRF, machine vision or another individual sensing method can provide sufficient nickel laterite classification accuracy at the proposed throughput. Future research should progress through a series of technical validation stages.

First, a three dimensional nickel laterite resource model should be classified using the proposed Class A, Class B and Class C suitability methodology.

Second, multiple life-of-mine schedules should be generated with explicit conveyor corridor, utilisation and relocation constraints.

Third, alternative truck excavator mining, hybrid and conveyor guided scenarios should be evaluated using comparable economic assumptions.

Fourth, discrete event simulation should assess interactions between continuous excavation, conveyor capacity, buffer storage and auxiliary mobile fleets.

Fifth, representative laterite material should undergo sensor testing to determine appropriate combinations of elemental sensing, moisture measurement and imaging.

Finally, the integrated framework should be evaluated using economic metrics including net present value, infrastructure utilisation, haulage distance exposure and schedule resilience.

Key future quantitative indicators should include:

- conveyor utilisation
- conveyor relocation frequency
- truck haul distance reduction
- classification accuracy
- dilution and ore loss
- material misrouting
- fixed system downtime
- sequence deviation
- discounted project value.

CONCLUSIONS

This paper presents the Adaptive Hybrid Continuous Mining Framework as a conceptual strategic planning architecture for large scale nickel laterite operations.

The framework does not propose universal replacement of conventional truck excavator mining. Instead, it begins by dividing the deposit into continuous, hybrid and conventional mining domains according to geological, material and operating characteristics.

Three principal conclusions are identified.

First, conveyor infrastructure can materially influence mine development and should therefore be considered during long-term sequencing rather than only after the production schedule has been established. Conveyor corridor and extraction sequence are interdependent strategic decisions.

Second, hybrid mining provides a practical planning mechanism for retaining selectivity. Continuous excavation and conveyor transport can be concentrated within suitable domains while conventional equipment remains available in areas of geological complexity and for development or contingency requirements.

Third, increasing continuity of material movement creates a requirement for stronger selective routing. The proposed AI assisted classification box provides a conceptual decision support layer that combines geological information with potential real time sensor inputs to support routing between direct feed, blending, marginal material and waste. Geological governance remains part of the decision process.

The assessment does not demonstrate a predetermined reduction in capital cost, operating cost, workforce or emissions. Instead, it identifies the conditions under which conveyor guided mining may warrant further detailed evaluation.

The strategic case strengthens as mine life, sustained material volume, geological continuity, truck haul distance exposure and conveyor utilisation increase. It weakens as selective mining requirements, sequence volatility and infrastructure relocation frequency increase.

The principal contribution of AHCMF is therefore not a particular excavator, conveyor or AI technology. It is the integration of geological suitability, mining system selection, fixed infrastructure, extraction sequence and material intelligence within a single strategic mine planning framework.

ACKNOWLEDGEMENTS

The author acknowledges the mining, planning, and operational personnel whose practical discussions and observations contributed to the development of the concepts presented in this paper.

The author also acknowledges the AusIMM Strategic Mine Planning Conference review committee for constructive feedback that assisted in refining the strategic focus of the study.

REFERENCES

Hartman, H L and Mutmansky, J M, 2002. *Introductory Mining Engineering*, second edition (John Wiley and Sons: New York).

Hustrulid, W A, Kuchta, M and Martin, R K, 2013. *Open Pit Mine Planning and Design*, third edition (CRC Press: Boca Raton).

Kamrani, A, Pourrahimian, Y and Askari Nasab, H, 2025. Semi mobile in pit crushing and conveying vs truck excavator mining systems: long term scheduling with road and conveyor networks integration, *Expert Systems with Applications*, 268:126122.

Lane, K F, 1988. *The Economic Definition of Ore*: *Cut off Grades in Theory and Practice* (Mining Journal Books: London).

Nehring, M, Knights, P F, Kizil, M S and Hay, E, 2018. A comparison of strategic mine planning approaches for in pit crushing and conveying and truck/excavator mining systems, *International Journal of Mining Science and Technology*, 28(2):205–214.

Robben, C and Wotruba, H, 2019. Sensor based ore sorting technology in mining – past, present and future, *Minerals*, 9(9):523.

Mining complexes and value chain optimisation

Integrating planning horizons and capital investments for mining complexes with uncertainty using multi-level stochastic optimisation

L Findlay[1] and R Dimitrakopoulos[2]

1. Research Assistant, COSMO Stochastic Mine Planning Laboratory, Montreal QC H3A 0E8, Canada. Email: liam.findlay@mail.mcgill.ca
2. Director, COSMO Stochastic Mine Planning Laboratory, Montreal QC H3A 0E8, Canada Email: roussos.dimitrakopoulos@mcgill.ca

ABSTRACT

Conventional mine planning makes capital investment decisions and long-term or strategic plans which set production targets for short-term or tactical planning. However, these targets may not be achievable because the strategic plan does not consider the level of operational detail that is present in tactical planning. Integrating time horizons and capital investments can ensure that targets are achievable in the short-term and long-term decisions are optimal. The proposed framework evaluates the first year of the long-term plan as a 12-month tactical schedule during the hyper-heuristic driven optimisation of strategic decisions. Specifically, the assessment of shovel movement feasibility and the characteristics of process material are more precise which allows the strategic optimisation to consider short-term feasibility. Risk is managed by accounting for uncertainty in the mineral deposits and processing performance. A case study at an open pit gold mining complex demonstrates the method, showing improved target compliance in the short-term and higher expected discounted cash flows in the long-term.

INTRODUCTION

Strategic planning of mining complexes aims to maximise net-present-value (NPV) by making strategic decisions including capital investments (capex) for integrated value chains while addressing uncertainty from key inputs to manage the associated risk (Montiel and Dimitrakopoulos, 2015; Goodfellow and Dimitrakopoulos, 2015; Dimitrakopoulos and Lamghari, 2022). Planning is typically hierarchical with capex decisions that dictate production rates being made first, then long-term planning outlines annual extraction sequences and targets based on the production rates, and finally medium- and short-term planning define more immediate schedules while aiming to meet long-term targets (Gershon, 1983; Hustrulid, Kuchta and Martin, 2013; Blom, Pearce and Stuckey, 2019). Long-term planning typically ignores operational details to keep the optimisation tractable for such a large decision-making scope. Medium and short-term planning include these details on shorter time frames to ensure schedules are feasible. However, this represents a misalignment where there is no guarantee that the targets set by the long-term plan are achievable (or optimal) in the short-term when those extra details are accounted for. The present work proposes an integrated framework where the first 12 months are evaluated while the capex and long-term plan decisions are being optimised. The monthly evaluation examines the first-year performance by considering the explicit movement of shovels to execute the schedule, and by simulating the processing performance of feed materials in monthly blended batches.

Existing work has demonstrated the general value of integrating mine planning timescales. However, some works are deterministic and fail to account for uncertainty (Nehring *et al,* 2010, 2012). Levinson and Dimitrakopoulos (2023) integrate yearly and monthly planning within a simultaneous stochastic optimisation framework for mining complexes and uses reinforcement learning to decide the monthly schedule. However, the monthly planning approach assumes linear blending and processing which may not represent reality accurately. None of the existing methods connects capex to the integrated framework even though these are the most impactful decisions made in a mining operation. For example, buying an additional shovel not only increases mining capacity in tonnes per annum but may also enable more mining areas to be open simultaneously which could favour a different extraction sequence. Therefore, the two primary contributions of this work are: i) improving the medium-term evaluation by considering non-linear blending and stochastic processing plant performance; and ii) extending the integrated framework to consider capex decisions.

An innovative approach to optimising this integrated planning formulation is necessary due to the scale of industrial mining complexes, the presence of non-linear components, and the fact that it aims to optimise several problems (ie the different timescales) simultaneously. To handle the multiple time levels, the problem can be viewed as a multi-level stochastic integer program, with each level being optimised by a suitable algorithm. An evolutionary algorithm (Coello Coello, Lamont and Van Veldhuizen, 2007) is used to search for optimal capex decisions because it can be implemented in parallel, allowing it to scale up in computing resources. Long-term decisions are optimised using a hyperheuristic strategy built on simulated annealing (Lamghari and Dimitrakopoulos, 2020), and the medium-term schedules are generated using a decision policy trained using reinforcement learning (Sutton and Barto, 2018).

The following section presents the proposed framework including the multi-level stochastic integer programming (MSIP) formulation and the solution method to handle the complex optimisation problem. Subsequently, an application using a gold mining complex is shown to demonstrate the proposed framework. Finally, conclusions and future work are discussed.

METHOD

Multi-level stochastic integer programming formulation

Selecting the optimal capex and long-term planning decisions while examining the short-term performance requires simultaneously optimising each set of decisions. However, there is a hierarchical relationship such that the solution space of the short-term plan cannot be defined unless the capex and long-term decisions are made, whereas the capex decisions can be made freely but require the long-term and short-term decisions to be optimised in order to be evaluated. To capture these dependencies between capital investments and production horizons, the system is conceptualised as a multi-level hierarchical optimisation problem under a single decision-maker. Let $\Omega, \mathrm{X},$ and Ψ represent the capex, long-term plan, and short-term plan decision variables, respectively. The global problem can be structured as:

$$\max_{\Omega,\mathrm{X},\Psi} \mathbb{E}_{\mathbb{S},\mathbb{R}}[\mathcal{F}(\Omega, \mathrm{X}, \Psi)] \tag{1}$$

Subject to:

$$\Omega \in \mathcal{C}_1 \tag{2}$$

$$\mathrm{X} \in \mathcal{C}_2(\Omega) \tag{3}$$

$$\Psi \in \mathcal{C}_3(\Omega, \mathrm{X}) \tag{4}$$

Where $\mathcal{F}$ is the global objective function, the expectation of which is being maximised over the combined set of stochastic scenarios that represent geostatistical supply $(\mathbb{S})$ and processing plant performance $(\mathbb{R})$ uncertainty. $\mathcal{C}_1, \mathcal{C}_2(\Omega),$ and $\mathcal{C}_3(\Omega, \mathrm{X})$ represent the hierarchical sets of constraints applied to the capex, long-term, and short-term decision levels, respectively, which are functions of decisions made in the upper-level(s). The remainder of this section provides the explicit definition of the capex level of the global MSIP formulation, beginning with Table 1 which introduces key notation.

TABLE 1

Notation used in the MSIP formulation.

Set	Set of long-term periods	$t \in \mathbb{T}$
Set	Set of capital investment options	$k \in \mathbb{K}$
Constant	Discounted capital cost per unit of option k in period t	p_{kt}
Variable	Decision variable representing the number of units of option k purchased in period t	w_{kt}
Constant	Upper limit of the number of units that can be purchased for option k in period t	U_{kt}

The global objective function from Equation 1 can be expanded to:

$$\max_{\Omega,\mathrm{X},\Psi} \mathbb{E}_{\mathbb{S},\mathbb{R}}[\mathcal{F}(\Omega,\mathrm{X},\Psi)] = \max_{\Omega\in\mathcal{C}_1}\left[-\sum_{k\in\mathbb{K}}\sum_{t\in\mathbb{T}} p_{kt}\cdot w_{kt} + \max_{\mathrm{X}\in\mathcal{C}_2(\Omega)}\left[\mathbb{E}_{\mathbb{S}}[\mathcal{G}(\mathrm{X})] + \max_{\Psi\in\mathcal{C}_3(\Omega,\mathrm{X})}\left[\mathbb{E}_{\mathbb{S},\mathbb{R}}[\mathcal{H}(\Psi)]\right]\right]\right] \quad (5)$$

where $\mathcal{G}(\mathrm{X})$ represents the objective function contribution of the long-term decisions, X, and $\mathcal{H}(\Psi)$ represents the contribution of the short-term decisions, Ψ. The capital investment decisions $\Omega = \{w_{kt} | k \in \mathbb{K}, t \in \mathbb{T}\}$ are subject to the following constraint which ensures that the purchased quantities are non-negative integers with defined upper limits:

$$w_{kt} \in \{0,1,\dots,U_{kt}\}\ \forall k \in \mathbb{K}, t \in \mathbb{T} \quad (6)$$

The placement of expectations in the objective function illustrates the treatment of uncertainty by the formulation. The capex decisions are made scenario-independently in the presence of uncertainty, while both the long-term and medium-term schedules include scenario-dependent decisions that react to uncertainty. Specifically, the long-term decisions include block extraction, cluster-based destination policy, and downstream process flow decisions, of which the latter two react to supply uncertainty. The medium-term level includes the same general decisions as the long-term, but more detailed analysis of the loading equipment movements and processing plant performance are examined during the medium-term evaluation. To accomplish this, the medium-term includes response variables which also react to the processing plant performance uncertainty.

Parallelised multi-level optimisation algorithm

To optimise the multi-level SIP formulation, a continuous loop structure (shown in Figure 1) is used such that the medium-term schedule is optimised for each iteration of the long-term schedule, which itself is optimised for each iteration of the capex decision search. With this structure, each decision level is handled by its own solving algorithm while finding a combined solution that maximises the unified multi-level objective function (Equations 1, 5).

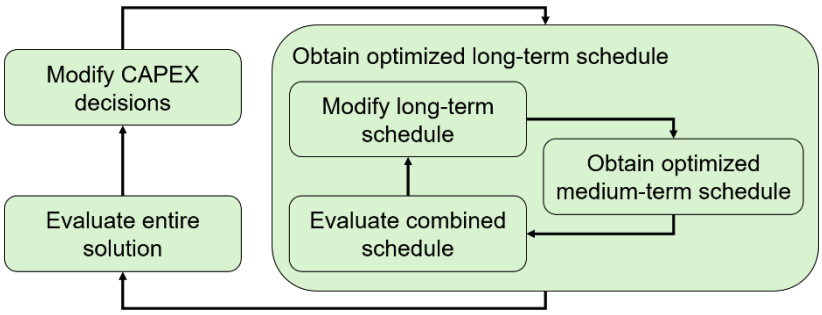


FIG 1 – Continuous loop structure for optimising capital investment, long-term schedule and short-term schedule decisions.

Designing the solving framework to work in parallel allows for the method to be implemented on multiple computing resources and get optimised solutions in less time than a single-threaded approach. To accomplish this, the capex decision level is optimised using an evolutionary algorithm (Coello Coello, Lamont and Van Veldhuizen, 2007). In each generation, a population of capex decision vectors (Ω) is evaluated and the next generation is created through selection, crossover, and mutation. The evaluation uses the unified objective function and requires optimising the long and medium-term schedules for each capex decision vector which is done in parallel.

The long-term schedule decisions (X) are optimised using multi-neighbourhood adaptive simulated annealing with a hyperheuristic strategy that selects low-level heuristics (Lamghari and Dimitrakopoulos, 2020). The medium-term schedules are found using a reinforcement learning policy based on an extension of the framework developed by Levinson, Dimitrakopoulos and Keutchayan (2023). The policy is pre-trained offline using a simulated multi-level environment. This exposes the agent to a variety of capex and long-term decisions, allowing it to learn generalised operational rules. During deployment, this pre-trained policy is evaluated directly to solve medium-term scheduling decisions at every iteration of the long-term optimisation that includes a perturbation of the first year.

APPLICATION AT A GOLD MINING COMPLEX

A gold mining complex (shown in Figure 2) includes two open pit areas on a single deposit and mineral processing streams that include a waste dump, two crushers, a screen separation facility, a grinding circuit, a flotation plant, and a leach pad. The screen separation facility directs fine materials to the grinding circuit and coarse materials to the leach pad. The leach pad and flotation plant produce gold concentrate that is sold to customers. Material supply uncertainty is quantified by a set of ten equiprobable geostatistical simulations that cover the two pit areas. Capex decisions are available to increase production capacities (in tons per unit time) at each facility in the mining complex. Mining capacity is shared across both pit areas because shovels can move between them.

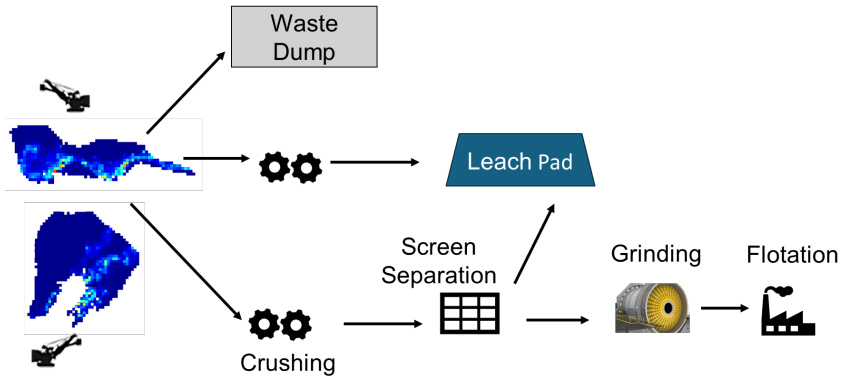


FIG 2 – Overview of the gold mining complex.

The application to demonstrate the proposed integrated framework is compared against a separated case where the capex and long-term decisions are optimised without considering the first-year performance and the first-year 12-month schedule is optimised subsequently. The following results are obtained by running each case for 48 hrs on a computer with a 6-core 3.60 GHz CPU and 32 GB of memory, using six threads in parallel with a matching population size. The reinforcement learning policy is pre-trained using the same instance as the application.

The first result which can be examined is the comparison of capex decisions made by each schedule which dictate the production rates throughout the mining complex. Figure 3 shows that the two cases both invested in extra production capacity for the mines, the second crusher (which feeds the grinding and flotation stream), and the mill which refers to the capacity of both the grinding and flotation facilities. However, the amounts and timing of the investments are different. A key difference is that the base case chose to increase the crushing and mill capacities in the first year whereas the integrated case made similar investments but waited until the second year where it also purchased additional mining capacity. The subsequent results showing the production forecasts of the two schedules will highlight the consequences of these different capex decisions.

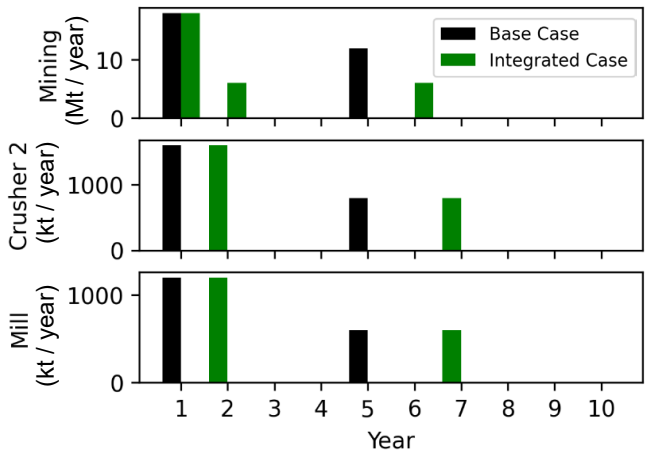


FIG 3 – Comparison of capex decisions to increase production rates in the integrated and base cases.

Figure 4 shows the production forecasts and risk profiles of the first year for the two cases. The dashed lines in the mill tonnage and cash flow plots represent the 10th and 90th percentiles of the risk profile, while the solid lines represent 50th percentiles for the related attributes. The base case has a higher mill capacity as determined by the capex investments but the short-term schedule clearly did not find enough high-grade material in the first seven months to fill the higher capacity. The integrated case with its lower mill capacity was able to provide more consistent feed material and ultimately achieved 3.4 per cent more cash flow than the base case.

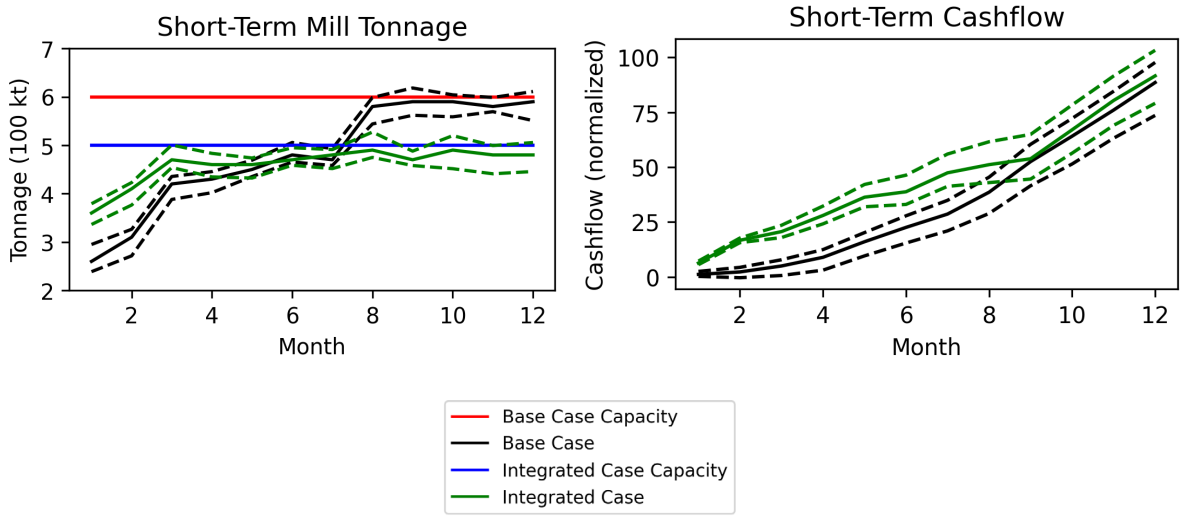


FIG 4 – Comparisons of scheduled mill tonnage (left) and cash flow (right) over the first year for the integrated and base cases.

The long-term forecasts show the differences between the two cases and also show the differences between the base case long-term forecast and the forecast achieved after optimising the 12-month schedule. On Figure 5, the black lines represent the initial forecast set by the long-term plan, and the yellow lines represent the forecast for the schedule after analysing the first year. In terms of mill tonnage, the lower total tonnage from the first year is evident where the yellow lines are much lower than the black lines for year one. The analysis also shows that the base case long-term plan overestimates total net-present-value (NPV) of the project by 2 per cent. By including the first-year analysis inside the multi-level optimisation framework, the integrated case shows a 1 per cent increase in total project NPV.

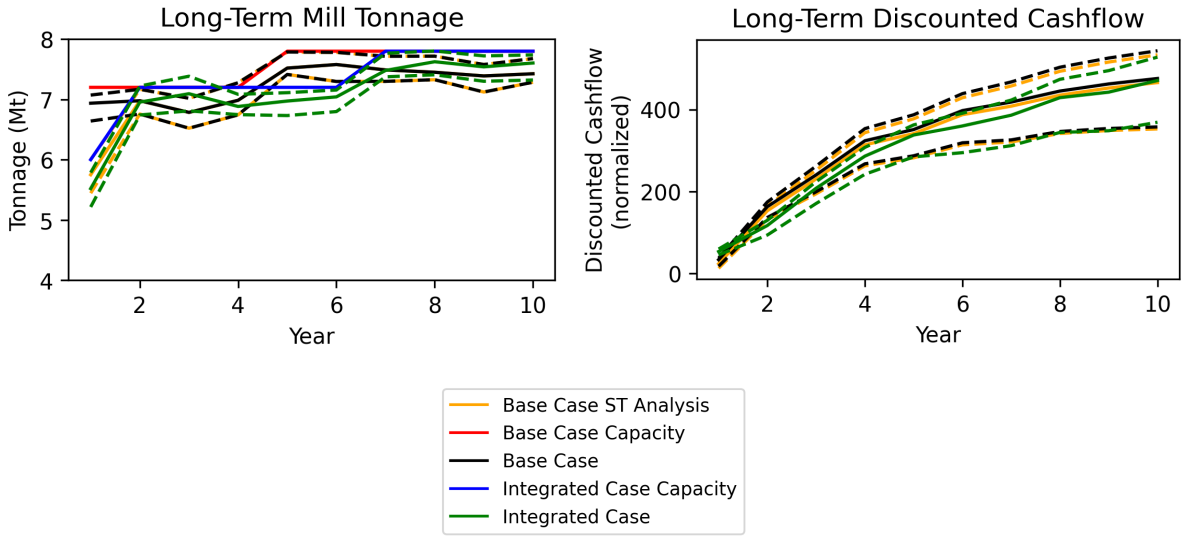


FIG 5 – Comparisons of scheduled mill tonnage (left) and cash flow (right) over the ten years for the integrated and base cases.

CONCLUSIONS

The present work proposes a multi-level stochastic optimisation framework for mining complexes that integrates decision-making levels including capex investments, long-term planning, and medium-term planning. The medium-term evaluation allows the schedule to consider the explicit equipment movement and performance in mineral processing using a finer level of granularity than long-term plans typically consider. Performing this analysis during the optimisation of major investments allows these decisions to be given the thorough analysis that is warranted by the magnitude of the investments and their impacts on the operations. Results show that integrating the decisions leads to better target compliance in the short-term and increased discounted value in the long-term. The 1 per cent increase in NPV quantifies the benefits of the framework.

Future work can examine environmental and emissions analysis that also helps align planning frameworks with sustainability objectives and aspects such as truck haulage which will provide even more accurate analysis of the operational feasibility of the integrated schedules.

ACKNOWLEDGEMENTS

The work in this paper was funded by the NSERC Discovery Grant 239019, and the industry consortium members of COSMO Stochastic Mine Planning Laboratory (AngloGold Ashanti, Anglo American/DeBeers, Agnico Eagle, BHP, IAMGOLD, Kinross Gold, Newmont, Vale SA and Vale Base Metals), and the Canada Research Chairs Program.

REFERENCES

Blom, M, Pearce, A R and Stuckey, P J, 2019. Short-term planning for open pit mines: a review, *International Journal of Mining, Reclamation and Environment,* 33:318–339. https://doi.org/10.1080/17480930.2018.1448248

Coello Coello, C A, Lamont, G B and Van Veldhuizen, D A, 2007. *Evolutionary algorithms for solving multi-objective problems* (Springer, New York).

Dimitrakopoulos, R and Lamghari, A, 2022. Simultaneous stochastic optimization of mining complexes - mineral value chains: an overview of concepts, examples and comparisons, *International Journal of Mining, Reclamation and Environment,* 36:443–460. https://doi.org/10.1080/17480930.2022.2065730

Gershon, M E, 1983. Optimal mine production scheduling: evaluation of large scale mathematical programming approaches, *International Journal of Mining Engineering*, 1:315–329. https://doi.org/10.1007/BF00881548

Goodfellow, R and Dimitrakopoulos, R, 2015. Stochastic Optimization of Open Pit Mining Complexes with Capital Expenditures: Application at a Copper Mining Complex, in Proceedings of the 37th International Symposium on Application of Computers and Operations Research in the Mineral Industry, pp 657–667 (Society for Mining, Metallurgy and Exploration: Fairbanks, Alaska).

Hustrulid, W A, Kuchta, M and Martin, R K, 2013. *Open Pit Mine Planning and Design*, two volume set and CD-ROM pack (CRC Press: London).

Lamghari, A and Dimitrakopoulos, R, 2020. Hyper-heuristic approaches for strategic mine planning under uncertainty, *Computers and Operations Research,* 115:104590.

Levinson, Z and Dimitrakopoulos, R, 2023. Connecting planning horizons in mining complexes with reinforcement learning and stochastic programming, *Resources Policy,* 86:104136. https://doi.org/10.1016/j.resourpol.2023.104136

Levinson, Z, Dimitrakopoulos, R and Keutchayan, J, 2023. Simultaneous stochastic optimization of an open-pit mining complex with preconcentration using reinforcement learning, *Applied Soft Computing,* 138:110180. https://doi.org/10.1016/j.asoc.2023.110180

Montiel, L and Dimitrakopoulos, R, 2015. Optimizing mining complexes with multiple processing and transportation alternatives: An uncertainty-based approach, *European Journal of Operational Research*, 247:166–178. https://doi.org/10.1016/j.ejor.2015.05.002

Nehring, M, Topal, E, Kizil, M and Knights, P, 2010. An investigation to integrate optimum long-term planning with short planning in underground mine production scheduling, in *Proceedings of the Mine Planning and Equipment Selection Conference,* pp 141–154.

Nehring, M, Topal, E, Kizil, M and Knights, P, 2012. Integrated short-and medium-term underground mine production scheduling, *Journal of the Southern African Institute of Mining and Metallurgy,* 112:365–378.

Sutton, R S and Barto, A G, 2018. *Reinforcement learning: An introduction* (MIT Press).

Strategic value optimisation for an integrated copper open pits and underground mining complex

B Gadelha[1], T Câmara[2], L Montiel[3] and J Todd[4]

1. RDP Manager, Vale Base Metals, Sudbury ON, Canada. Email: bruno.gadelha@vale.com
2. Specialist Engineer, Vale Base Metals, Belo Horizonte MG, Brazil. Email: tais.camara@vale.com
3. Principal Engineer, Vale Base Metals, Toronto ON, Canada. Email: luis.petro@vale.com
4. Head of Strategic Mine Planning, Vale Base Metals, Toronto ON, Canada. Email: jody.todd@vale.com

ABSTRACT

This study examines the integration of two copper deposits—one developed as an open pit and the other combining open-pit and underground mining, for a total of three operating sources—located within the same region. The objective was to evaluate potential synergies through shared use of a centralised processing plant and tailings storage facilities. The integrated mine plan was built around the proposed plant capacity of the most advanced project and assessed possible mill expansions supported by incremental capital expenditure (CAPEX) to process additional ore, considering licensing timelines and operational constraints (ie total primary mining movement). Multiple start-year scenarios were modelled, accounting for underground development requirements and alternative open pit/underground transition strategies. A strategic planning tool enabled ore blending through cut-off grade optimisation from all three operations to optimise copper grades and plant utilisation. Sensitivity analyses on CAPEX, operating costs (OPEX), and cost assumptions were conducted to validate outcomes. Results show that integration increases net present value (NPV) relative to stand-alone cases, creating significant value and improving resource and infrastructure use. In some scenarios, NPV peaked at levels that were operationally infeasible due to life-of-mine (LOM) limitations. Across all approaches, plant expansion to four times the base capacity was only viable under optimistic cost and price assumptions, while a more conservative view supports expansion up to three times capacity. The methodology also allowed the use of variable cut-off grades to determine the optimal plant size, balancing financial returns and operational feasibility.

INTRODUCTION

Strategic mine planning has evolved significantly over the past decades, becoming a key component in maximising value in mining projects, particularly in integrated systems involving multiple deposits and shared infrastructure (Newman *et al,* 2010; Dimitrakopoulos, 2011). In this context, open pit optimisation and cut-off grade strategies play a fundamental role in defining economically viable extraction limits and production schedules (Whittle, 1989; Lane, 2016). More recently, the integration of mining and processing decisions has been recognised as essential for improving resource utilisation and overall system performance, particularly under uncertainty and operational constraints.

The integration of multiple mining operations within a shared infrastructure has therefore become an increasingly relevant approach in strategic mine planning, especially in the context of capital-intensive copper projects. In such systems, coordinating production from different sources enables improved resource utilisation, enhanced plant performance, and increased overall project value.

In multi-source mining hubs, decisions related to processing capacity, production scheduling, and capital allocation are highly interdependent. Optimising these decisions requires a comprehensive framework capable of evaluating trade-offs between economic performance and operational feasibility. In particular, the definition of optimal plant capacity and the sequencing of deposits play a critical role in ensuring sustainable long-term value generation.

The definition of optimal processing capacity is a key challenge in strategic mine planning, as production rates directly affect both operational costs and overall economic performance. Previous studies have shown that production systems typically exhibit diminishing returns beyond an optimal

capacity region, highlighting the importance of balancing throughput and cost efficiency when defining expansion strategies (Souza *et al,* 2019).

Within integrated mining systems, the simultaneous optimisation of multiple components—including open pit and underground operations, processing decisions, and blending strategies—has been shown to significantly enhance project value when compared to traditional sequential approaches (Montiel, Dimitrakopoulos and Kawahata, 2016). These integrated approaches are particularly relevant when accounting for geological uncertainty, as they enable the generation of production schedules that both maximise net present value (NPV) and reduce associated risks.

This study investigates the integrated development of three mining sources, comprising an open pit operation and a second deposit combining open pit and underground mining methods. The analysis focuses on assessing the synergies associated with shared processing infrastructure, while evaluating alternative expansion strategies under varying economic and operational constraints.

Within this context, defining optimal processing capacity and production sequencing becomes a complex problem, requiring the evaluation of trade-offs between capital investment, operating flexibility, and long-term value creation (Newman *et al,* 2010; Lane, 2016). Special emphasis is placed on how plant capacity expansion, project scheduling, and cut-off grade optimisation interact within an integrated system. The objective is to identify configurations that maximise NPV while maintaining operational continuity and practical feasibility over the life-of-mine (LOM).

MATERIALS AND METHODS

The study was developed based on an integrated mine planning framework considering three mining sources sharing a centralised processing plant and tailings storage facility. The methodology was structured in sequential stages, allowing the progressive evaluation of the system from its theoretical maximum potential to more realistic operational conditions.

The first stage consisted of the optimisation of mining shapes for each deposit. For the open pit operations, pit optimisation was carried out using varying economic parameters in order to identify optimal pit limits under different economic assumptions. For the underground operation, stope optimisation was performed following a similar approach, also incorporating variations in economic parameters to define the optimal extraction shapes. This analysis ensured that both open pit and underground resources were evaluated consistently under a range of possible economic conditions.

In the second phase, integrated production sequencing was developed for all deposits without capital constraints. The objective of this step was to evaluate the full potential of the mining hub by maximising value through unrestricted scheduling. This approach allowed the identification of the theoretical maximum NPV and provided a benchmark for subsequent constrained scenarios.

Following this unconstrained analysis, operational and external constraints were progressively introduced into the model. In particular, restrictions on project start dates were incorporated to reflect realistic conditions related to licensing timelines, development requirements, and operational readiness. This step enabled the assessment of the impact of scheduling limitations on the overall system performance.

In the final stage, processing plant feed was varied by incrementally increasing the amount of material processed. This expansion was associated with additional capital expenditure (CAPEX), representing the cost of increasing plant capacity. By linking processing capacity increments to corresponding CAPEX investments, it was possible to evaluate the economic trade-offs between increased throughput and capital intensity. This approach allowed the identification of the optimal hub configuration, balancing value generation, operational feasibility, and investment constraints.

In addition to the base analysis, dynamic cut-off grade optimisation was applied across all mining sources to support ore blending strategies and improve plant utilisation. Sensitivity analyses were also conducted to evaluate the robustness of the results under variations in CAPEX, operating costs, and commodity price assumptions.

Scenario definition

Three main scheduling approaches were evaluated, including delayed integration (10-year lag), moderate delay (5-year lag), and an unconstrained start scenario with underground development lag.

For each approach, scenarios were constructed varying:

- Processing plant capacity (base case and incremental expansions up to 3–4× capacity).
- Expansion CAPEX intensity (US$30–50 million per 1 Mt/a increment).
- Operating cost assumptions.
- Commodity price scenarios.

Optimisation and planning tools

A strategic mine planning tool was used to generate production schedules and evaluate economic performance. The optimisation process incorporated key operational and economic constraints, including open pit and underground sequencing requirements, processing capacity limits, and overall mining movement restrictions. Project performance was assessed through discounted cash flow analysis, using NPV at a 9 per cent discount rate as the primary economic indicator.

Sensitivity analysis

Sensitivity analyses were conducted to evaluate the impact of key economic variables, including CAPEX associated with plant expansion, mining and processing operating costs (OPEX), and commodity price assumptions. These analyses enabled the validation of results under a range of economic conditions, providing insight into the robustness of the proposed strategies and their sensitivity to different risk factors.

RESULTS AND DISCUSSION

The results demonstrate that the integration of multiple mining sources within a single processing system leads to a significant increase in economic value when compared to stand-alone project configurations. This improvement is primarily driven by enhanced plant utilisation and the ability to optimise ore blending across different deposits, which is consistent with previous studies on integrated mine planning and resource allocation.

The evaluation of plant capacity expansion scenarios reveals a clear trend in the relationship between processing capacity and economic performance, as illustrated in Figure 1. Increasing plant capacity results in higher NPVs across all scheduling strategies; however, the marginal economic gains tend to decrease beyond 36 Mt/a, particularly for delayed scenarios.

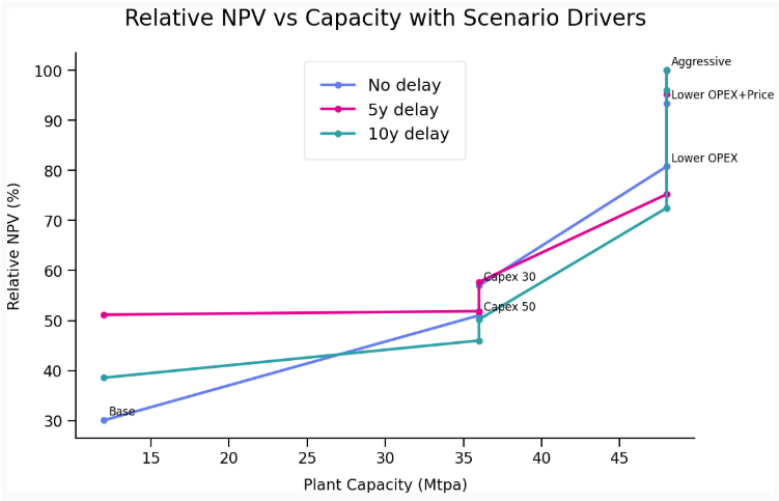


FIG 1 – NPV as a function of plant capacity for different scheduling strategies, including unconstrained and delayed project start scenarios.

As shown in Figure 1, multiple scenarios are represented at the same plant capacity due to variations in economic assumptions. Differences in capital expenditure, operating costs, and commodity prices lead to significant variations in NPV, even for identical processing capacities. In particular, reduced OPEX and higher price assumptions result in substantial value increases, highlighting the importance of economic parameters in capacity expansion decisions. These results demonstrate that plant capacity alone does not determine project value. Instead, the interaction between capacity, cost structure, and market conditions plays a fundamental role in defining the optimal configuration of the mining system.

The incremental contribution of different economic and operational drivers to project value is illustrated in Figure 2. Starting from the base scenario at 12 Mt/a, successive improvements are introduced, including plant capacity expansion, reductions in operating costs, and increases in commodity prices. The most significant increase in NPV is associated with the expansion to 36 Mt/a, which establishes the structural basis for value creation. Additional gains are subsequently achieved through improvements in economic parameters, particularly through reduced operating costs and higher commodity prices. Figure 2 demonstrates that the full potential of the mining system is not achieved through capacity expansion alone, but rather through the combined effect of technical and economic optimisations. This highlights the importance of integrated decision-making in strategic mine planning. The integrated scenarios consistently demonstrated higher NPVs compared to stand-alone project evaluations. The observed increase can be attributed to improved plant utilisation, enhanced ore blending and grade control practices, and more effective allocation of high-value resources.

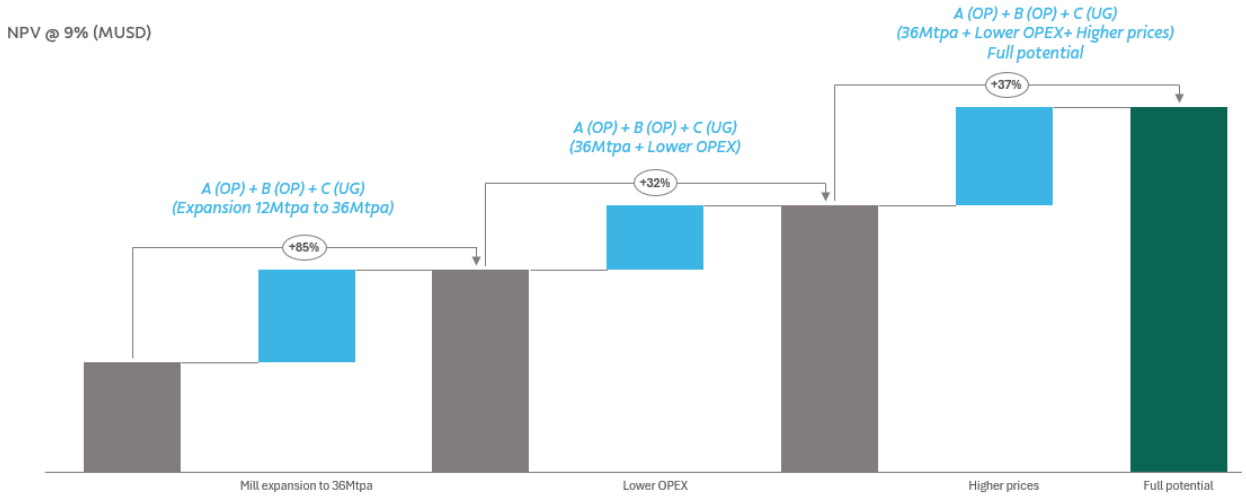


FIG 2 – Incremental NPV contributions from plant capacity expansion, operating cost reductions, and commodity price variations.

Under conservative cost and commodity price assumptions, expansion up to approximately three times the base capacity (around 36 Mt/a) provides a robust and well-balanced solution. This configuration enables sustained plant utilisation throughout a large portion of the LOM, as shown in Figure 3, while maintaining operational feasibility.

In contrast, scenarios involving expansion to four times the base capacity (approximately 48 Mt/a) are only economically viable under more optimistic assumptions, such as reduced operating costs and higher commodity prices. Although these scenarios may yield higher NPVs, they frequently present important operational challenges. In several cases, the production schedules required to sustain higher throughput result in shorter periods of full plant utilisation, increased mining rates, and reduced continuity in the LOM profile.

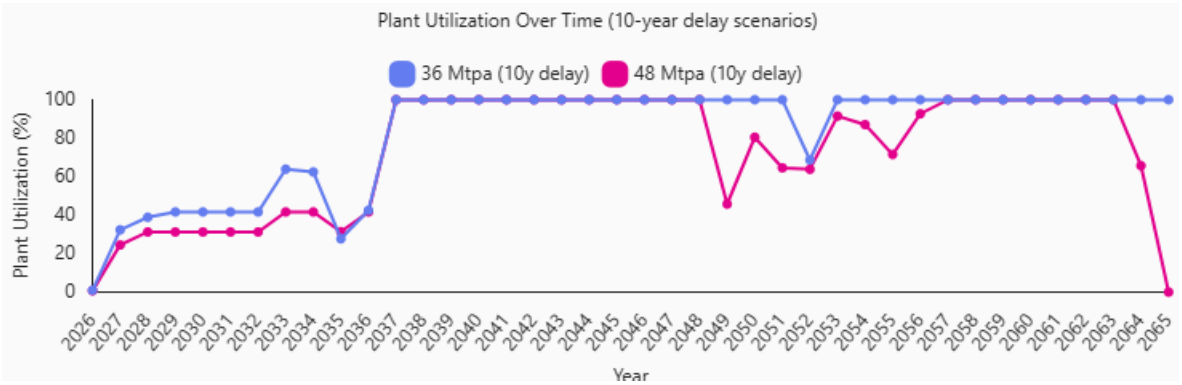


FIG 3 – Plant utilisation over the LOM for selected capacity scenarios, highlighting differences in operational sustainability.

The results indicate that expansion is robust up to approximately three times the base capacity under conservative economic assumptions; however, expansion to four times the base capacity becomes viable only under optimistic conditions characterised by lower costs and higher commodity prices.

Although higher capacities yield higher NPVs, such scenarios often present operational limitations, including reduced LOM continuity and constraints in mining sequence execution. These effects can be further observed in the production profiles, where the integration of sources and sequencing constraints directly impact the sustainability of high-capacity scenarios.

The analysis of scheduling strategies further indicates that variations in project start dates, such as delays of five or ten years, have a limited impact on overall project value. This suggests that the system retains a certain degree of flexibility to accommodate external constraints, including licensing timelines and capital allocation requirements. However, earlier integration of mining sources generally contributes to improved plant utilisation and more efficient system performance.

The results also highlight the importance of evaluating operational feasibility alongside economic performance. While some scenarios achieve higher NPVs, they are associated with constraints that limit their practical implementation. These include insufficient periods of full plant utilisation, excessively high mining rates, sequencing constraints, and elevated capital intensity. As such, the identification of an optimal solution requires a balanced consideration of both financial and operational criteria, in line with integrated optimisation approaches that consider multiple decision variables simultaneously (Montiel, Dimitrakopoulos and Kawahata, 2016). In contrast, scenarios involving moderate expansion demonstrated more balanced and operationally feasible outcomes. These alternatives were characterised by more stable production profiles, improved alignment between mining and processing capacities, and a consequent reduction in overall operational risk.

Finally, the application of dynamic cut-off grade optimisation proved to be a key element in maximising value within the integrated system. By enabling flexible ore allocation and blending between mining sources, this approach supports improved control of plant feed grade and allows for more efficient use of available capacity. Consequently, cut-off grade optimisation plays a fundamental role in defining the optimal configuration of the mining hub.

CONCLUSIONS

The results of this study confirm that integrated mine planning represents a powerful approach for maximising value in multi-deposit mining systems. By enabling the coordinated development of multiple sources within a shared processing infrastructure, the methodology allows for improved resource utilisation and increased economic performance.

From a strategic perspective, the evaluation of plant capacity expansion indicates that increasing processing capacity up to approximately three times the base case provides the most robust solution. This configuration achieves a favourable balance between economic performance and operational feasibility.

Although further expansion may lead to additional increases in NPV, such scenarios typically rely on optimistic economic conditions and are associated with higher levels of risk. In particular, these

configurations may compromise operational continuity and require more aggressive mining schedules, which can limit their practical implementation.

The study also demonstrates that flexibility in project scheduling, including variations in start dates, has a relatively limited impact on overall value. This provides an opportunity to adapt development strategies to external constraints without significantly compromising economic performance.

Overall, the integration of optimisation techniques, scenario analysis, and dynamic cut-off grade strategies provides a robust framework for supporting strategic decision-making in complex mining systems. The methodology presented in this work can be applied to similar multi-source operations to identify optimal configurations under uncertainty, balancing financial returns with operational constraints.

ACKNOWLEDGEMENTS

Support for this work was provided by Vale Base Metals, including the data and technical environment necessary for its development. The contributions of all colleagues involved in the project are gratefully acknowledged.

REFERENCES

Dimitrakopoulos, R, 2011. Stochastic optimization for strategic mine planning: A decade of developments, *J Min Sci,* 47:138–150.

Lane, K F, 2016. *The Economic Definition of Ore*: *Cut-Off Grades in Theory and Practice,* COMET Strategy Pty Ltd.

Montiel, L, Dimitrakopoulos, R and Kawahata, K, 2016. Globally optimising open-pit and underground mining operations under geological uncertainty, *Mining Technology*, 125(1):2–14.

Newman, A M, Rubio, E, Caro, R, Weintraub, A and Eurek, K, 2010. A review of operations research in mine planning, *Interfaces*, 40(3):222–245.

Souza, F R, Câmara, T R, Torres, V F N, Nader, B and Galery, R, 2019. Optimum mine production rate based on price uncertainty, *REM Int Eng J Ouro Preto*, 72(4):625–634.

Whittle, J, 1989. The facts and fallacies of open pit optimization, pp 1–7 (Whittle Programming Pty Ltd: North Balwyn).

Strategic optimisation of the Voisey's Bay mining and processing value chain

D Galloway[1], R Mallett[2], B Gadelha[3], L Montiel[4], T Câmara[5] and J Todd[6]

1. Principal Engineer – Strategic Planning, Vale, Sudbury ON P0M 1N0, Canada.
   Email: derek.galloway@vale.com
2. Principal Engineer – Strategic Planning, Vale, Sudbury ON P0M 1N0, Canada.
   Email: robert.mallett@vale.com
3. Manager – Strategic Planning, Vale, Sudbury ON P0M 1N0, Canada.
   Email: bruno.gadelha@vale.com
4. Principal Engineer – Strategic Planning, Vale, Toronto ON M5J 2K2, Canada.
   Email: luis.petro@vale.com
5. Principal Engineer – Strategic Planning, Vale, Belo Horizonte MG 34006-270, Brazil.
   Email: tais.camara@vale.com
6. Head – Strategic Mine Planning, Vale, Toronto ON M5J 2K2, Canada.
   Email: jody.todd@vale.com

ABSTRACT

Vale's Newfoundland and Labrador portfolio comprises the Voisey's Bay (VB) mining complex, concentrator, and the Long Harbour Processing Plant (LHPP). VB is a major nickel–copper–cobalt operation with three active mines (two underground and one open pit) and multiple proximal and regional resource opportunities that could be developed through mine expansions or new operations. All ore mined at VB is processed on-site, with concentrates feeding both external markets and LHPP for refining into finished nickel, cobalt, and copper products.

Following depletion of the Ovoid open pit, the remaining mines no longer generate sufficient concentrate to operate LHPP at nameplate capacity, necessitating the purchase of external feed and materially reducing portfolio margins. The strong interdependencies across the mining, milling, and processing assets, combined with a large inventory of future development options, highlight the need for integrated, enterprise-level optimisation of the entire value chain (Whittle, 2010; Goodfellow and Dimitrakopoulos, 2017).

This paper presents the application of advanced mathematical optimisation to develop a comprehensive strategic model encompassing all current and potential mining areas at VB, operational constraints, external feed obligations, capital requirements, and downstream value-chain linkages. The model was calibrated to Vale's official Newfoundland and Labrador financial model and used to optimise both the existing base-case mine plan and a long-term Resource Development Plan.

Optimisation of the base-case plan identified significant opportunities to increase value through improved mining sequences, debottlenecking constrained processes, and targeted capital deployment. These changes enable VB to fully supply LHPP from internal feed for an extended period, significantly defer reliance on external feed and increase portfolio net present value by approximately 90 per cent. Extension of the optimisation to include all future resource opportunities demonstrates the potential to extend the life of the VB complex by more than 20 years and deliver a further approximately 70 per cent, representing a total uplift of approximately 160 per cent.

This case study demonstrates how enterprise-level optimisation and integrated value-chain planning translate strategic intent into executable, value-accretive development pathways across a complex, vertically integrated mining operation.

INTRODUCTION

Vale's Newfoundland and Labrador portfolio is a vertically integrated nickel-copper-cobalt operation in which mine production, concentrator throughput and downstream processing performance are tightly linked (Whittle, 2010; Newman *et al,* 2010). Following depletion of the world-class Ovoid open pit, the operation entered a structurally different phase in which the remaining mines no longer

generate sufficient internal feed to fill the capacity of the downstream processing at optimal operating levels. This reduced portfolio production and increased dependence on external nickel concentrate to sustain downstream utilisation.

In this operating context, the planning problem extends beyond conventional mine planning and scheduling. Decisions relating to mining sequence, cut-off, throughput and expansion timing influence not only mine-level cash flow, but also concentrator loading, finished nickel production and the overall value generated across the entire portfolio. As a result, planning the operation on a mine-by-mine basis is insufficient to identify the highest-value outcome for the system and requires consideration of the mining and processing value chain (Whittle, 2010; Goodfellow and Dimitrakopoulos, 2017).

At the same time, the Voisey's Bay (VB) district contains a large inventory of future development opportunities including extensions to existing operations and additional regional mineralisation identified through exploration and resource evaluation activities (Vale Base Metals, 2025). These include extensions of the current mines, nearby opportunities that may leverage existing infrastructure, and broader regional opportunities that could materially extend the life of the operation. The challenge is therefore not simply to identify which resources add value, but to determine how those resources should be staged so that value is maximised across the full mining and processing chain.

This paper applies integrated mathematical optimisation to the VB mining and processing value chain to evaluate that challenge. The primary objective was to quantify the value unlocked by optimising the existing base plan and by allowing coordinated increases in mine and concentrator throughput. A secondary objective was to assess how future inventory should be staged, progressing from the current operating base to proximal endowment and then to the broader regional endowment. The intent was to demonstrate how integrated value-chain optimisation can support both immediate planning decisions and longer-term strategic development pathways.

ASSET AND VALUE CHAIN OVERVIEW

The current operating base at VB comprises the Reid Brook, Eastern Deeps and Discovery Hill mines, feeding a concentrator with a nominal base throughput of 2.5 Mt/a. The concentrator produces copper concentrate for direct sale and nickel concentrate for downstream processing at the Long Harbour Processing Plant (LHPP), which produces finished nickel, copper and cobalt and has a maximum finished nickel production capacity of approximately 50 kt/a (Vale Base Metals, 2025). In the base plan, Long Harbour is underfilled in all years, requiring the purchase of external nickel concentrate at much lower margins to fill the available processing capacity. The principal material flows and processing linkages represented in the optimisation model are summarised in Figure 1.

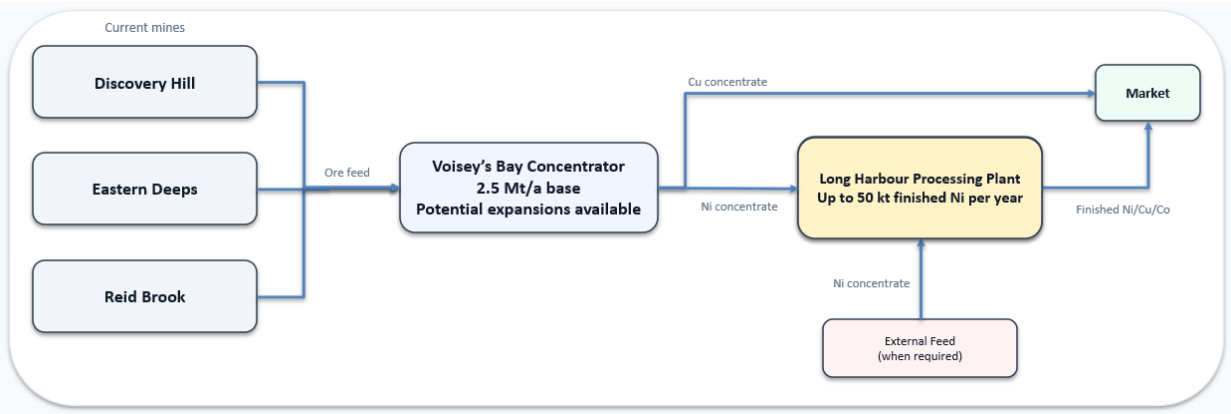


FIG 1 – Simplified Voisey's Bay and Long Harbour mining and processing value chain.

The economic behaviour of the system is controlled by the interaction of upstream and downstream capacities. For each of the individual mines at VB, production is constrained by material handling capacity, lateral and vertical development, and the timing of mine access and infrastructure

completion. At concentrator level, throughput is limited by plant capacity unless expansion capital is introduced. At processing level, finished nickel production is capped by the capacity of Long Harbour, meaning that value is not created simply by maximising ore tonnes or contained nickel in isolation. Instead, value depends on delivering the right combination of tonnes and grade at the right time so that the downstream plant is supplied effectively without exceeding its limit.

This creates a strongly coupled planning system. A mine plan for an individual mine at VB that appears optimal in isolation may still be suboptimal at portfolio level if it leads to an incorrect combination of material feeding the concentrator, fails to support Long Harbour, or introduces expansion capital at the wrong time. Conversely, a plan that sustains downstream utilisation for longer may create substantially more value even where it requires different sequencing or different cut-off and throughput choices than a conventional plan for any of the individual mines.

Future opportunities can also be distinguished by their relationship to the existing system. Proximal endowment includes lateral and down-plunge extensions of the current mines, together with nearby potential satellite orebodies that can utilise existing infrastructure. Full regional endowment includes a broader set of opportunities that would require stand-alone infrastructure or step-change expansions of the current mines. These differences matter because the value of future inventory is shaped not only by size and grade, but also by proximity, capital intensity and the timing with which each opportunity can be introduced into the value chain.

To evaluate these trade-offs consistently, a mathematical optimisation framework was developed to represent current operations, future inventory options, capacity constraints and capital timing across the full value chain.

METHODOLOGY

Inventory development

The first stage of the work was to develop a consistent inventory basis for both the current operation and future development opportunities.

The base-case inventory comprised Reid Brook, Eastern Deeps and Discovery Hill. Mining method assumptions were aligned with the current operating basis, with open stoping at Reid Brook and Eastern Deeps and open pit mining at Discovery Hill. All endowment opportunities were evaluated individually, and mining methods considered appropriate for strategic assessment were selected for each.

Mineable inventories for all endowment opportunities were then generated using all available geological information and industry standard stope optimisation software consistent with accepted underground mine planning practice (Alford and Hall, 2009). The purpose of this work was not to produce fully engineered project schedules for every future ore source, but to convert geological potential into mineable inventory representations that could be evaluated consistently within the strategic optimisation framework.

Once mineable inventories had been generated, internal automated tools were used to produce strategic-level mine designs and schedules for each endowment opportunity. These strategic schedules were then used to estimate capital requirements, establish indicative development timing, and generate potential production profiles suitable for long-range portfolio optimisation. This provided a consistent basis for comparing different endowment opportunities, even where detailed engineering and project definition had not yet been completed.

Cut-off was defined on a Net Smelter Return (NSR) basis and varied by mining area, consistent with established cut-off optimisation theory and NSR-based underground inventory evaluation practices (Lane, 1988; Alford and Hall, 2009). This was an important feature of the study because cut-off affected not only mineable tonnes and grade, but also concentrator head grade, nickel concentrate production and the ability to support Long Harbour at preferred operating levels. In this system, cut-off is therefore both a mine planning variable and a value-chain variable. Lower cut-offs increased inventory and supported higher throughput, while higher cut-offs improved grade but could reduce the volume of feed available to sustain downstream utilisation. The relative net present value (NPV) of each alternative depended on how it interacted with the full mining and processing chain.

For strategic comparison, future inventory was grouped into two categories beyond the base case. Proximal endowment consisted of extensions of the current mines both laterally and at depth, as well as nearby potential orebodies able to leverage existing infrastructure. Full regional endowment consisted of a broader set of potential orebodies that would require either new stand-alone infrastructure, such as a shaft or decline, or step-change expansions of the current mines.

This classification allowed the study to distinguish between inventory that could reasonably be brought forward as a medium-term extension of the current operation and inventory that should be treated as longer-term strategic optionality.

Integrated optimisation model

A mathematical optimisation model was developed to represent the VB mining and processing value chain over annual time steps. The model maximised pre-tax NPV of the full portfolio using a 9 per cent discount rate, consistent with accepted strategic mine planning optimisation approaches (Newman *et al,* 2010).

The optimisation framework included the following major constraints and decision variables:

- Annual production capacity of each mine.
- Capital investment required to increase mine capacity.
- Lateral and vertical development.
- Mine ventilation and infrastructure requirements.
- Material handling constraints.
- Concentrator throughput limitations.
- Capital investment required to increase concentrator throughput.
- Downstream processing capacity.
- Start dates for new projects.
- Construction dependencies.
- External nickel concentrate requirement when internal feed was insufficient to support downstream utilisation.

This formulation allowed the model to evaluate the combined effect of mining sequence, cut-off, throughput growth, and project timing on portfolio value. The intent was not simply to generate a mine plan with the highest local value at each operation, but to identify the sequence of decisions that produced the best overall result for the integrated system, reflecting enterprise optimisation and mining-complex optimisation principles (Whittle, 2010; Goodfellow and Dimitrakopoulos, 2017).

The model also treated capital as an active strategic variable. Capital was estimated for mine expansions, concentrator expansions and future deposit development. Capital had to be incurred before any associated production increase or access to new mining areas could occur, and all capital was represented on an annual basis. This allowed the optimisation to determine not only whether an expansion should occur, but also the timing and scale at which it should occur.

This was particularly important for throughput growth. Concentrator expansion can unlock significant value, but only if enough mine inventory exists to utilise the higher throughput and only if the timing of expansion aligns with the availability of feed. Similarly, future ore sources only create value when introduced at a point in the schedule where they strengthen the overall system rather than simply displacing already adequate internal supply.

Model calibration

The optimisation model was calibrated against the official financial model for the portfolio to ensure that the base economic logic of the optimisation model aligned with the business planning basis.

Calibration was undertaken by running the approved base schedule with all tasks and material flows locked in time and comparing the resulting outputs against the reference financial model. The calibration process considered annual cash flow, overall NPV and annual production numbers. Once the model reproduced the reference case to an acceptable strategic planning standard, the schedule was progressively unlocked so that sequence, throughput expansion and future inventory options could be tested.

This approach provided confidence that relative NPV uplifts identified by the optimisation model reflected genuine strategic opportunities rather than artefacts of model structure. It also ensured that results from different scenario families could be compared within a common economic framework.

Scenario definition

Five scenario families were evaluated.

Scenario 1 – Base case locked

The approved base plan was run without any changes to schedule, throughput limits or inventory basis. This case established the baseline against which all subsequent value uplift was measured.

Scenario 2 – Base case optimised

The same base inventory was retained, but activities were rescheduled while still adhering to all existing schedule dependencies, mine production capacities, and mill throughput limits. This case isolated the value created by sequencing improvements alone.

Scenario 3 – Base case optimised with increased throughput

The base inventory was retained, but both mine production and concentrator throughput were allowed to increase where justified by capital investment. This case measured the combined value of improved sequencing and coordinated throughput growth.

Scenario 4 – Proximal endowment

The base plan was expanded to include proximal endowment, representing approximately an 80 per cent increase in total available inventory relative to the base case. The schedule was optimised and throughput growth was permitted where justified by capital investment. This case tested whether relatively accessible future inventory was value accretive and could provide an executable bridge to unlocking the full regional endowment.

Scenario 5 – Full regional endowment

The base plan was expanded by the addition of the full regional endowment, representing approximately a 200 per cent increase in total inventory relative to the base case. Scheduling, throughput increase and capital timing were all optimised. Start timing for each of the regional opportunities was assumed based on their depth and distance from existing infrastructure, consistent with the strategic nature of the case and the likely need for further study and infrastructure development.

This scenario structure allowed the value contribution of sequencing, throughput growth, proximal additions and full endowment development to be separated and compared clearly.

RESULTS

Base-case optimisation

The first important result is that substantial value existed in the current operating plan before any throughput increases or future endowment was introduced. Optimisation of the base schedule without changing throughput or inventory generated a relative NPV uplift of approximately 15 per cent. This showed that a meaningful portion of value leakage was caused by suboptimal sequencing alone.

When the model was then allowed to optimise the base plan while also increasing mine and concentrator throughput through capital investment, the relative uplift increased to approximately 90 per cent. This established two important points. First, sequencing matters. Second, throughput matters even more when it is increased in a coordinated way across the mines and the concentrator.

The improvement was driven by a combination of better sequencing of available mineable inventory, more effective alignment between mine production and concentrator throughput, stronger support for downstream finished nickel production, and better capital timing for expansions required to unlock higher rates.

Before optimisation, Long Harbour was underfilled in all years based on the approved planning case and available reserve inventory. After optimisation of the base case, downstream utilisation improved materially and the need for external nickel concentrate was significantly deferred. However, the plant still became underfilled from approximately 2035 onwards because the current inventory base alone was insufficient to sustain full production over the long-term.

One of the clearest strategic insights from the base-case results is that increasing throughput adds value only up to the point where internal supply becomes the limiting factor. Once the highest-value internal inventory is depleted, further throughput increase ceases to add material value because the system is then constrained by the underlying resource base rather than by plant capacity. In practical terms, throughput stopped adding meaningful value once Long Harbour capacity was fully utilised, which occurred at varying production rates, depending on the case.

The optimisation also provided clear mine planning guidance within the existing operations. In both Reid Brook and Eastern Deeps, higher-priority areas were those containing a combination of stronger grades and larger tonnage, and therefore more total nickel per unit of development. These areas were preferentially scheduled because they contributed more effectively to sustaining concentrator feed quality and downstream finished nickel production. By contrast, areas with the lowest total nickel per unit of development were deprioritised as they contributed less value within the integrated system and were less effective at supporting Long Harbour utilisation. This outcome reinforces that the highest-value plan is not necessarily the one that maximises local mine extraction at every point in time, but the one that best allocates available inventory across the full value chain.

The base-case optimisation therefore provides a strong planning message: major value can be unlocked before new resources are developed, but the full value of higher throughput can only be realised if the right mining areas are prioritised and if future inventory is introduced on a timely basis. The resulting reprioritisation of mining areas at Discovery Hill and Reid Brook is illustrated in Figure 2, while the corresponding reprioritisation at Eastern Deeps is shown in Figure 3.

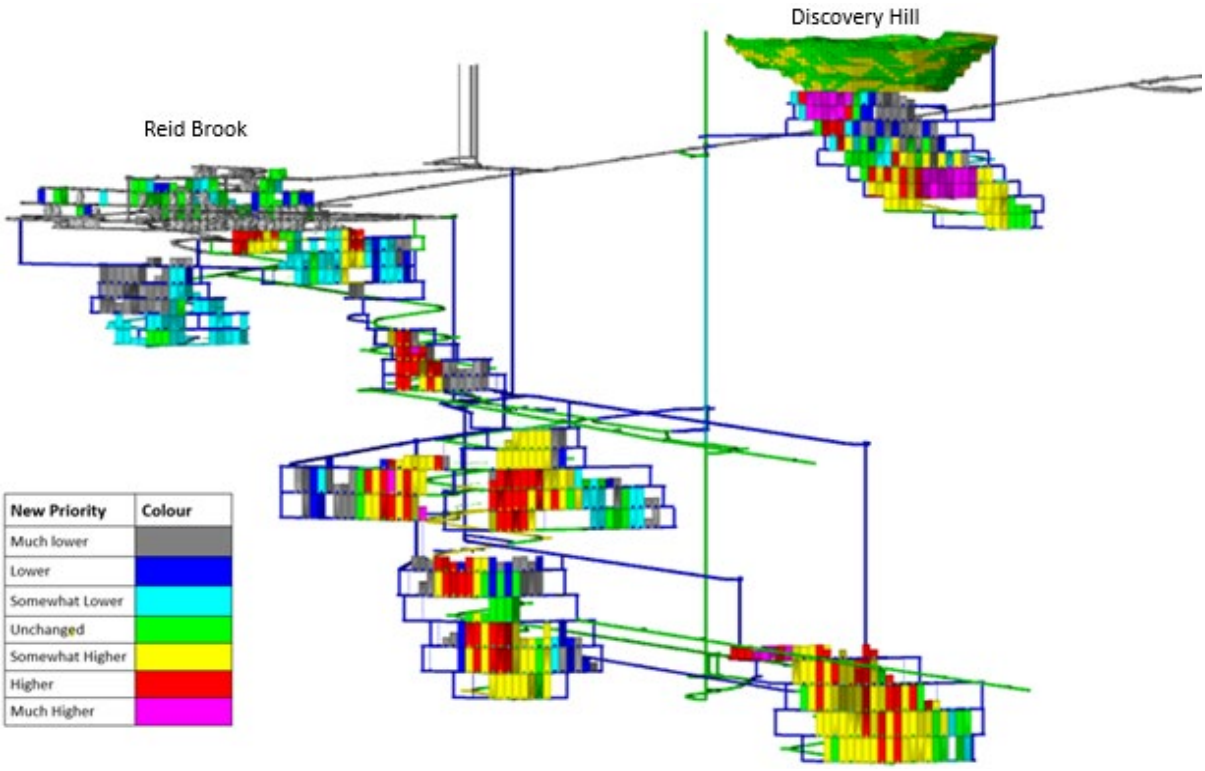


FIG 2 – Discovery Hill and Reid Brook reprioritisation based on optimisation results.

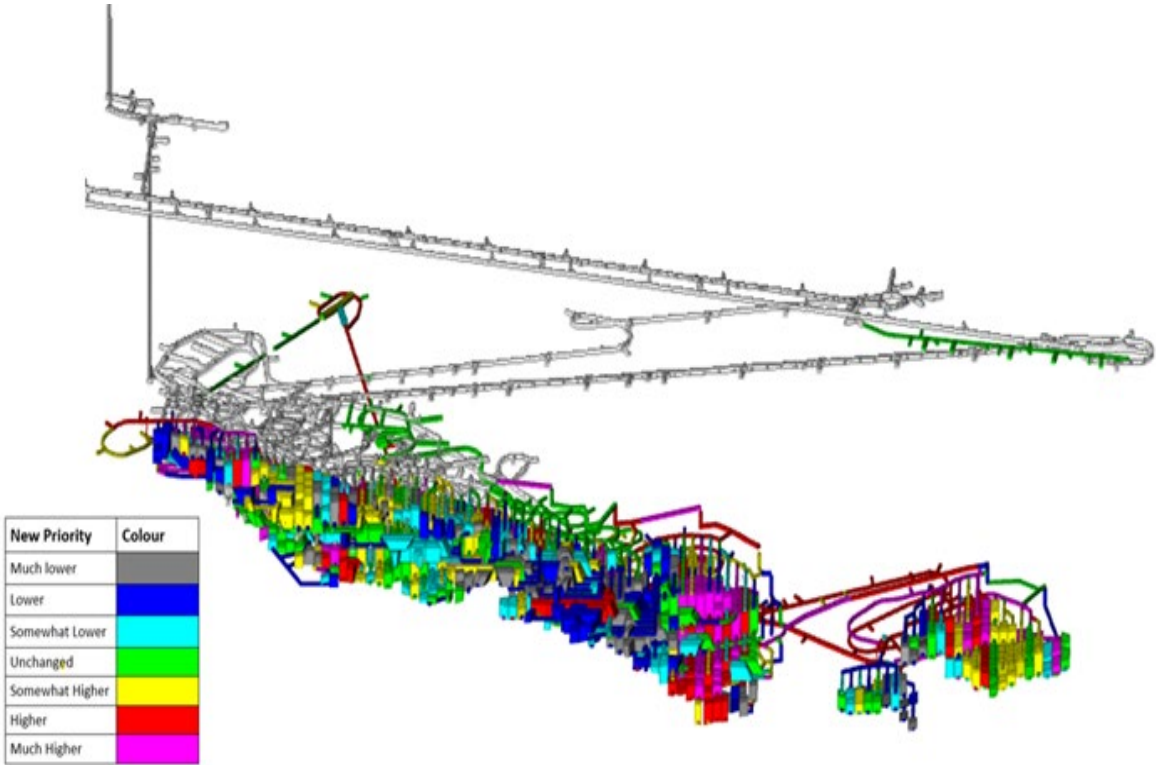


FIG 3 – Eastern Deeps reprioritisation based on optimisation results.

Proximal endowment

The proximal endowment scenario tested whether extensions of the current mines and nearby orebodies that could leverage existing infrastructure would add value and provide a practical bridge between the current plan and longer-term regional development.

These opportunities are strategically distinct from the broader regional endowment because they are close to the current operating footprint, require less new infrastructure, and can therefore be considered at lower capital intensity and with shorter lead times.

When proximal endowment was added, the relative NPV uplift increased to approximately 120 per cent when compared to the base case. This indicates that the value of proximal inventory lies not only in the addition of tonnes, but in the timing of those tonnes within the value chain.

The main contribution of proximal endowment was to sustain Long Harbour with internal feed beyond the point at which the optimised base case began to weaken. Full finished nickel production could be maintained until the mid to late 2040s, extending downstream self-sufficiency by approximately 10 years relative to the base case pathway. In parallel, the need for purchasing external nickel concentrate was deferred until a similar time frame.

The proximal case therefore acts as a medium-term bridge. It materially improves portfolio value, extends the useful life of current infrastructure, and postpones the need for major regional capital. For strategic mine planning, this is a significant outcome because it demonstrates that full value does not need to be pursued in one step. A staged pathway exists in which lower-risk and more executable proximal opportunities are brought forward first, creating a better platform for later, larger developments.

Full regional endowment

The full regional endowment scenario tested the maximum strategic value that could be created by introducing the broader regional inventory into the optimised portfolio.

This case expanded the available inventory by approximately 200 per cent relative to the base case and treated future opportunities as strategic development projects requiring longer lead times and more extensive infrastructure. Earliest project start timing for these opportunities was evaluated on a case-by-case basis.

When full regional endowment was introduced, relative NPV uplift increased to approximately 160 per cent. This was the highest-value case considered and confirms that the broader district endowment has the potential to materially improve the long-term economics of the operation.

The principal advantage of this case is duration. Whereas the proximal case fills Long Harbour using internal feed until the mid- to late-2040s, the full regional case extends internal-feed support significantly longer. In effect, the broader endowment transforms the operation from a portfolio that would eventually underfill the downstream plant even after optimisation into one capable of sustaining downstream production over a much longer time horizon.

Capital timing was central to this case. The optimisation identified a staged capital pathway in which:

- Near-term capital supports concentrator expansion.
- Mid-term capital supports proximal endowment development.
- Long-term capital supports regional endowment development.

Each of these phases serves a different strategic purpose. Concentrator expansion enables the system to make better use of current and future mine supply. Proximal capital supports the medium-term bridge. Regional capital extends overall mine life and long-term downstream self-sufficiency.

The full regional case should therefore be viewed as a strategic plan to optimise value of the portfolio over the entire life of the region. It is the highest-value pathway, but its success depends on preserving the sequencing logic established by the base and proximal cases rather than attempting to shortcut directly to a district-scale development outcome. The relative NPV uplifts generated by the five scenarios are summarised in Figure 4.

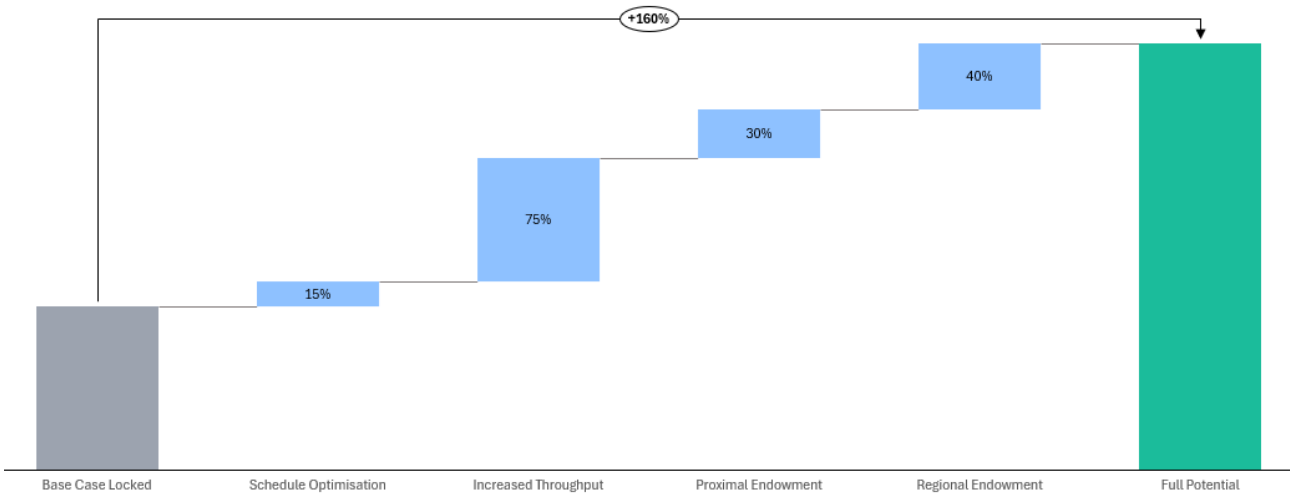


FIG 4 – Relative NPV uplift results by scenario.

DISCUSSION

Taken together, the scenario results define a practical sequence for maximising value across the VB mining and processing portfolio, consistent with enterprise optimisation concepts (Whittle, 2010). The first step is to unlock value within the existing operation through improved sequencing and prioritisation of the current inventory. The second is to increase mine and concentrator throughput where the required capital investment is justified and where sufficient internal feed exists to support Long Harbour. The third is to introduce proximal endowment as a medium-term bridge that extends downstream self-sufficiency and delays the need for major new infrastructure. The final step is to develop the broader regional endowment, which maximises long-term value and extends the life of the integrated operation significantly. This sequence is important because it aligns value creation with execution maturity rather than treating all future opportunities as if they should be pursued at the same time. The resulting staged development pathway is summarised in Figure 5.

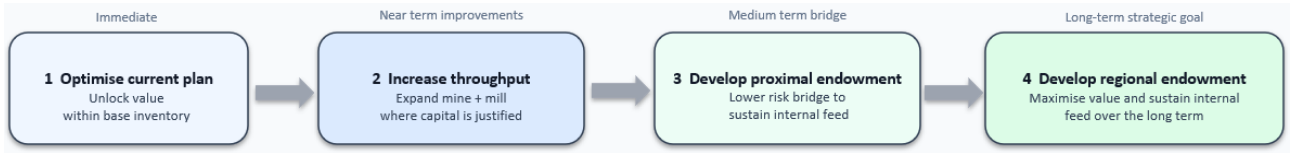


FIG 5 – Practical staged development pathway identified through integrated optimisation.

This case study demonstrates that integrated value-chain optimisation can materially alter strategic mine planning decisions in a vertically linked mining and processing portfolio (Goodfellow and Dimitrakopoulos, 2017). At VB, the dominant value driver is not the performance of any individual mine in isolation, but the extent to which the system can sustain enough internal nickel units to keep Long Harbour fully utilised over time. This has important implications for long-range planning, because it means that decisions on cut-off, mining sequence, throughput and expansion timing must be evaluated in terms of their effect on the full value chain rather than on local mine economics alone (Lane, 1988; Whittle, 2010; Goodfellow and Dimitrakopoulos, 2017).

One of the most important findings is that throughput growth is a major source of value, but only when combined with the right sequencing and future inventory support. Expansion of mine and concentrator throughput in isolation does not capture the full benefit, and can even destroy value, if the portfolio later becomes constrained by depletion of high-value mineable inventory. Conversely, sequence optimisation alone can unlock material value even without throughput growth, demonstrating that the original operating schedule contained meaningful unrealised upside. The study therefore shows that throughput should be treated as a conditional value lever rather than an automatic growth objective.

The work also provides practical planning insight into how the existing mine base should be managed. Within Reid Brook and Eastern Deeps, areas containing the strongest combination of

grade and tonnage, and therefore the greatest total nickel per unit of development, were consistently prioritised by the optimisation. Lower-grade and/or lower-tonnage areas contributed less value in the integrated system and were therefore deprioritised. This reinforces a broader planning lesson: in a constrained downstream system, the highest-value mine plan is not necessarily the one that extracts the greatest local value from each mining area independently, but the one that best allocates scarce tonnes, grade and development effort across the portfolio.

The distinction between proximal and regional endowment is also strategically important. Proximal endowment does not simply add more resources to the inventory. Its value lies in its ability to be introduced earlier, with lower capital intensity and shorter lead times, to bridge the point at which the base-case operation would otherwise begin to weaken. By extending full internal-feed support to Long Harbour into the mid to late 2040s, proximal endowment improves the medium-term outlook of the portfolio and delays the need for larger regional capital. The broader regional endowment then becomes the long-term strategic phase, extending internal-feed support significantly and maximising value across the full life of the region.

In terms of maturity, concentrator expansion is the highest-confidence and most immediately executable component of the pathway. The endowment cases remain strategic in nature. Their purpose is to define a robust and value-accretive path forward rather than to present definitive project designs. Even so, the study influenced practical planning decisions, including mine sequencing, resource prioritisation, cut-off thinking, throughput assumptions, study pipeline development and capital prioritisation. This indicates that enterprise-level optimisation can function not only as a strategy-screening tool, but also as a decision-making framework that shapes real planning outcomes.

Sensitivity work on nickel price, capital escalation, discount rate and inventory uncertainty confirmed that the overall value pathway remained robust across a range of tested conditions. This increases confidence that the preferred development sequence is structurally sound rather than dependent on a single narrow planning assumption. In that sense, the principal contribution of the work is not just the identification of high-value scenarios, but the demonstration that strategic mine planning in a vertically integrated operation can be materially improved by treating mining, milling and downstream processing as one optimisation problem.

CONCLUSIONS

Integrated optimisation of the VB mining and processing value chain identified substantial value uplift across all scenarios and confirmed that mining, concentrator throughput and downstream processing must be evaluated as a single strategic planning problem (Whittle, 2010; Goodfellow and Dimitrakopoulos, 2017).

Optimisation of the current plan alone created material value, showing that sequencing is a significant source of unrealised upside within the existing mining base.

The largest uplift within the base inventory was achieved through coordinated increases in mine and concentrator throughput. Value increased until Long Harbour capacity was effectively filled, which occurred at varying production rates depending on the scenario. Within the existing operation, higher-grade and higher-tonnage areas in Reid Brook and Eastern Deeps were prioritised, while lower-grade and/or lower-tonnage areas were deprioritised.

Proximal endowment provided a lower-risk medium-term bridge, extending internal-feed support to Long Harbour into the mid- to late-2040s and delaying the need for major regional capital.

Full regional endowment maximised long-term portfolio value and extended internal-feed support significantly, defining the long-term strategic opportunity for the operation.

Overall, the preferred planning strategy is staged: optimise the current operation, expand throughput where justified, develop proximal endowment as the executable bridge, and then introduce regional endowment to maximise long-term value.

ACKNOWLEDGEMENTS

The authors acknowledge the contributions of the Long-Range Planning team and the Voisey's Bay and Long Harbour Processing Plant site teams. The authors also acknowledge Vale for approval to publish this work.

REFERENCES

Alford, C and Hall, B E, 2009. Stope optimisation tools for selection of optimum cut-off grade in underground mine design, in *Proceedings of the Project Evaluation Conference 2009*, pp 137–144 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Goodfellow, R and Dimitrakopoulos, R, 2017. Simultaneous stochastic optimization of mining complexes and mineral value chains, *Mathematical Geosciences*, 49(3):341–360.

Lane, K F, 1988. *The Economic Definition of Ore: Cut-Off Grades in Theory and Practice* (Mining Journal Books Ltd: London).

Newman, A M, Rubio, E, Caro, R, Weintraub, A and Eurek, K, 2010. A review of operations research in mine planning, *Interfaces*, 40(3):222–245.

Vale Base Metals, 2025. Voisey's Bay Operations, Newfoundland & Labrador, Canada: Technical Report Summary [online]. Report current as at 31 December 2025. Available from: <https://valebasemetals.com/wp-content/uploads/2026/04/Voiseys-TRS-2026-V12-website-1.pdf> [Accessed: 1 May 2026].

Whittle, G, 2010. Enterprise Optimisation, in *Proceedings of the Mine Planning and Equipment Selection Conference 2010*, pp 105–117 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Assessing strategic production schedules of mining complexes and project evaluation impacts relating to uncertainty

W Patton[1] and B Marriott[2]

1. Principal Engineer Mine Planning, BHP, Perth WA 6000. Email: will.patton@bhp.com
2. Lead Project Studies, BHP, Perth WA 6000. Email: blake.c.marriott@bhp.com

INTRODUCTION

Mining project valuation is fundamentally dependent on the ability to extract, blend and deliver ore products that consistently meet market specifications over the life-of-mine. In iron ore mining complexes, where multiple open pits feed integrated processing and logistics systems, production outcomes are governed by complex interactions between geological variability, operational constraints and market conditions. Suboptimal sequencing decisions can lead to degradation in product quality, inefficient utilisation of capacity and ultimately reduced project value.

Conventional long-term mine planning approaches rely on deterministic orebody models and sequential optimisation workflows. These methods simplify the spatial variability and uncertainty inherent in mineral deposits and typically optimise components of the mining value chain in isolation. Consequently, the resulting production schedules often fail to represent the true distribution of material properties and are exposed to supply risks arising from geological uncertainty. In addition, the treatment of commodity prices through fixed assumptions or subjective scenarios introduces further exposure to demand uncertainty, limiting the robustness of project evaluations.

Recent developments in stochastic optimisation enable simultaneous consideration of multiple equiprobable realisations of geological and market conditions. By integrating these uncertainties directly into the optimisation process, it is possible to generate production schedules that better reflect the range of outcomes likely to be encountered during mining and improve the reliability of valuation metrics.

This study examines the impact of incorporating both geological and commodity price uncertainty into the optimisation of a Pilbara iron ore mining complex. Production schedules generated using a deterministic optimisation workflow are compared against those obtained from a simultaneous stochastic optimisation framework, with a focus on value, risk and operational practicality. A mining sequence and cut-off grade strategy that is practical in execution achieved a net present value approximately 20 per cent greater than that of the deterministic sequence with significant up-side potential was generated using the stochastic framework.

METHODOLOGY

The case study considers a mining complex comprising two open cut mines supplying a centralised processing and export system producing a single direct shipping fines product stream, with no lump production. Materials extracted from each mine are blended to satisfy product specifications prior to shipment. These specifications include minimum iron grades that vary by period, ranging from Fe ≥ 60.0 per cent to Fe ≥ 61.4 per cent. Silica targets also vary by period from SiO_2 ≤ 4.5 per cent to SiO_2 ≤ 4.0 per cent whilst other impurities alumina and phosphorous are fixed at Al_2O_3 ≤ 2.25 per cent and P ≤ 0.12 per cent for all periods. Adherence to ore quality targets is important considering that value-in-use (VIU) penalties (Honeyands and Jelenich, 2011) on revenues are built into the optimisation model objective function.

The mining complex is designed to operate at a nominal production rate of approximately 20 Mt/a, constrained by processing throughput, mining fleet capacity, and downstream logistics limits (Figure 1). Material characteristics are defined by a set of ore types, including supergene martite-goethite (M-G) ores of the Brockman Iron Formation (Harmsworth *et al*, 1990) and minor weathered redeposition products (Canga) of limited economic potential (Kneeshaw and Morris, 2014). Spatial variability in grade and metallurgical properties captured through multiple geostatistical simulations representing equiprobable realisations of the orebody represent supply uncertainty (Goovaerts,

1997). Construction of these ore deposit models followed the methods outlined in a previous publication (Patton, 2023).

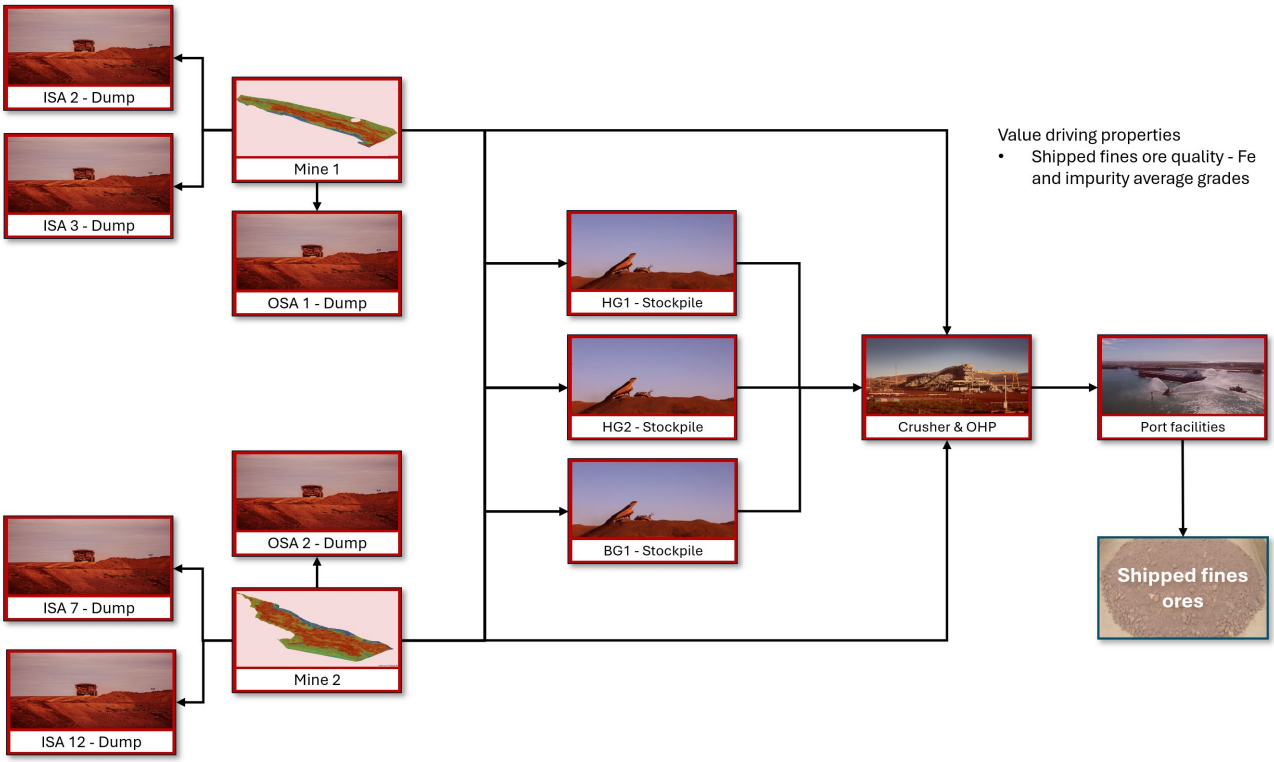


FIG 1 – The iron ore mining complex consisting of two open pit mines. Each mine has two in-pit (ISA) and one external (OSA) waste dumps and three stockpiles for various quality ores before blending occurs at the crusher and ore handling facility. From the ore handling facility, the crushed iron ore fines travels via rail to a port facility for ship loading and export.

Commodity price uncertainty is incorporated using stochastic simulations of iron ore fines prices derived from a mean reverting process (Ornstein, 1930) calibrated to historical monthly spot price data (MarketIndex, 2025) between July 2008 and June 2025. Key parameters were estimated following Dixit and Pindyck (1994) and include the mean reversion speed (η = 0.0918), long-run mean price in US dollars per dry metric tonnage unit ($\bar{S}$ = 102.66 US$/dmtu), volatility ($\sigma$ = 0.1653) and initial price (S_0 = 94.17 US$/dmtu). Monte Carlo simulation is used to generate multiple (100) equiprobable price paths representing uncertainty in future market conditions (Suarez and Fernandez, 2009). Five price forecast realisations representing the 10th, 25th, 50th, 75th and 90th percentiles of the total aggregated iron ore price over the forecast period were drawn from the set (Figure 2) prior to upscaling each realisation to annual periods by averaging for use in the optimisation framework.

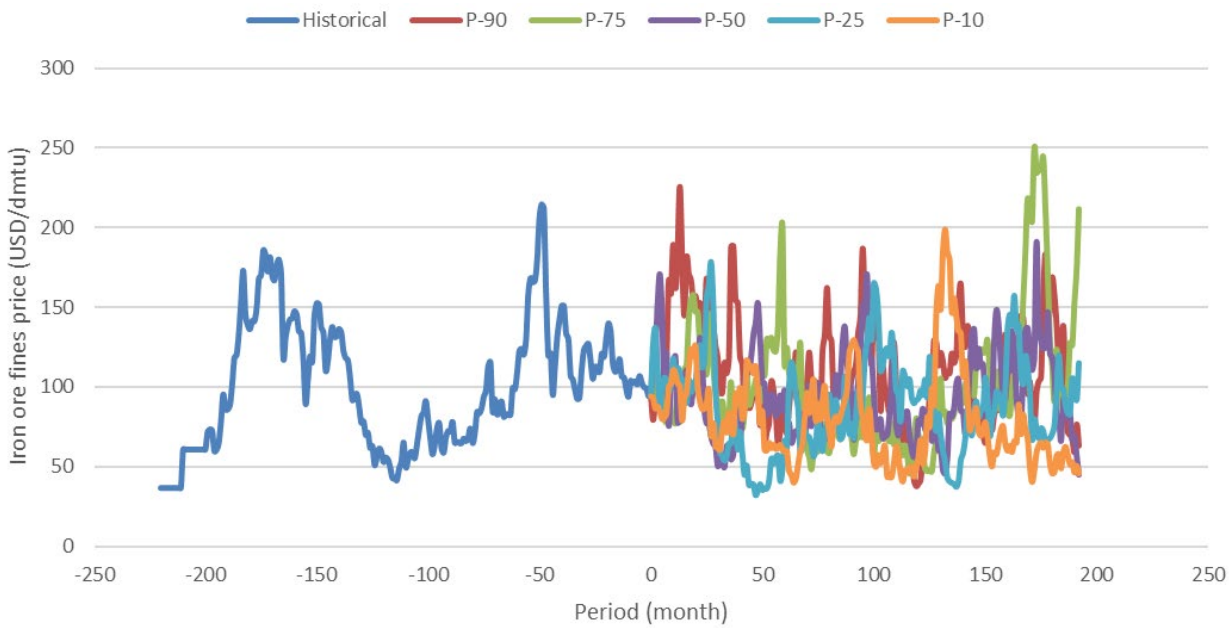


FIG 2 – Historical monthly iron ore fines spot price data between July 2008 and June 2025 and a representative subset of five price forecasts using the stochastic mean reverting process model.

The optimisation framework integrates these uncertainties within a simultaneous stochastic formulation (Goodfellow and Dimitrakopoulos, 2017; Saliba and Dimitrakopoulos, 2019). In total, 125 scenarios of the mining complex are evaluated when market price and geological uncertainties are combined within this framework. Mining extraction sequences, destination policies and processing decisions are optimised collectively to maximise project value while respecting operational, processing and product quality constraints. Hypothetical opening capital costs have been included to allow some assessment of supply and demand risk on the payback period metric. A risk-analysis of the sequence generated using the company's in-house solution (Stone *et al*, 2004), following a traditional deterministic workflow based on estimated orebody inputs is also presented.

RESULTS

The deterministic and stochastic optimisation approaches produce materially different production schedules and project outcomes (Figure 3). The sequence generated using the deterministic workflow exhibits more sensitivity to geological variability when evaluated across stochastic realisations, with increased risk of deviations from ore quality (Fe, SiO_2 and Al_2O_3) targets. These deviations in product grades allow the operation to meet production targets for the crusher and ore handling plant (mass) in most periods. On the contrary, in the stochastic sequence, production targets for mass are slightly more flexible to ensure ore quality targets are achieved in all periods.

From an economic perspective, the stochastic sequence achieves a significantly higher net present value (NPV). The uplift in NPV relative to the deterministic sequence is approximately 20 per cent, with significant reductions in downside risk (P_{10}) of 90 per cent over the deterministic base case. This is largely achieved through maximising returns in early periods by minimising mining costs relating to the stochastic sequences ability to meet production targets whilst mining considerably less mass. Reduced product quality penalties in periods 4 through to 11 achieved by the stochastic sequence also support improved revenues. The deterministic sequence also demonstrates significant exposure to commodity price uncertainty and results in risk that the project may not recoup opening capital cost until period 5. In contrast, the stochastic optimisation framework produces a production schedule that is likely to achieve this by period 4.

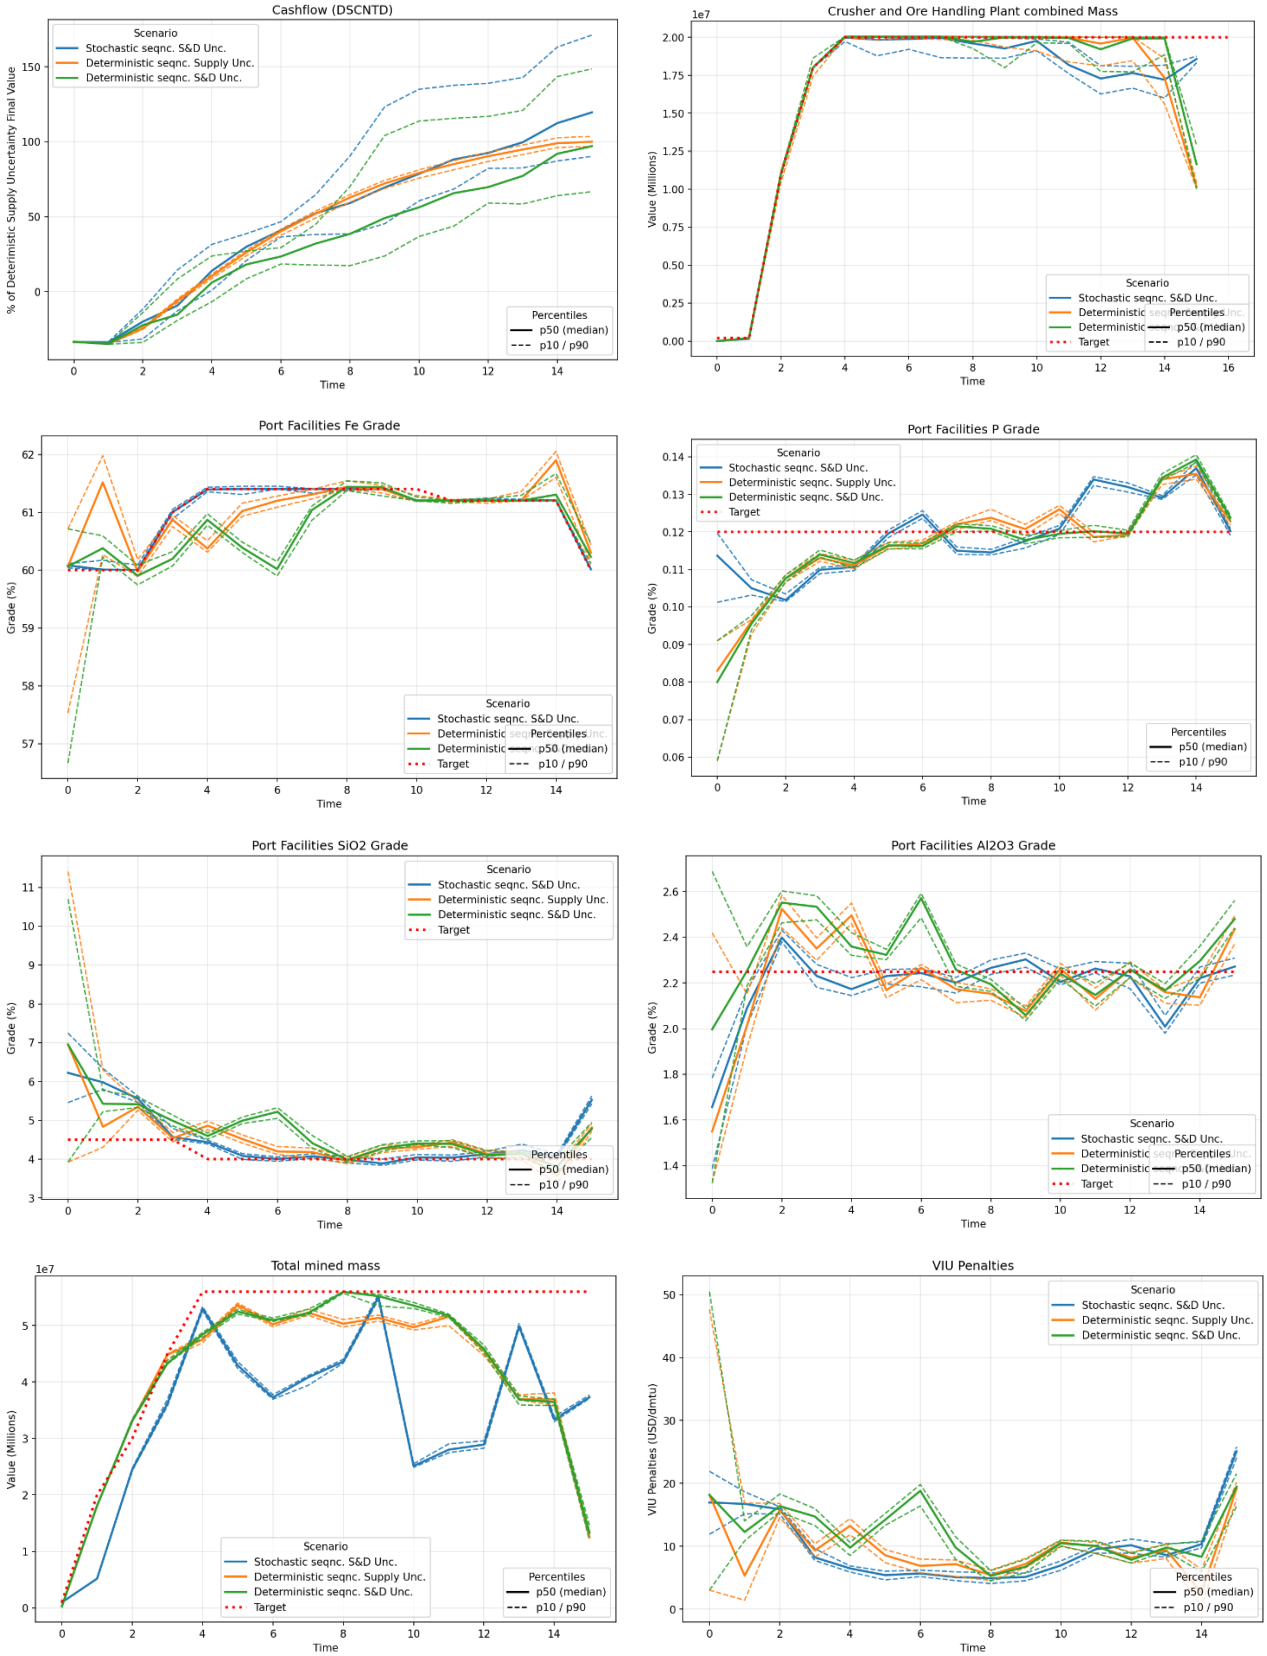


FIG 3 – Comparison plots for key metrics comparing the median (solid lines) from the deterministic sequence with supply uncertainty (orange), the deterministic sequence with supply and demand uncertainty (green) and of the stochastic solution (blue). The dotted lines represent the 10th and 90th percentiles showing the range of possible outcomes. Red dotted lines represent production targets determined from the deterministic sequence.

The deterministic and stochastic mining sequences are compared in Figure 4. The two methods apply mining constraints differently. The deterministic sequence advances in horizontal layers or panels, whereas the stochastic approach produces more irregular extraction shapes. This reflects the underlying optimisation methods. The in-house solution aggregates blocks into larger panels and sequences those panels, whilst the stochastic framework determines the optimal sequence block-by-block. Higher risk zones are visibly deferred in the stochastic sequence, represented by dark-red and orange zones mined in later periods to minimise operational impact. Minor artefacts can be observed in both mining sequences. Thin elongated strips in a north–south orientation of questionable mineability can be observed in the deterministic sequence. Whereas isolated volumes being mined in different periods to other neighbouring blocks can be observed in the stochastic sequence.

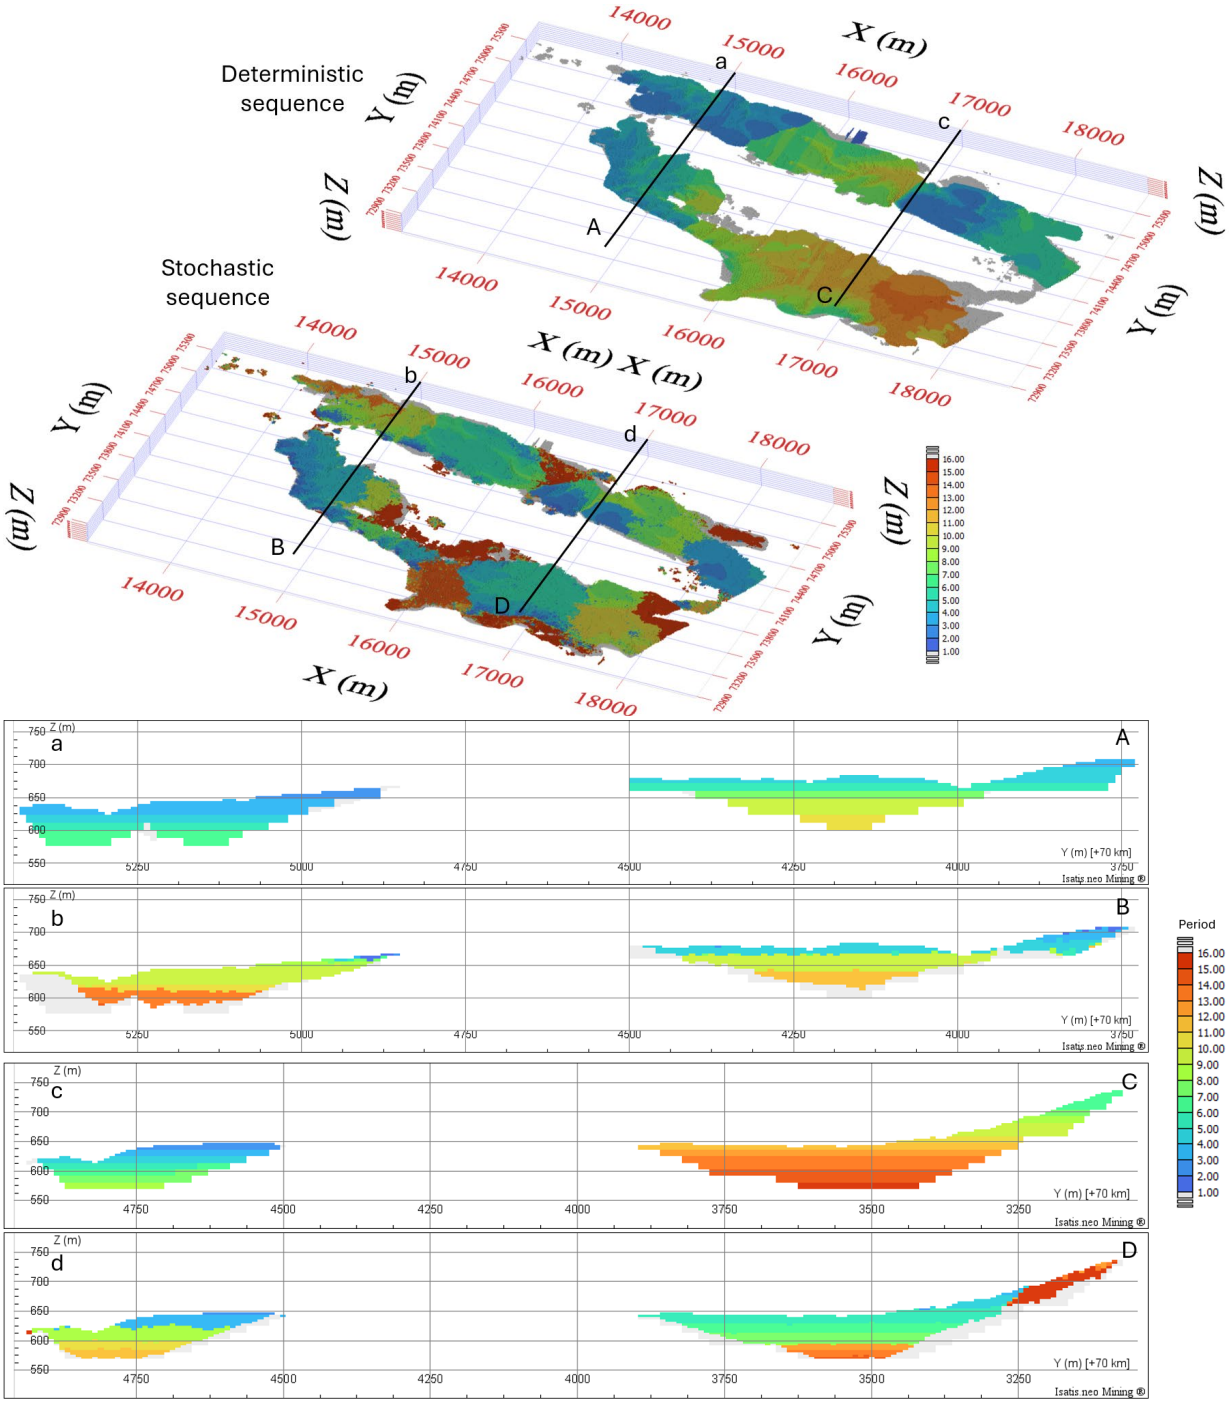


FIG 4 – Oblique view and cross-sections of the extraction sequences generated using deterministic and stochastic workflows coloured by mined period. Blocks not mined in the featured sequence (ie stochastic) but are mined in the alternate sequence (deterministic) are coloured grey.

The cut-off strategy required to achieve schedules for each scenario recalled from Figure 3 are presented in Figure 5. In all scenarios cut-off grades on Fe grade bins sent to ore destinations are determined by the objective function and vary by pit and from period to period. The deterministic sequence requires a complex cut-off strategy for each pit whereby Fe cut-off grades change greatly between periods to deal with variability in the grades of the ore deposit to ensure mass targets are met. The cut-off strategy for the stochastic sequence is more stable ranging from 57.0 per cent to 58.0 per cent Fe for the North pit and 56.8 per cent to 58.0 per cent Fe for the South pit with a limited amount of long-term strategic stockpiling required to buffer against variability in the ore deposit.

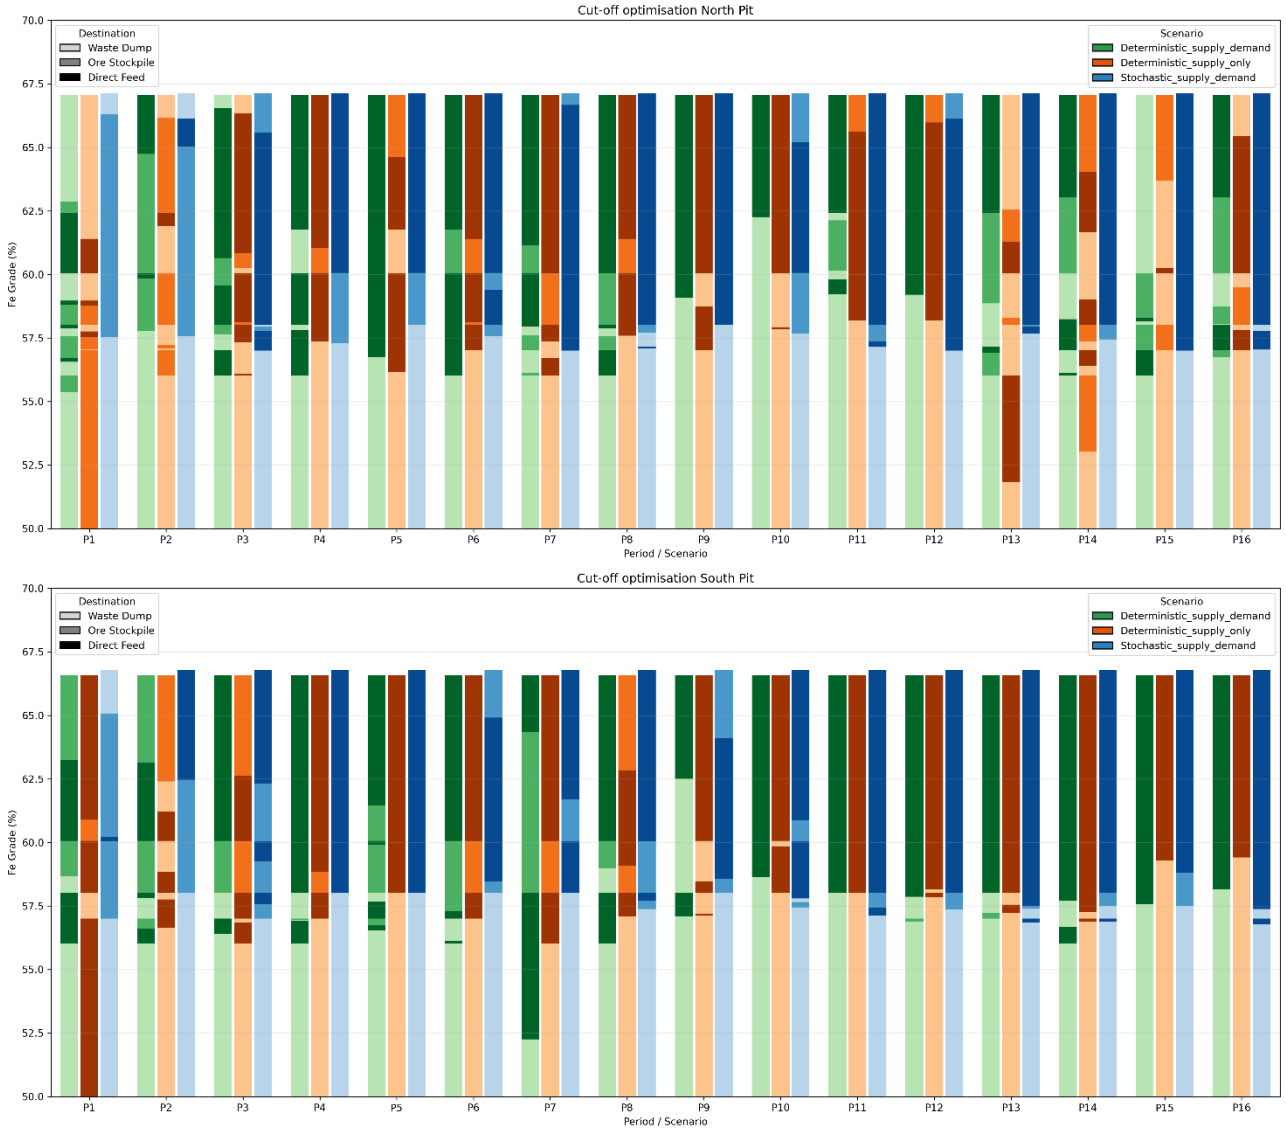


FIG 5 – Comparison of cut-off strategy for North (top) and South (bottom) pits from risk analysis of the deterministic sequence under supply only (orange), supply and demand (green) risk and the stochastic sequence generated under both supply and demand risk (blue).

CONCLUSION

The results demonstrate the importance of incorporating geological and commodity price uncertainty into strategic mine planning optimisation. It was revealed that commodity price uncertainty had a more significant impact on the viability of the project than that resulting from uncertainty in the layout and variability of grades in the ore deposits. Without properly accounting for this risk, forecast revenues would be substantially overestimated for much of the mine life following the traditional deterministic framework.

Whilst the stochastic optimisation framework produces production schedules that are higher in value and operationally realistic at a strategic level, further mine planning work is required to produce a

comprehensive sequence of phases that incorporates complexities that may introduce suboptimal factors such as haulage infrastructure, staged regulatory approval or construction timelines. This work can be done using the stochastic sequence as a starting point and completed using many industry standard mine design software. The impacts of any additional ore-loss or dilution attributed to modifications to the stochastic sequence would reduce the presented NPV uplift (+20 per cent) and should be evaluated in a similar manner used to risk assess the deterministic sequence demonstrated in this case study.

The deterministic sequence was evaluated assuming a business strategic decision to maximise ore mass processed through the crusher and ore handling plant. Alternately, a decision to maximise ore quality was employed for the stochastic sequence. How sensitive the project is to these different operating modes would require further investigation using the stochastic framework. This is the focus of ongoing research and may even higher value sequences than presented here.

Deterministic workflows generate sequences that are lower in value and exposed to significant risk when evaluated against realistic representations of the orebody and market environment. The integration of stochastic orebody modelling with probabilistic commodity price simulation provides a more robust basis for project evaluation in iron ore mining systems. This information can in turn enable project teams and leaders to make informed risk-based strategic business decisions.

ACKNOWLEDGEMENTS

Access to the data set and permission to publish was provided by BHP Pty Ltd Western Australia Iron Ore. The authors would like to acknowledge the BHP Pty Ltd Western Australia Iron Ore Strategic Mine Planning department and its leadership team for supporting the publication of this article. The authors would also like to thank Mr Gabriel Gariba for technical advice and support throughout the study.

REFERENCES

Dixit, A K and Pindyck, R S, 1994. *Investment under uncertainty* (Princeton University Press).

Goodfellow, R and Dimitrakopoulos, R, 2017. Simultaneous Stochastic Optimization of Mining Complexes and Mineral Value Chains, *Mathematical Geosciences*, 49(3):341–360.

Goovaerts, P, 1997. *Geostatistics for Natural Resources Evaluation* (Oxford University Press).

Harmsworth, R A, Kneeshaw, M, Morris, R, Robinson, C J and Shrivastava, P K, 1990. BIF derived iron ores of the Hamersley Province, in *Geology of the Mineral Deposits of Australia and Papua New Guinea*, pp 617–642 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Honeyands, T and Jelenich, L, 2011. Calculating the value of iron ores in ironmaking and steelmaking, in *Proceedings of the Iron Ore Conference 2011*, pp 431–436 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Kneeshaw, M and Morris, R C, 2014.The Cenozoic detrital iron deposits of the Hamersley Province, Western Australia, *Australian Journal of Earth Sciences*, 61(4):513–586.

MarketIndex, 2025. Iron Ore Price, MarketIndex.com.au, [Online]. Available from: <https://www.marketindex.com.au/iron-ore> [Accessed: 23 August 2025].

Ornstein, L S, 1930. On the theory of the Brownian motion, *Physical Review*, 36:823–841.

Patton, W, 2023. Risk assessment of iron mineral resources using conditional simulations, in *Proceedings of the Mineral Resource Estimation Conference*, pp 203–230 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Saliba, Z and Dimitrakopoulos, R, 2019. Simultaneous stochastic optimization of an open pit gold mining complex with supply and market uncertainty, *Mining Technology*: *Transactions of the Institute of Mining and Metallurgy*, 128(4):216–229.

Stone, P, Froyland, G, Menabde, M, Law, B, Rasyar, R and Monkhouse, P, 2004. Blasor – Blended iron-ore mine planning optimisation at Yandi, in *Proceedings of the International Symposium on Orebody Modelling and Strategic Mine Planning*, pp 285–288 (AIMM).

Suarez, M M and Fernandez, P L, 2009. Commodities prices modeling using Gaussian Poisson exponential stochastic processes, A practical implementation in the case of copper, in *Proceedings of the 23rd European Conference on Modelling and Simulation, ECMS 2009*, pp 433–438.

Production planning under geological uncertainty

Optimisation of selective mining unit (SMU) size through misclassification and economic analysis using sequential Gaussian simulation (SGS) – application to a copper deposit

N Bouhlali[1], A Elghali[2], A Moradi[3], Y Taha[4] and M Benzaazoua[5]

1. PhD Student, Geology and Sustainable Mining Institute (GSMI), Mohammed VI Polytechnic University, Ben Guerir 43150, Morocco. Email: noaman.bouhlali@um6p.ma
2. Assistant Professor, Geology and Sustainable Mining Institute (GSMI), Mohammed VI Polytechnic University, Ben Guerir 43150, Morocco. Email: abdellatif.elghali@um6p.ma
3. Assistant Professor, Department of Mining Engineering, University of Kentucky, Lexington KY 40506, USA. Email: ali.moradi@uky.edu
4. Associate Professor, Geology and Sustainable Mining Institute (GSMI), Mohammed VI Polytechnic University, Ben Guerir 43150, Morocco. Email: yassine.taha@um6p.ma
5. Full Professor, Geology and Sustainable Mining Institute (GSMI), Mohammed VI Polytechnic University, Ben Guerir 43150, Morocco. Email: mostafa.benzaazoua@um6p.ma

INTRODUCTION

The selective mining unit (SMU) is the smallest discretised volume of rock for which extraction decisions are made, linking geological block models to production planning (Vasylchuk and Deutsch, 2018). Its dimensions govern the trade-off between ore recovery, dilution and economic performance. Simulation-based, economically weighted classification of ore and waste is well established in grade-control and mine-planning practice (Dimitrakopoulos and Godoy, 2014; Rossi and Deutsch, 2014), and simulation has also been used to investigate SMU size (Leuangthong, Neufeld and Deutsch, 2003). This study therefore does not introduce a new classification or simulation method. Its contribution is applying this established approach specifically to SMU-size selection by quantifying the relationship between misclassification risk and economic outcomes in a porphyry copper deposit.

METHODOLOGY

The classification and economic-translation procedure follow established methods for probability-weighted, economically defined ore/waste classification from conditional simulation (Dimitrakopoulos and Godoy, 2014; Rossi and Deutsch, 2014); it is applied here to selective mining unit (SMU)-size selection using the following workflow, illustrated for a porphyry copper case study. The study is conducted at an open pit porphyry copper operation with chalcopyrite-bornite mineralisation hosted in quartz monzonite porphyry, with sericite (SER) and unaltered (UNA) alteration domains. Two complementary data sets support the analysis: 6178 exploration drill hole composites at 2 m intervals, and 55 119 production blasthole assays on a 5 × 5 m grid covering mined-out benches from 2450 m to 2555 m elevation. The blasthole assays are interpolated using inverse distance weighting (IDW) at 5 × 5 × 15 m resolution to construct a grade control model used as the base case for comparison against the sequential Gaussian simulation (SGS)-derived candidate SMUs. Because the base case and candidate SMUs differ in data density, estimation method and block support, the reported trends reflect the combined influence of these factors rather than the isolated effect of SMU size. A pure block-size comparison would require a common high-resolution reference model and consistent estimation conditions across all candidates supports. Figure 1 summarises the methodological workflow.

Contact analysis between the SER and UNA domains reveals a sharp grade discontinuity (median SER/UNA grade ratio of 4.5) and supports the use of hard boundaries (Krige, 1999). Declustering is performed using the cell-centre nearest-neighbour method on a 50 × 50 × 15 m grid (de Carvalho, Dimitrakopoulos and Minniakhmetov, 2019). Omnidirectional variograms are modelled with nested spherical structures: for SER, nugget 0.19 with ranges of 33 m and 180 m; for UNA, nugget 0.28 with ranges of 35 m and 227 m. One hundred SGS realisations are generated per domain (Manchuk and Deutsch, 2012) and aggregated into candidate SMU sizes from 5 × 5 × 15 m to 65 × 65 × 15 m, with bench height held constant at 15 m to examine changes in horizontal support while keeping the vertical mining-block dimension constant.

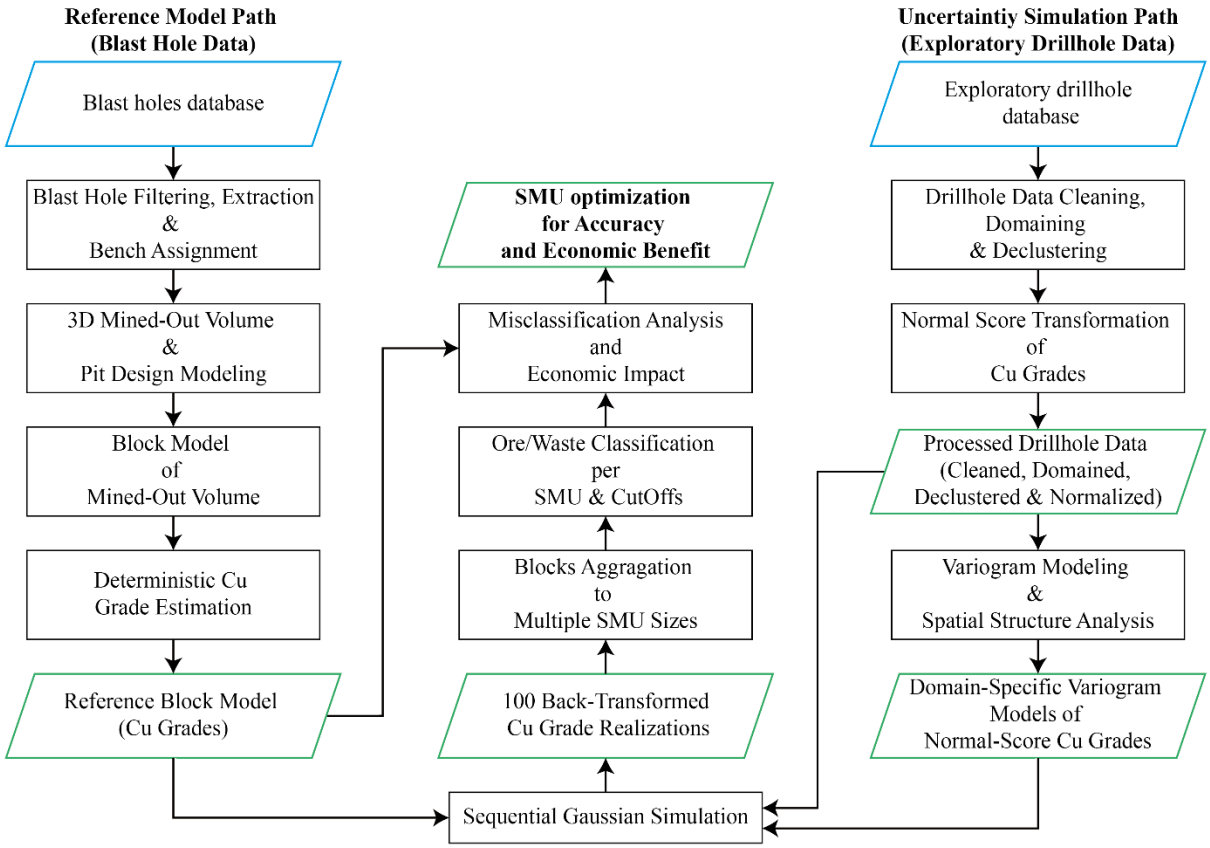


FIG 1 – Repeatable workflow for SMU-size classification and economic evaluation.

Each candidate SMU in each realisation is classified relative to the corresponding base-case grade-control support. The 5 × 5 × 15 m base-case cells are aggregated using volume-weighted average grades to match each candidate SMU before classification. For a cut-off grade (0.15, 0.25 and 0.35 per cent), correct ore is defined when both the base case and the simulated SMU are classified as ore; correct waste is defined when both are classified as waste; ore loss occurs when the base case is ore but the simulated SMU is waste; and dilution occurs when the base case is waste but the simulated SMU is ore. Classification is probability-weighted across the 100 realisations per domain. Economic translation uses probability-weighted tonnages and grades for each class (Rossi and Deutsch, 2014), combined with confidential site-specific copper price, metallurgical recovery, and mining, processing and refining cost assumptions. The numerical economic inputs are withheld because they are commercially confidential. Four indicators are computed: net profit, lost profit, profit per tonne moved and profit per tonne processed.

KEY RESULTS

At the baseline 0.25 per cent Cu cut-off, the correct-to-misclassification ratio rises from 1.84 at 10 × 10 × 15 m to a peak of 2.04 at 55 × 55 × 15 m, then declines slightly; total misclassification probability falls only from 35.2 to 32.9 per cent over this range, driven mainly by a reduction in dilution probability (12.9 to 9.3 per cent), while ore loss remains stable (21.9–24.0 per cent) and ore/waste proportions are preserved across sizes. This narrow overall movement indicates that the observed change in misclassification is modest and cannot be attributed to SMU size alone because the comparison also incorporates differences in data density and estimation method. At the 0.35 per cent cut-off the classification ratio reaches 3.11 at 35–45 m SMUs, showing that a higher cut-off sharpens classification confidence at the expense of overall ore inclusion.

These classification trends translate into tonnage trade-offs: larger SMUs produce higher total ore tonnage (~79 to ~93 Mt) but lower average ore grade (0.643 to 0.609 per cent Cu), consistent with the volume-variance effect (Frykman and Deutsch, 1999); dilution tonnage falls from ~42 to ~30.8 Mt by 60 × 60 × 15 m, while ore loss tonnage rises from ~72 to over 83 Mt, so larger SMUs trade fine-scale ore recovery for dilution control.

This trade-off resolves two distinct economic optima rather than three unrelated ones. Net profit peaks at 40 × 40 × 15 m, with a flat plateau at larger sizes; profit per tonne moved and per tonne processed both instead peak at 30 × 30 × 15 m, because unit profitability is total profit divided by tonnage, and tonnage keeps rising past the point where total profit growth flattens. The two metrics therefore answer different questions rather than disagreeing: 40 × 40 × 15 m maximises absolute project value, whereas 30 × 30 × 15 m maximises profit per unit of material moved or processed. Sensitivity analysis confirms this pattern is robust to cut-off grade, which shifts profit levels vertically without moving either optimum.

CONCLUSIONS

A central insight of this work is that total project value and unit profitability are maximised at different SMU sizes, 40 × 40 × 15 m and 30 × 30 × 15 m respectively, a distinction that follows mechanically from dividing net profit by tonnage rather than reflecting two independent optima. Similar dual-optimum behaviour has been reported in mine planning studies that explicitly account for geological uncertainty (Camara, Scheffer Leal and De Lemos Peroni, 2020). This trade-off is governed by the competing effects of grade averaging, which homogenises ore and waste at larger scales and reduces dilution, and selectivity, which discriminates finer at the cost of higher uncertainty per block. Operations focused on overall asset value should therefore favour 40 × 40 × 15 m, while operations constrained by processing throughput or capital intensity may prefer 30 × 30 × 15 m.

The framework links geological uncertainty to financial risk through probability-weighted economic calculations based on 100 conditional realisations (Rossi and Deutsch, 2014). However, the base-case grade-control model is not ground truth or perfect selective mining. Because it is based on dense blasthole data and IDW, while the test realisations use sparser exploration data and SGS, the comparison combines support size, data density and estimation method. The results should therefore be interpreted as a case-specific decision framework rather than as an isolated causal estimate of block-size effects.

In conclusion, this study applies an established, uncertainty-aware simulation and economic classification framework specifically to the problem of SMU-size selection, integrating geostatistical simulation, misclassification analysis, and mine economics. Applied to a copper deposit, it identifies 40 × 40 × 15 m as the value-maximising SMU and 30 × 30 × 15 m as the unit-profitability optimum. The methodology provides a framework that could be adapted to other commodities and deposit types, subject to appropriate data, estimation and economic assumptions.

REFERENCES

Camara, T R, Scheffer Leal, R and De Lemos Peroni, R, 2020. Accounting for operational dilution by incorporating geological uncertainties in short-term mine planning, *DYNA*, 87(213):178–183. https://doi.org/10.15446/dyna.v87n213.83661

de Carvalho, J P, Dimitrakopoulos, R and Minniakhmetov, I, 2019. High-Order Block Support Spatial Simulation Method and Its Application at a Gold Deposit, *Mathematical Geosciences*, 51(6):793–810. https://doi.org/10.1007/s11004-019-09784-x

Dimitrakopoulos, R and Godoy, M, 2014. Grade control based on economic ore/waste classification functions and stochastic simulations: Examples, comparisons and applications, *Transactions of the Institutions of Mining and Metallurgy, Section A: Mining Technology*, 123(2):90–106. https://doi.org/10.1179/1743286314Y.0000000062

Frykman, P and Deutsch, C V, 1999. Geostatistical Scaling Laws Applied to Core and Log Data, paper presented to SPE Annual Technical Conference and Exhibition, Houston. https://doi.org/10.2118/56822-ms

Krige, D G, 1999. Essential basic concepts in mining geostatistics and their links with geology and classical statistics, *South African Journal of Geology*, 102(2):147–151.

Leuangthong, O, Neufeld, C and Deutsch, C V, 2003. Optimal Selection of Selective Mining Unit Size for Geostatistical Modeling, Centre for Computational Geostatistics report, University of Alberta. https://api.semanticscholar.org/CorpusID:198111940

Manchuk, J G and Deutsch, C V, 2012. Implementation aspects of sequential Gaussian simulation on irregular points, *Computational Geosciences*, 16(3):625–637. https://doi.org/10.1007/s10596-012-9274-6

Rossi, M E and Deutsch, C V, 2014. *Mineral Resource Estimation* (Springer: Netherlands). https://doi.org/10.1007/978-1-4020-5717-5

Vasylchuk, Y, V and Deutsch, C V, 2018. Improved grade control in open pit mines, *Mining Technology: Transactions of the Institutions of Mining and Metallurgy*, 127(2):84–91. https://doi.org/10.1080/14749009.2017.1363991

Unlocking opportunities through a stochastic production plan optimisation of an iron ore greenfield mining project

J Carvalho[1], A Mayrink[2], A Amaral[3], D Barbosa[4], A Ditlef[5], G Santos[6] and A Silva[7]

1. Mining Engineer, Vale S.A, Santa Luzia, MG, Brazil. Email: joao.carvalho@vale.com
2. Mine Planning Consultant, KPI Mining, Montreal QC, Canada. Email: alisson.mayrink@kpimining.com
3. Geologist, Vale S.A, Santa Luzia, MG, Brazil. Email: aline.amaral@vale.com
4. Mining Engineer, Vale S.A, Santa Luzia, MG, Brazil. Email: diego.barbosa@vale.com
5. Principal Mining Engineer, Vale S.A, Santa Luzia, MG, Brazil. Email: alexandre.ditlef@vale.com
6. Manager, Vale S.A, Santa Luzia, MG, Brazil. Email: gesner.santos@vale.com
7. General Manager, Vale S.A, Santa Luzia, MG, Brazil. Email: alessandro.silva@vale.com

ABSTRACT

Geological uncertainty and variability are inherent to mine planning, largely due to budget constraints that prevent systematically sampling and thorough characterisation of mineral deposit. Recognising uncertainty from different sources, particularly related to the mineralisation, is essential to more-informed decision-making. This is especially important during the exploration stage, when project parameters and aspects can still be adjusted at relatively low cost. This work presents a case study consisting of a single mine where two mineral processing routes are available for generating iron ore concentrates: direct shipping ore (DSO) and pellet feed (PF). The DSO plant involves only a comminution circuit and requires an important control over the feeding ore's iron grade. The PF plant incorporates milling and flotation circuits, which increases processing costs but generates more valuable products. To evaluate this project, stochastic simulations of the orebody model were generated to quantify uncertainty in iron grades and deleterious elements, such as silica, phosphorous and manganese. These simulations were inputted into a stochastic production planning optimiser to generate risk-resilient production schedules to maximise the project's net profit value (NPV) and mitigate the risk of not meeting production targets. Three scenarios are evaluated in attempt to capture the full economic potential of the project: Case 1 considers a stochastic optimisation using only measure and indicated resources are treated as candidates to feed the processing plant. Next, inferred resource are also incorporated into the model to investigate extending the life-of-mine (LOM) (Case 2). Lastly, optimistic prices for iron ore concentrates are applied to explore additional benefits of a best-case scenario (Case 3). A deterministic production plan was also developed and used as baseline for comparison. In this workflow, the grades of interest were estimated into the orebody model as this was used as input for the optimiser. Case 1 outperformed the baseline, increasing the overall production of iron ore concentrate by 6.1 per cent, having additional 16 per cent DSO product, and extending the LOM by two periods. This is explained by the smoothing effect of estimation methods that underestimated the high-grades mostly associated with hematite ore. Additionally, Case 2 further extended the LOM and increased the production of iron ore concentrate by 6.2 per cent compared to Case 1. The production schedule presented in Case 3 do not differ drastically from Case 2, in terms ore processed and products generated; however, the increase in NPV is substantial, highlighting the project's potential for stronger financial performance under favourable market conditions.

INTRODUCTION

Industrial mining complexes can be viewed as several interconnected activities and operations that started from the extraction of mineralised elements from one or several mines and transported through one among multiple processing routes, in which the material undergoes transformations to generate mineral products (Whittle and Whittle, 2007; Pimentel, Mateus and Almeida, 2010; Goodfellow and Dimitrakopoulos, 2016; Dimitrakopoulos and Lamghari, 2022). In this context, long-term planning of mining complexes involves defining strategic decisions, such as extraction sequence and material destination definitions, pit limits, control of deleterious elements at processing plants and net profit value (NPV) maximisation (Hustrulid, Kuchta and Martin, 2013). All these decisions must be optimised simultaneously while the risk of not delivering the forecasted production

and respecting targets is minimised. Industry practices optimise the long-term production schedule based on an estimated model and involves several optimisations steps that are performed sequentially. There are important drawbacks of this approach. First, estimated models are no able to capture geological uncertainty, which is an inherent characteristic of mineral deposits due to limited samples available. Additionally, any estimation approach incurs a smoothing effect, in which the forecasted value underestimates high values and overestimate low values (Goovaerts, 1997; Chilès and Delfiner, 1999). This is particularly important in mineral deposits as it generates models misrepresenting tonnages and grades that drive production schedule. The result is a production schedule that often deviates from production targets. Another important aspect is that industrial software optimises the model in a stepwise fashion, preventing to reach a more profitable solution.

Conversely, the field of stochastic mine planning has been developed for a few decades, and several studies have documented important gains compared to deterministic and traditional optimisations approaches (Dimitrakopoulos, 2011, 2018; Montiel, Dimitrakopoulos and Kawahata, 2016; Goodfellow and Dimitrakopoulos, 2017). In addition to geological uncertainty the field of simultaneous stochastic optimisation of mining complexes (Dimitrakopoulos and Lamghari, 2022) have incorporated multiple components in a single mathematical formulation, for example, operating modes (Montiel and Dimitrakopoulos, 2015), facility placement (Dimitrakopoulos, Fernanda and Castillo, 2026), acid drainage and waste dump management (Levinson and Dimitrakopoulos, 2019; Guimaraes and Dimitrakopoulos, 2026), in-pit crushing and conveying systems (Findlay and Dimitrakopoulos, 2024), cut-off grade (Menabde *et al,* 2007; Cutler and Dimitrakopoulos, 2025), other uncertainty sources, such as market (Saliba and Dimitrakopoulos, 2019), equipment (Matamoros and Dimitrakopoulos, 2016; Both and Dimitrakopoulos, 2020), geometallurgical attributes (de Carvalho and Dimitrakopoulos, 2025; Jiang and Dimitrakopoulos, 2025). Due to these studies, it has becoming more commonly applied in industry.

This work presents a case study applying the simultaneous stochastic optimisation approach (Goodfellow and Dimitrakopoulos, 2016, 2017) to an iron ore mining complex. The project under study is at a greenfield stage, and the goal is to evaluate more thoroughly different opportunities in this phase to unlock profitable opportunities, such as the extension of the life-of-mine (LOM) or generation of more products. As the project moves forward there is less room for modifications, which may prevent reaching its full potential. Three stochastic optimisation cases are evaluated: considering measured and indicated resource to feed the processing plant, allowing inferred resources, and applying optimistic iron concentrate prices. The following section describes the case study in consideration. Conclusions follow.

CASE STUDY

The operation considers a single mine, where the different material types can be extracted and transported to one of the corresponding facilities: waste dump, stockpile, direct-shipping ore (DSO) plant or pellet feed (PF) plant. Figure 1 displays the mining complex possible routes. The DSO plant only comprises only a comminution stage to generate direct shipping ore DSO products. This process requires a minimum iron grade must be guaranteed, which implies only hematite ore can be processed. Regarding the PF plant, more processing stages, such as milling and flotation, are in place. Thus, lower grade ore can be processed, which means that itabirites and hematite rocks are possible candidates. Once the ore is fed into this processing facility, the mass recovery is calculated based a non-linear function given iron and deleterious content. Additionally, the flotation step can be impacted by a large concentration of manganese; therefore, the manganese content should be monitored and an upper bound is applied. The PF plant generates two products, a direct-reduction (DR) and a blast-furnace (BF) iron concentrate. The former is a more valuable and richer product, while the latter is generated by inefficiencies of the process.

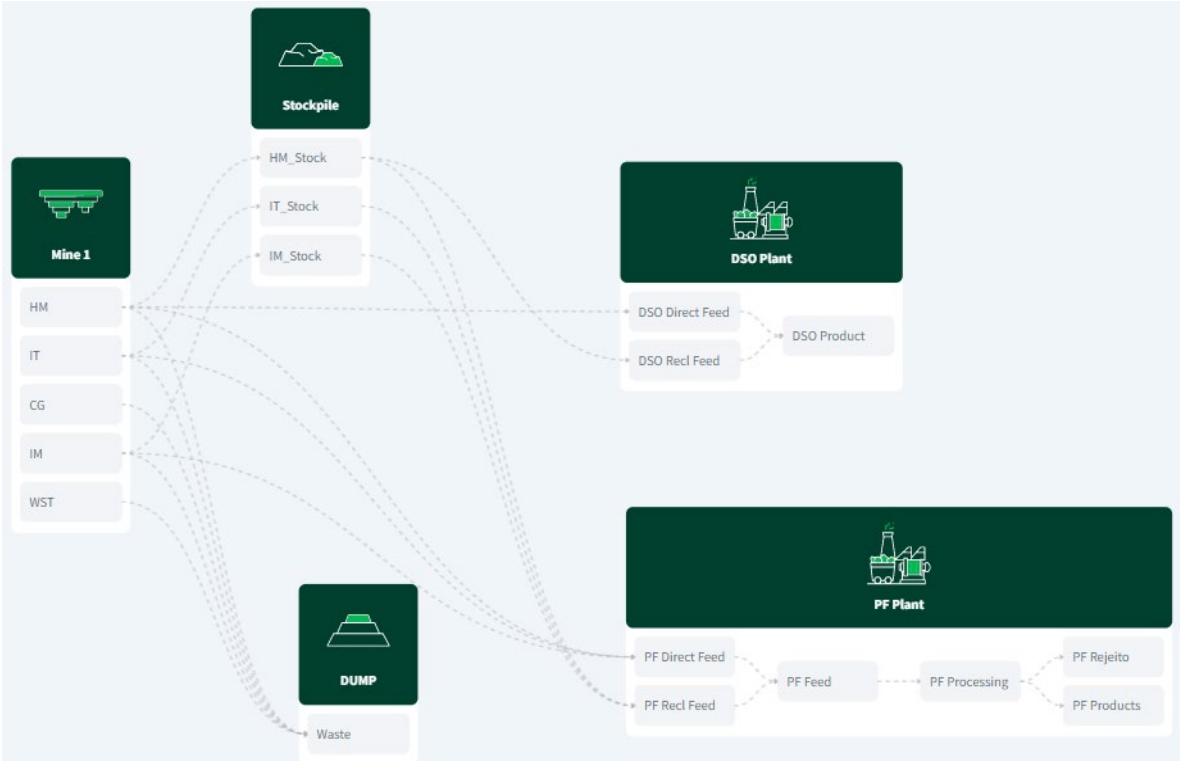
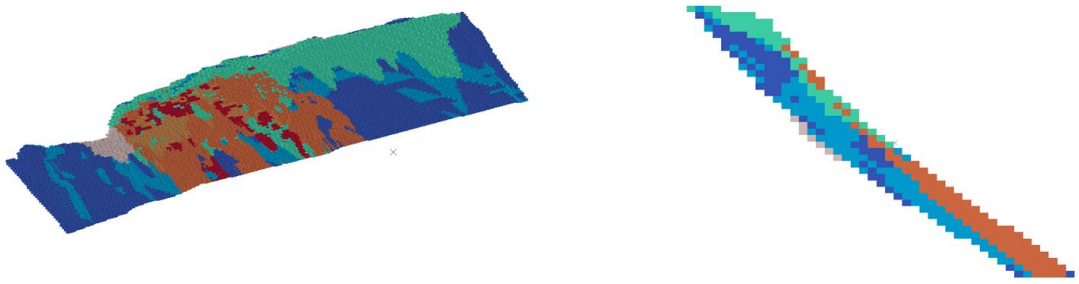


FIG 1 – Diagram of the iron ore mining complex.

The source of material is a typical iron ore deposit composed of canga, hematite rocks, itabirite and manganese-rich itabirite, which is illustrated by Figure 2. There is also weathering classification subdividing hematites and itabirites into friable and compact rocks. Other formations are also present and classified as waste, which is mostly characterised as quartzite. This iron ore deposit is the source of material for the mining complex in consideration.



(a) (b)

FIG 2 – Representation of the orebody model displaying the lithologies: (a) isometric view; (b) cross-section.

Geostatistical simulations of the orebody models

The lithologies were conventionally modelled considering the geologist expertise and were not simulated in this study. However, future work shall address this. Regarding the grade of the different elements, several elements were considered, such as iron, silica, phosphorous and manganese. To maintain the inherent correlation between these elements, it was applied a Projection Pursuit (Barnett, Manchuk and Deutsch, 2014) approach to decorrelated at lag 0 and min/max autocorrelation factors (Desbarats and Dimitrakopoulos, 2000) to continue the attributes decorrelation at all lags. Next, each factor was simulated individually using the Turning Bands approach (Journel and Huijbregts, 1978). Figure 3 displays three realisations illustrating the spatial distribution of iron and manganese grades.

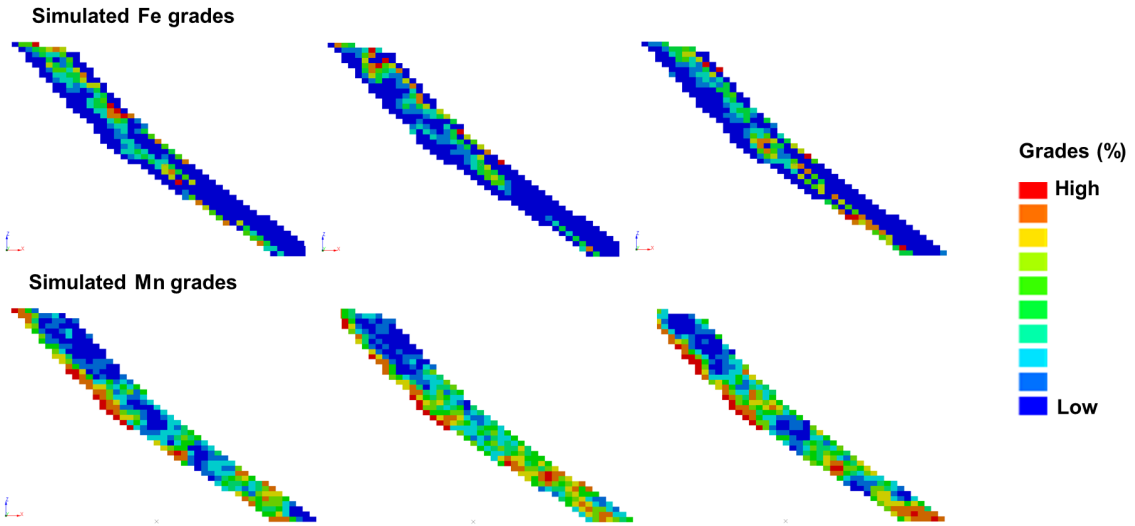


FIG 3 – Cross-section of the generated simulations showing the spatial distribution of Fe and Mn.

Figure 4 shows the grade tonnage curves per lithologies. In all cases it can be observed the smoothing effect, in which as the cut-grade grows, less tonnages are available in the estimated (kriged) model. In the other hand, the simulations respect the data statistics and better reproduce the spatial distribution of grades throughout the orebody model. Additionally, the smoothing effect brings large values towards the mean, which means that the overall grade above cut-off drops below the lines representing the simulations. Note that the smoothing effect is more pronounced for hematite rocks, as these contain higher iron grades. This aspect is going to have an important impact on the production schedule evaluated in the 'Simultaneous Stochastic Optimisation' Section.

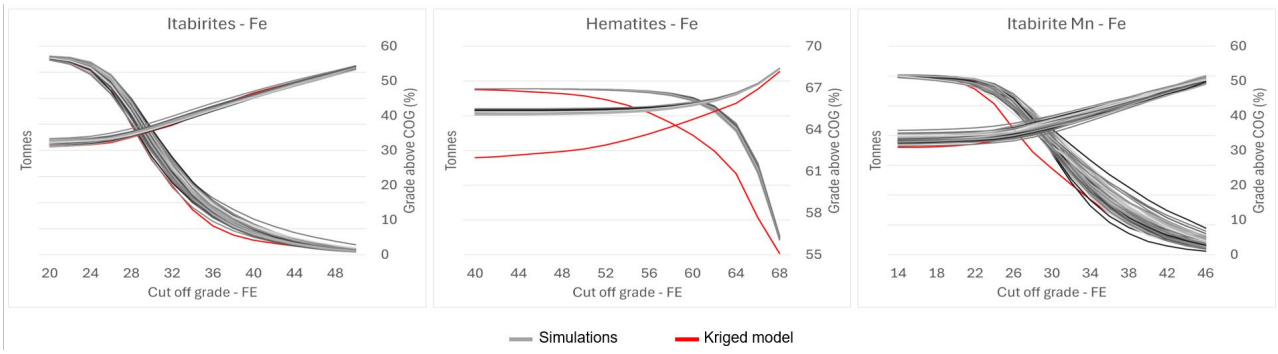


FIG 4 – Grade-tonnage curves for hematite and itabirite rocks.

Risk analysis of contaminants

A kriged model was also generated, however the material destination obtained from a deterministic optimisation was not available, which prevented generating a proper risk profile. Thus, to perform an initial analysis regarding the phosphorous content per annum, it was adopted the following strategy. Wireframes representing the portions mined per period were used as a filter to evaluate the overall content of phosphorous, which is presented by Figure 5. The simulations revealed that estimated model has considerable chances of underestimating the amount of phosphorous mined throughout the life-of-mine (LOM).

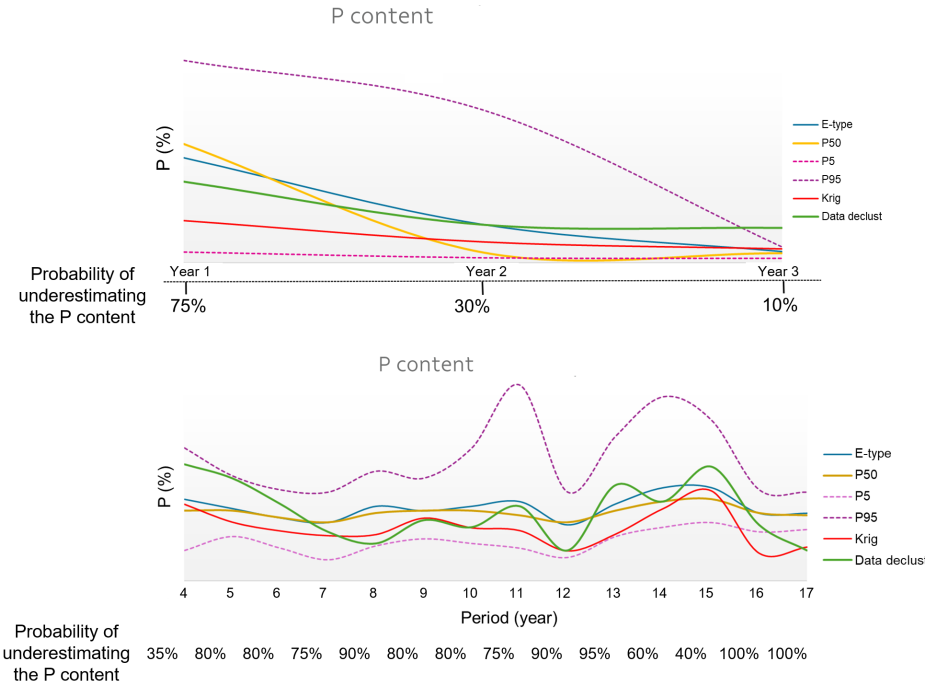


FIG 5 – Evaluation of phosphorous content mined per period by the deterministic optimisation.

Simultaneous stochastic optimisation

Next, the model is optimised based on the method described by (Goodfellow and Dimitrakopoulos, 2016, 2017). Three stochastic optimisation cases were generated:

- **Case 1** considers that only material classified as measured or indicated can feed the processing plant.
- **Case 2** allows material classified as inferred as potential ore.
- **Case 3** expands Case 2 and applied optimistic future iron ore prices.

Additionally, there is also available a production plan generated using a deterministic orebody model and a conventional optimisation framework, which is referred to as Base Case. However, the extraction sequence and material destination per block was not available at the time; thus, it was not possible to perform a risk analysis of this deterministic model.

Figure 6 shows the extraction sequence obtained in Case 1. This is the basis for the production plan presented in Figures 7 and 8, and are compared to Base Case. As pointed by past studies (Dimitrakopoulos, 2011), the results from the deterministic plan are very likely to deviate from the production targets because they disregard geologic uncertainties, which misrepresents tonnages and grade. On the other hand, Case 1 account for grade uncertainties and still provides a more stable ore throughput at the PF plant over the LOM, as shown in Figure 7a. Additionally, the LOM is extended by two years in both PF and DSO plants. Figure 7b displays the cumulative production over the LOM. Case 1 produces considerably more DSO product than the Base Case. The smoothing effect of deterministic models underestimates the high Fe grades especially the ones present in hematite rocks. Thus, the deterministic optimisation identifies the block model having less tonnages that what is observed in the simulations. Conversely, the stochastic simulations capitalise in this better characterisation of the orebody model and indicate that it is possible to generate 16 per cent more DSO concentrate. Overall, Case 1 produces less PF Fe concentrate (DR+BF), even though it provides more ore to the processor. This is explained by the non-linear grade-recovery function. The grades in Case 1 are smaller than initially forecasted by the Base Case, as shown in Figure 8, which resulted in a drop in mass recovery, Figure 8b. However, if a risk analysis was performed over the deterministic model, it is very likely to also observe a considerable drop in this performance indicator. Another important aspect is the Mn grades at the feeding ore, which is presented in Figure 8c. The Base Case forecasted that will respect the Mn targets, but it has no control over its uncertainty and variability. Case 1 ensures an adherence to this target and postpone

riskier periods to the future. Still, the Mn targets are always respected in at least 50 per cent of the scenarios.

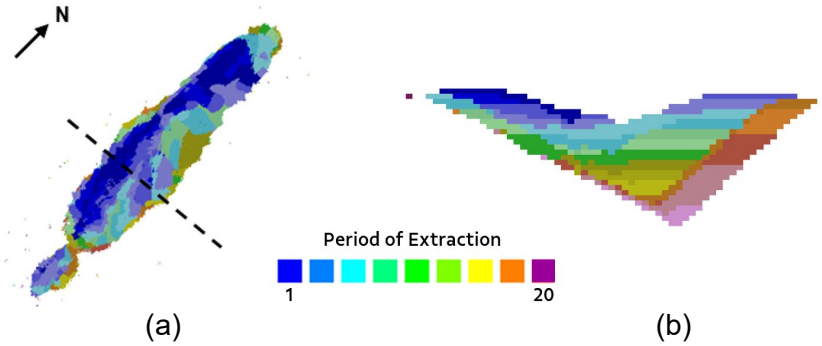


FIG 6 – Extraction sequence of Case 1: (a) horizontal plane; (b) cross-section.

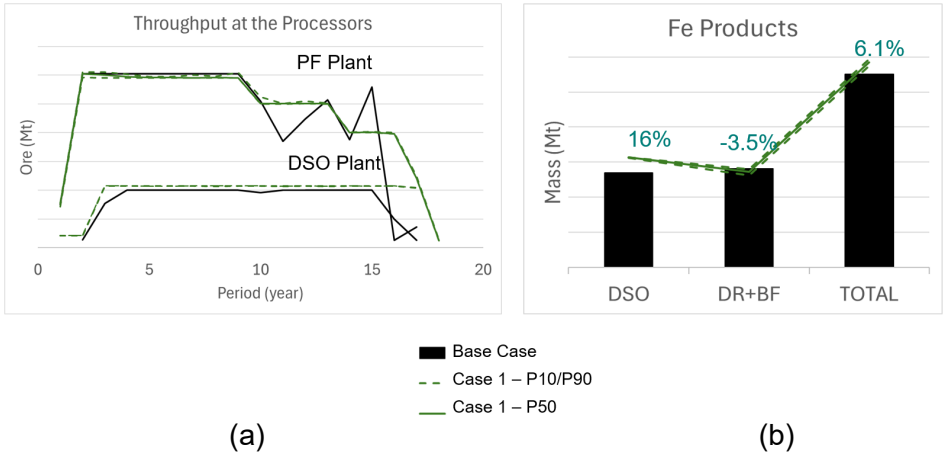


FIG 7 – Comparison of: (a) ore throughput at the processing plants; and (b) iron products generated for Case 1 and Base Case.

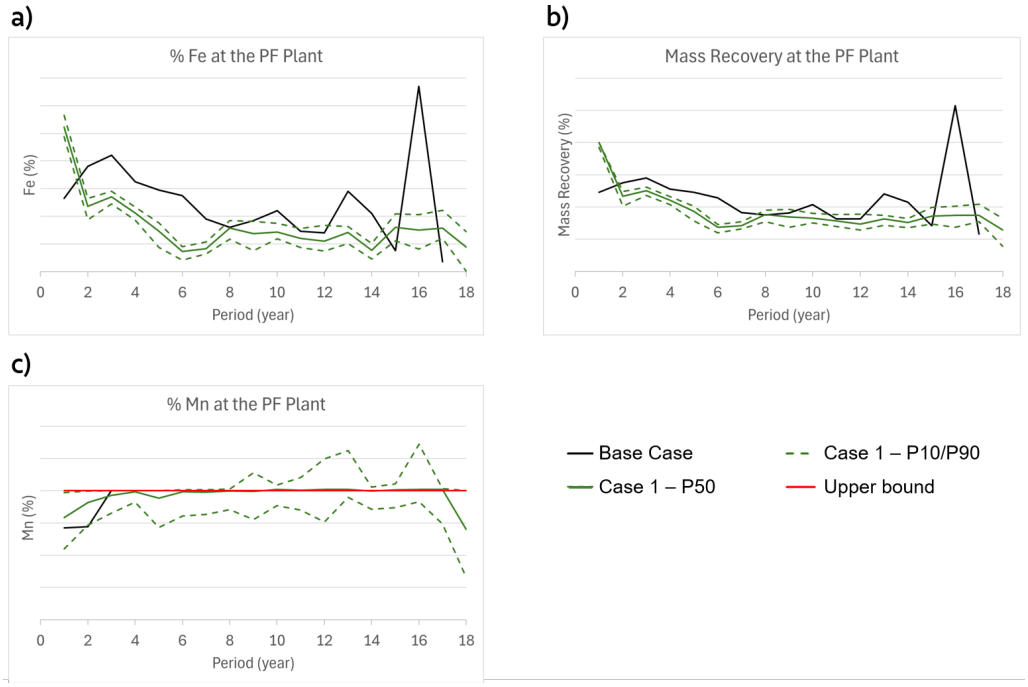


FIG 8 – Comparison of: (a) Fe head grade at the PF plant; (b) mass recovery at the PF plant; and (c) Mn grade at the PF plant for Case 1 and Base Case.

Case 2 investigates the potential that inferred resource could add into the project while also accounting for geological uncertainty. The optimised extraction sequence is displayed in Figure 9. There are some similarities with the one obtained by Case 1, Figure 6. But it holds differences as more material can be processed as ore in this scenario. Figures 10 and 11 present production schedule graphs in which the comparison is performed against Case 1. Case 2 extends the LOM by two years, and can feed the PF plant at a higher rate for longer, as shown in Figure 10a. This additional ore is translated into 6.4 per cent more products, Figure 10b. This is particularly higher for PF products, which is more valuable. Figures 11a and 11b present similar values for Fe head-grades and mass recoveries for both Cases 1 and 2. This is also not a problem in terms of Mn content at the PF plant, in which both cases perform similarly, postponing riskier periods. This provides interesting insights to decision-makers as the mine could be still expanded and result in more profit as additional information is incorporated into the model.

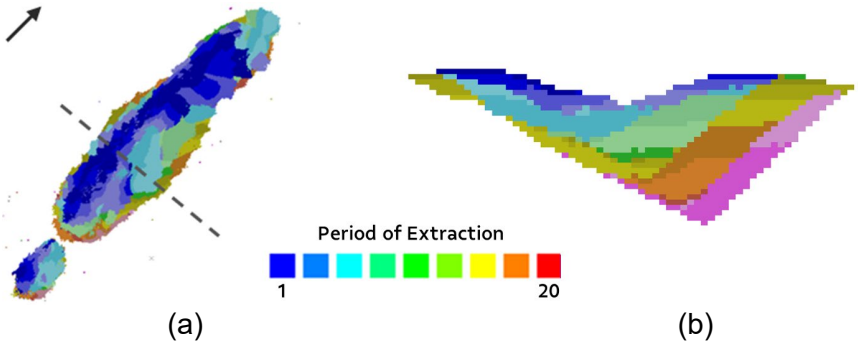


FIG 9 – Extraction sequence of Case 2: (a) horizontal plane; (b) cross-section.

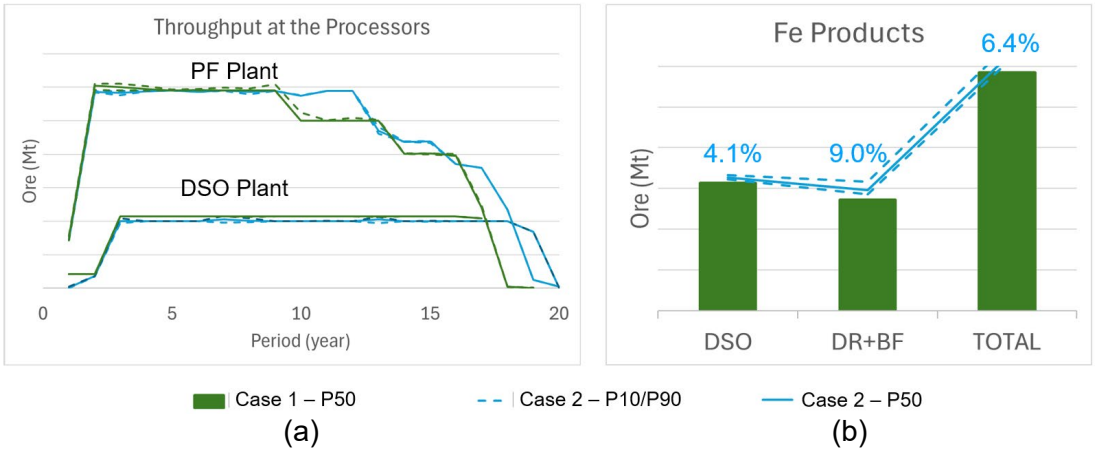


FIG 10 – Comparison of: (a) ore throughput at the processing plants; and (b) iron products generated for Case 2 and Case 1.

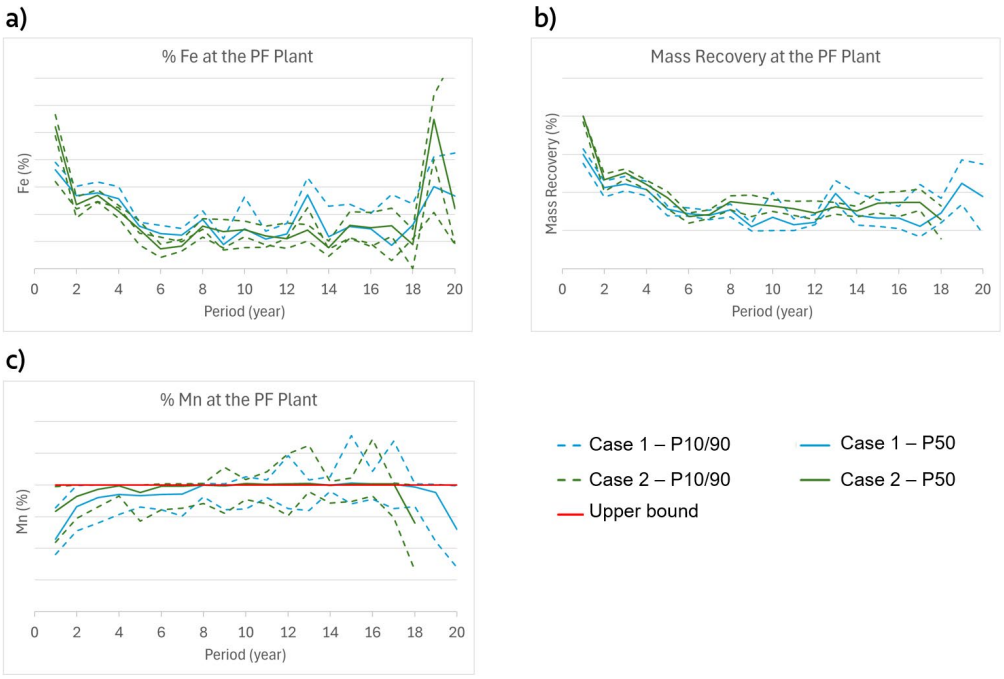


FIG 11 – Comparison of: (a) Fe head grade at the PF plant; (b) mass recovery at the PF plant; and (c) Mn grade at the PF plant for Case 2 and Case 1.

Case 3 investigates how profitable the project would be if higher iron ore prices are applied, which could be due to a higher demand of this product. Figure 12 displays the extraction sequence that differs from Cases 1 and 2. Regarding the production schedule, Figures 13 and 14 compare the results obtained for Cases 3 and 1. The improvements over Case 1 are slightly higher than the one presented by Case 2, which means that prices beyond the ones applied in Case 1 and 2 increase project value but do not modify drastically the production schedule.

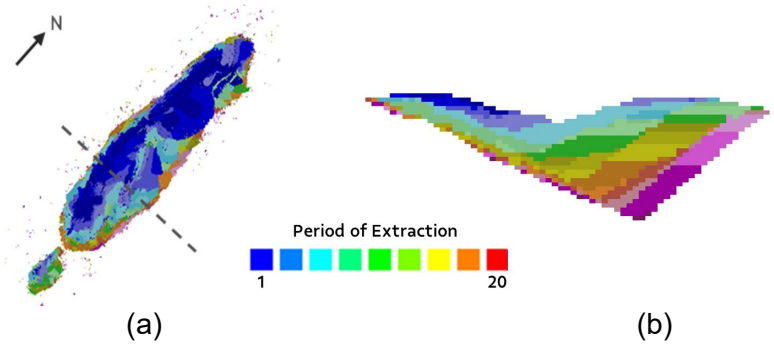


FIG 12 – Extraction Sequence of Case 2: (a) horizontal plane; (b) cross-section.

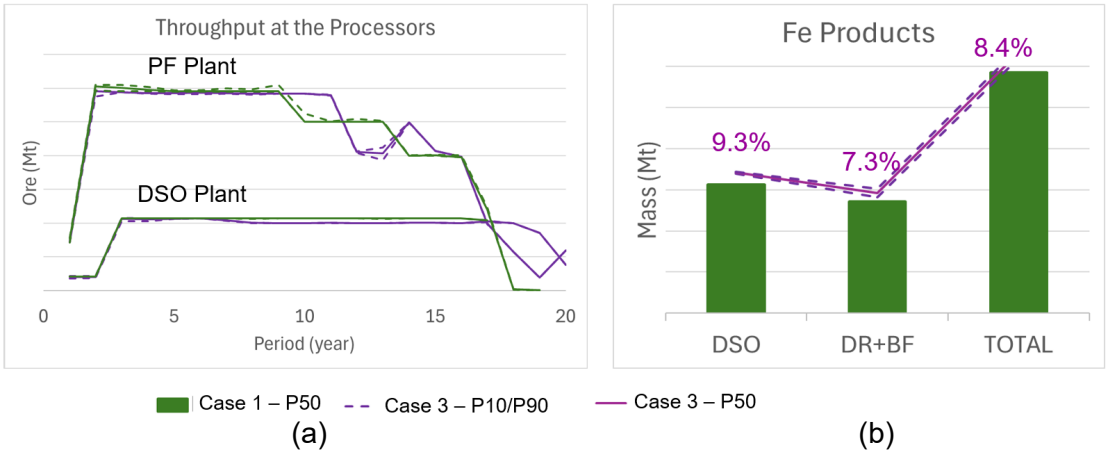


FIG 13 – Comparison of (a) ore throughput at the processing plants and (b) iron products generated for Case 3 and Case 1.

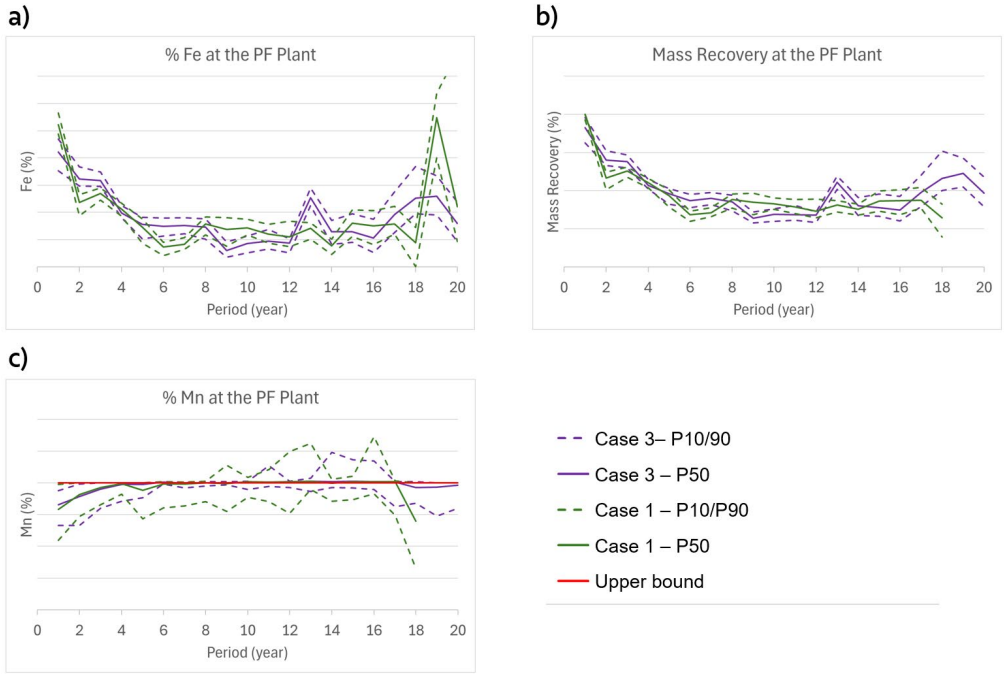


FIG 14 – Comparison of: (a) Fe head grade at the PF plant; (b) mass recovery at the PF plant; and (c) Mn grade at the PF plant for Case 2 and Case 1.

Figure 15 presents a summary of the results obtained. The percentage values displayed in the graphs are increments over the case in the left. Figure 15a shows the overall run-of-mine (ROM) forecasted and Figure 15b the cumulative iron ore products to be obtained in both processing plants. In the DSO plant, the material quality entering the plant is the same as one in the product. Regarding the PF plant, Fe grades and other deleterious impact the mass recovery. Thus, disregarding geological uncertainty and variability can lead to important deviations when mining, which is the result of estimation methods due to the smoothing effect. All stochastic optimisations presented, accounted for that and increase considerably the products to be obtained. Figure 15c presents the comparison in terms of NPV among the optimised cases. Due to confidentiality, it was not obtained the NPV forecasted by the Base Case. Thus, the comparison is based on Case 1. Case 2 improves the project only by 0.7 per cent, even though the LOM and ore throughput have increased by a larger margin. Regarding Case 3, the other performance indicators have not been increased substantially, however, the NPV grows 44.9 per cent. This shows that the project can be even more lucrative than expected if the market conditions are favourable. If Fe concentrate prices grows the production schedule does not be modified substantially to profit from this scenario.

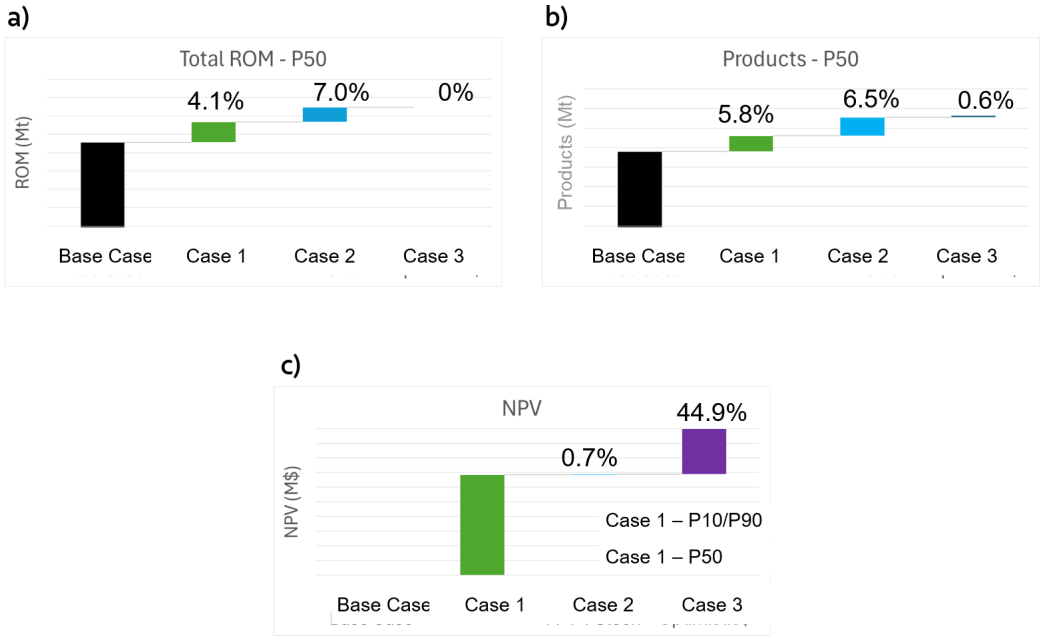


FIG 15 – Summary comparison of all cases regarding: (a) total ore processed; (b) products generated; and (c) NPV.

CONCLUSIONS

Mineral deposits are characterised by geological uncertainty due to limited samples that can be collected. Being able to quantify the uncertainty and managing the risk of not meeting production targets are crucial factors when developing a mineral project. Several studies over the past decades have associated geologic uncertainties with projects performing below expectations or not reaching its full potential. This is especially important at the exploration stage, when there is more room for changes and more potential to profit from. The work presented here evaluates the economic potential of an iron ore mining complex project. Geostatistical simulations are generated to characterise the spatial uncertainty and variability associated with the grades of iron and other deleterious elements. These simulations were input for a stochastic optimisation model that unlocked the economic potential of the iron ore mineral deposit. The mining complex in consideration is composed of a single mine, a waste dump and two mineral processing routes to generate iron ore concentrates. Three stochastic optimisation cases were investigated: Case 1 considered measured and indicated resource feeding the processing plant; Case 2 allowed inferred material; and Case 3 extended Case 2 by applying optimistic future iron ore prices. The results are also compared to a Base Case which uses a deterministic orebody model and traditional optimisation approach. Overall, the cases applying stochastic optimisations forecasted more iron ore products to be sold and extended the LOM by 2–4 years compared to the Base Case. The PF processing plant can operate at full capacity for longer periods, while controlling Mn grades, which could have been a potential concern if not managed. Including the inferred ore into the stochastic optimisation model led to gains in several performance indicators, even though the impact in NPV is marginal. On the other hand, increasing the Fe product prices did not present great variations compared to Case 2, but substantially improved the NPV.

Future work will focus on simulating lithologies as well as this defines the possible mineral processing routes and therefore impacts considerably the project financial performance. Including different mines nearby can also add substantial value to the project by capturing synergies in a larger mining complex.

ACKNOWLEDGEMENTS

The work in this article was supported by the Vale S.A.

REFERENCES

Barnett, R M, Manchuk, J G and Deutsch, C V, 2014. Projection Pursuit Multivariate Transform, *Math Geosci*, 46(3):337–359. https://doi.org/10.1007/s11004-013-9497-7

Both, C and Dimitrakopoulos R, 2020. Joint stochastic short-term production scheduling and fleet management optimization for mining complexes, *Optim Eng,* 21(4):1717–1743. https://doi.org/10.1007/s11081-020-09495-x

Chilès, J P and Delfiner, P, 1999. *Geostatistics: modeling spatial uncertainty* (Wiley).

Cutler, J and Dimitrakopoulos, R, 2025. Optimising multi-element cut-off grades for a strategic production plan under geological uncertainty, *International Journal of Mining, Reclamation and Environment,* 39(7):531–545. https://doi.org/10.1080/17480930.2025.2455567

de Carvalho, J P and Dimitrakopoulos, R, 2025. Integrating geometallurgy in the simulation of processing plant responses and applying policy gradient for shovel allocation in operating mining complexes, *International Journal of Mining, Reclamation and Environment*, 39(10):740–767. https://doi.org/10.1080/17480930.2024.2415059

Desbarats, A J and Dimitrakopoulos, R, 2000. Geostatistical simulation of regionalized pore-size distributions using min/max autocorrelation factors, *Math Geol,* 32(8):919–942. https://doi.org/10.1023/A:1007570402430

Dimitrakopoulos, R and Lamghari, A, 2022. Simultaneous stochastic optimization of mining complexes - mineral value chains: an overview of concepts, examples and comparisons, *International Journal of Mining, Reclamation and Environment,* 36(6):443–460. https://doi.org/10.1080/17480930.2022.2065730

Dimitrakopoulos, R, 2011. Stochastic optimization for strategic mine planning: A decade of developments, *J Min Sci,* 47(2):138–150. https://doi.org/10.1134/S1062739147020018

Dimitrakopoulos, R, 2018. *Advances in applied strategic mine planning* (Springer: Cham).

Dimitrakopoulos, R, Fernanda, M and Castillo, D, 2026. Strategic adaptations of infrastructure / facilities in an operating mining complex under geological uncertainty, *International Journal of Mining, Reclamation and Environment* 40(2):107–120. https://doi.org/10.1080/17480930.2025.2494110

Findlay, L and Dimitrakopoulos, R, 2024. Stochastic optimization for long-term planning of a mining complex with in-pit crushing and conveyance systems, *Mining Metall Explor,* p 15. https://doi.org/10.1007/s42461-024-01005-2

Goodfellow, R and Dimitrakopoulos, R, 2016. Global optimization of open pit mining complexes with uncertainty, *Appl Soft Comput*, 40:292–304. https://doi.org/10.1016/j.asoc.2015.11.038

Goodfellow, R and Dimitrakopoulos, R, 2017. Simultaneous stochastic optimization of mining complexes and mineral value chains, *Math Geosci*, 49(3):341–360. https://doi.org/10.1007/s11004-017-9680-3

Goovaerts, P, 1997. *Geostatistics for natural resources evaluation* (Oxford University Press: New York).

Guimaraes, V F and Dimitrakopoulos, R, 2026. Simultaneous stochastic optimisation of mining complexes: integrating waste management and progressive reclamation with encapsulation, *International Journal of Mining, Reclamation and Environment,* 38(9):1–14. https://doi.org/10.1080/17480930.2026.2654147

Hustrulid, W A, Kuchta, M and Martin, R K, 2013. *Open pit mine planning and design*, 3rd edn (CRC Press: London).

Jiang, Y and Dimitrakopoulos, R, 2025. An application of simultaneous stochastic optimisation on an open-pit copper mining complex with supply, recovery and market uncertainties, *International Journal of Mining, Reclamation and Environment,* 39(1):74–92. https://doi.org/10.1080/17480930.2024.2381904

Journel, A G and Huijbregts, C J, 1978. *Mining geostatistics* (Blackburn Press: New York).

Levinson, Z and Dimitrakopoulos, R, 2019. Simultaneous stochastic optimisation of an open-pit gold mining complex with waste management, *International Journal of Mining, Reclamation and Environment,* pp 1–15. https://doi.org/10.1080/17480930.2019.1621441

Matamoros, M E V and Dimitrakopoulos, R, 2016. Stochastic short-term mine production schedule accounting for fleet allocation, operational considerations and blending restrictions, *European Journal of Operational Research*, 255(3):911–921. https://doi.org/10.1016/j.ejor.2016.05.050

Menabde, M, Froyland, G, Stone, P and Yeates, G A, 2007. Mining schedule optimisation for conditionally simulated orebodies, in *Orebody Modelling and Strategic Mine Planning* (ed: R Dimitrakopoulos), Spectrum Series 14, pp 379–384 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Montiel, L and Dimitrakopoulos, R, 2015. Optimizing mining complexes with multiple processing and transportation alternatives: An uncertainty-based approach, *European Journal of Operational Research*, 247(1):166–178. https://doi.org/10.1016/j.ejor.2015.05.002

Montiel, L, Dimitrakopoulos, R and Kawahata, K, 2016. Globally optimising open-pit and underground mining operations under geological uncertainty, *Min Technol,* 125(1):2–14. https://doi.org/10.1179/1743286315Y.0000000027

Pimentel, B S, Mateus, G R and Almeida, F A, 2010. Mathematical models for optimizing the global mining supply chain, in *Intelligent Systems in Operations: Methods, Models and Applications in the Supply Chain* (ed: B Nag), pp 133–163 (IGI Global).

Saliba, Z and Dimitrakopoulos, R, 2019. Simultaneous stochastic optimization of an open pit gold mining complex with supply and market uncertainty, *Transactions of the Institution of Mining and Metallurgy: Mining Technology*, 128(4):A216–A229. https://doi.org/10.1080/25726668.2019.1626169

Whittle, J and Whittle, G, 2007. Global long-term optimization of very large mining complexes, in Proceedings of 33rd International Symposium on Application of Computers and Operations Research in the Mineral Industry, APCOM, Santiago.

Generating operational mine schedule using a stochastic block-based simultaneous optimisation framework

A Mayrink[1], M Faria[2], R Senécal[3], L Silva[4] and A Lioret[5]

1. Senior Mining Engineer, KPI Mining Solutions, Belo Horizonte MG, Brazil. Email: alisson.mayrink@kpimining.com
2. Director of Technical Services, KPI Mining Solutions, Montreal QC, Canada. Email: matheus.faria@kpimining.com
3. Mine Optimization Technical Lead, KPI Mining Solutions, Montreal QC, Canada. Email: renaud.senecal@kpidigital.com
4. Mine Optimization Software Developer, KPI Mining Solutions, Montreal QC, Canada. Email: luiz.silva@kpidigital.com
5. Mine Optimization Software Developer, KPI Mining Solutions, Montreal QC, Canada. Email: alain.lioret@kpidigital.com

ABSTRACT

Conventional strategic mine planning relies on deterministic resource models and an optimisation workflow supported by phase-bench scheduling assumptions, in which massive aggregations of blocks into benches or panels are homogenously extracted across periods. This approach misrepresents grade variability, recoverable tonnages and grades, and an asset's full economic potential since it departs from the selective mining unit (SMU) typically represented at block scale. Over the past decades, advanced simultaneous stochastic optimisation of mining complexes (SSOMC) allowed to model the complexity of block granularity decisions, global optimisation and stochastic mathematical formulations, as well as the ability to solve the extremely large models. However, existing formulations, despite incorporating constraints such as mining width and sink rates, still lack the operational practicality ensured by mine design, leaving a gap between block-scale optimisation and practical execution.

This paper presents a three-step stochastic framework bridging this gap by integrating geological uncertainty into practical stochastic mine scheduling while preserving block-level selectivity. The methodology: (1) applies existing SSOMC guided by geostatistical simulations to generate optimal and risk-resilient block-based extraction sequences; (2) defines phases by grouping yearly periods subsequently mine-designed, accounting for dilution, ore loss, and practical mining constraints; and (3) applies an adapted SSOMC that reintroduces designed phases as hard precedence constraints while maintaining block granularity and uncertainty management.

The framework is demonstrated on an open pit iron ore case study benchmarked against a conventional deterministic plan. Risk analysis of the conventional schedule revealed a probable (P_{50}) net present value (NPV) decrease of -8 per cent and 30 per cent probability of missing the mill target. The proposed phase-bench stochastic scheduler (PBS) delivered a probable (P_{50}) NPV improvement of +2 per cent over the conventional expected value (+11 per cent relative to risk-adjusted baseline), while consistently meeting the mill target across all periods and satisfying critical operational constraints. These results demonstrate the framework produces risk-resilient and practical long-term mine plans.

INTRODUCTION

A traditional strategic mine plan consists in producing a long-term production schedule for the life-of-mine (LOM) of a mine complex, usually driven by a single 'best-estimate' resource model of the materials-metals of interest and with the primary objective to maximise the net present value (NPV) while constrained by operational requirements at extraction and processing levels. Conventional mine planning methods are based on sequential optimisation workflow to define pit limits, pushback design, mine sequence of extraction and product-blending materials. Those methods have traditionally relied on aggregating blocks into panels based on pre-defined and not-scheduling driven pit limit and phases to achieve computational tractability, supported by simplified homogeneous phase-bench scheduling assumptions (Hustrulid, Kuchta and Martin, 2013; Lerchs and Grossmann,

1965; Johnson, 1968, 1969; Gershon, 1987). While effective from a computational standpoint, this traditional approach has significant limitations: leads to suboptimal solutions due to the incapacity to integrate synergies between its components and to misleading representation of recoverable ore tonnages and grades within the production volumes as well as production forecasts vulnerable to the inherent variability and uncertainty from mineral deposits (Dimitrakopoulos, Farrelly and Godoy, 2002).

Over the past decades, numerous studies have been published on block-based sequencing methodologies (Caccetta and Hill, 2003; Bley *et al,* 2010; Jelvez *et al,* 2016; Letelier *et al,* 2020), commercially known as direct block scheduling (DBS), to address some of the scheduling limitations. However, these methods have not addressed the risk associated with the deterministic mine planning. To overcome this limitation, block-based two-stage stochastic integer programming (SIP) formulation has been proposed (Ramazan and Dimitrakopoulos, 2013; Dimitrakopoulos, 2011). This method integrates geological uncertainty and variability through geostatistical simulations of the mineral deposits as inputs for the scheduling optimisation models. This formulation aims to maximise the NPV while minimise the risk of not meeting production targets throughout a risk-discount factor.

These earlier formulations were unable to optimise mining complexes – systems comprising multiple mines, processing plants, waste dumps, stockpiles, and other facilities operating in conjunction (Monkhouse and Yeates, 2007; Pimentel, Mateus and Almeida, 2010). To address this limitation, the SIP model was extended to accommodate multi-mine extraction, dynamic destination and processing-stream decisions, and non-linear transformations, shifting the basis of optimisation from the economic value of individual blocks to the economic value of final products. This evolution gave rise to what is now known as the simultaneous stochastic optimisation of mining complexes (Montiel and Dimitrakopoulos, 2015; Goodfellow and Dimitrakopoulos, 2016, 2017; Dimitrakopoulos and Lamghari, 2022).

To generate more practical block-based sequences, the aforementioned methods integrate mineability constraints such as mining width and sink rate, and have proven to be valuable guidance for the mine design process (Benndorf and Dimitrakopoulos, 2013; Dimitrakopoulos and Ramazan, 2004; Saliba and Dimitrakopoulos, 2019). However, these works do not address another recurrent requirement of long-term planning: reporting forecasts within designed phases while accounting for the dilution and ore loss incorporated during the mine design process. Aras, Dagdelen and Johnson (2019) demonstrated that a production schedule generated within phases designed from a previously optimised block-based schedule can deliver higher NPV compared to a schedule generated within phases formed by grouping nested pit-shells. These works, however, were deterministic and did not attempt to integrate geological uncertainty and variability. Farmer and Dimitrakopoulos (2018) proposed a methodology to form phases directly from a SIP schedule; however, this study did not evaluate the further impact of the mine design process on the schedule and the related yearly forecasts.

This work presents an innovative framework that adapts the existing stochastic optimisation model from Goodfellow and Dimitrakopoulos (2016, 2017) to account for mine designing considerations while maintaining block-based decisions in practical mine phase-bench oriented. The proposed methodology closes the ‘operational loop’ by, first, employing stochastic block scheduler – guided by a set of equiprobable geostatistical simulations – to define optimal mining sequences and end-year volumes. These stochastic outputs are then used to design operational phases, which are subsequently reintroduced into the adapted fine-tuning stochastic scheduling model. Unlike conventional tools, this scheduler operates on a phase-bench logic for vertical precedencies while maintaining block selectivity and incorporating uncertainty.

While the proposed approach still does not provide an algorithmic way to group optimised yearly production periods into phases, and relies on a degree of subjectivity to form and design them, the methodology nonetheless enables an end-to-end stochastic workflow for long-term mine planning. Its effectiveness is demonstrated through a case study of a large-scale iron ore deposit, comparing a conventional deterministic plan with phase-bench granularity against the proposed stochastic block-based scheduler with phase-bench orientation. Finally, conclusions and future work are presented.

METHOD

The proposed framework for a generic end-to-end LOM mine planning optimisation that integrates stochastic approach with practicality of results is summarised in Figure 1 and its three steps are described in the subsections. First, a simultaneous stochastic optimisation of mining complex (SSOMC) as per the original method proposed by Goodfellow and Dimitrakopoulos (2016, 2017) is employed to generate a stochastic block-based sequence. Following this, a grouping periods strategy defines phases that are submitted to a subsequent mine design. Finally, a modified version of the SSOMC model is used to generate a risk-resilient stochastic-practical schedule accounting for the designed shapes and phase-bench sequencing constraints, while maintaining the selective mining units (SMU) as the scheduling decision unit.

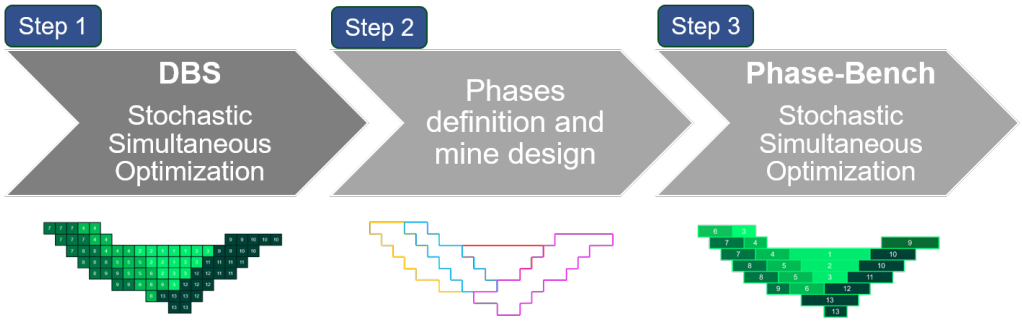


FIG 1 – Proposed stochastic strategic mine planning workflow.

SSOMC basic mathematical model

A mine complex can integrate C components (mines, stockpiles, plants, concentrators, dumps) and each one of them is denoted by $i \in C$. The only source of raw materials is described as any mine $m \in C$ and it is discretised into volumes of material called blocks $\mathbb{B}_m$, where each block is $b \in \mathbb{B}_m$. Any given material while flowing through those components (eg a block extracted from the mine m and sent to the plant P, stockpile K or to waste dump W) can have its A attributes or properties (such as mass, grade, metal, costs, revenue) quantified along the value chain and denoted by $a \in A$. The uncertainty in the pertinent attributes of the mineral deposit (eg grades, lithologies, commodity price) can be quantified using a set of joint scenarios, where each scenario $s \in S$ defines a realisation from all sources of uncertainty. The LOM optimisation has T production periods (t) time frame $(t \in T)$.

From the assumptions stated above, the general optimisation formulation in Equation 1 describes the objective function of the two-stage SIP applied for the open pit SSOMC in the first step of the proposed framework in this work.

$$\underbrace{max\ \frac{1}{|S|}\sum_{t\in T}\sum_{s\in S}\sum_{i\in C}\sum_{a\in A} p_{a,i,t} * v_{a,i,t,s}}_{\text{Part 1}} \underbrace{-\frac{1}{|S|}\sum_{t\in T}\sum_{s\in S}\sum_{i\in C}\sum_{a\in A}\left(c^{+}_{a,i,t} * u_{a,i,t,s} + c^{-}_{a,i,t} * l_{a,i,t,s}\right)}_{\text{Part 2}} \quad (1)$$

The first segment accounts for maximisation of discounted net value (NPV) from products generated. The term $p_{a,i,t}$ denotes revenues and expenses attributes time-discounted (based on a given discount rate) and $v_{a,i,t,s}$stands for the optimised attribute values. The second segment expresses the risk minimisation-deferral quantified by the deviation from targets (production attributes or 3D block aggregation components), where $u_{a,i,t,s}$ and $l_{a,i,t,s}$ are variables related to the upwards and downwards deviations, respectively, from production targets and $c^{+}_{a,i,t}$ and $c^{-}_{a,i,t}$ are the cost-penalty of deviation, defined by the user to prioritise the targets.

The optimiser can modify three decision variables that impact the mining complex. A binary decision $x_{b,t} \in \{0,1\}$ denotes the Extraction Sequence, defining whether (1) or not (0) a block $b \in \mathbb{B}_m$ is extracted in a period $t \in T$. The Destination Policy decisions $z_{g,j,t} \in \{0,1\}$ define whether a material of cluster g is sent to destination j in period t. Finally, the Processing Streams decisions $y_{i,j,t,s} \in$

$\{0,1\}$ define what proportion of material is sent from location $i \in P \cup K$ to destination j in period t and scenario s.

The objective function formulation is subjected to the mine reserves and to block open pit precedence constraints, aiming to enforce slope stability and one-time extraction of blocks, as described by Equations 2 and 3. The mining of block b is conditioned to the entire extraction of a set of overlying blocks denoted by $\mathcal{O}(b)$. This set of overlapping blocks can be obtained through a 3D slope-cone preprocessing algorithm (Khalokakaie, Dowd and Fowell, 2000).

$$\sum_{t \in T} x_{b,t} \leq 1 \; \forall b \in \mathbb{B}_m;\; m \in C \tag{2}$$

$$x_{b,t} \leq \sum_{t'=1}^{t} x_{u,t'} \; \forall b \in \mathbb{B}_m;\; m \in C;\; u \in \mathcal{O}(b);\; t \in T. \tag{3}$$

In these formulations, slope angle and reserve constraints are modelled as hard constraints. In addition, to ensure a degree of operational practicality in the optimised extraction sequences, smoothness constraints and sink rates are modelled as soft constraints. This means that blocks within a mining-width window, or within an overlying stack of blocks, are penalised if they are not scheduled to the same period as the block being evaluated.

This model is solved using a Multi-Neighbourhood Adaptive Simulated Annealing algorithm, which ranks and calls multiple low-level heuristics to perturb the block-based extraction sequence, material destinations, and processing streams, generating new solutions for the SSOMC model until convergence to a high-quality solution is reached. The details of this formulation and solution method are beyond the scope of this paper; the reader is referred to Goodfellow and Dimitrakopoulos (2016, 2017) for the original formulation.

As a result, this approach generates an optimised extraction sequence for each mine in the complex with a certain level of operational practicality. Although local violations of these mineability constraints can still occur, the resulting sequence can be used as guidance for the subsequent mine design step.

Phases definition and mine design

Once the optimised extraction sequence for each mine in the complex is obtained from Step 1, the end-of-year production volumes can be grouped to form scheduling-driven mining phases. The number of periods to be grouped – for example, every three or five years – can be defined based on multiple operational criteria, such as the need for three- or five-year phases to feed medium-term mine planning, operational permit constraints, plant start-up or shutdown schedules, relocation of existing infrastructure, and other site-specific considerations.

This grouping of periods inherently integrates the time value of money into the definition of phases. As a result, the formed phases exhibit a decreasing NPV contribution in sequential order: Phase 1 contributes the highest NPV, while the final phase contributes the lowest. This also implies that phases have potentially decreasing grades, since the optimisation performed in Step 1 (with its embedded cut-off grade optimisation) tends to prioritise high-grade areas in earlier periods. In addition, each end-of-period volume already satisfies blending requirements, and so, potentially, do the phases formed from them. This approach differs substantially from the traditional method of generating phases by grouping nested pit-shells.

Once the grouping of periods is defined, phase design can be performed on any commercial mine-design CAD platform, taking into account operational parameters such as ramp access, bench angle and width, and bottom width – the latter depending on equipment size and geotechnical considerations specific to the mine operation – as well as ensuring alignment across the multiple phases.

Modified phase-bench SSOMC

To accommodate the stochastic reschedule constrained by the phase design the main required modifications on the described SSOMC model are summarised by:

- The phase-by-phase precedence is determined based on whether one phase spatially overlies another.

- A new set of block precedences, denoted by $\mathcal{O}(b)$ in Equation 3, replaces the slope-conical algorithm with a hierarchical phase-bench-block precedence structure.
- The sink rate constraints were changed from soft to hard constraints, controlling the maximum number of benches mined in a period rather than the maximum number of blocks in a vertical stack.
- Additional phase-related scheduling constraints were added as hard constraints: the bench-lag and the maximum number of active phases controls.
- The Multi-Neighbourhood Adaptive Simulated Annealing implementation from the original SSOMC model was adapted to incorporate a new set of low-level heuristics that account for the new hard scheduling constraints.

A 'phase-bench' unit is defined by a group of blocks *b* that have the same horizontal attribute (Z coordinate) and lies inside the same phase design. The new set of blocks precedences is now generated per 'phase-bench' unit, considering its vertical superposition (ie any block *b* in a phase-bench *Z* can only be extracted in a period *t* if all blocks from an overlapping phase-bench *Z-1* were mined in period *t* or *t-1)*. This relationship is later translated to each block, since the block granularity was kept in the SSOMC model. The removal of the previous conical logic is possible because the design and the phase-bench unit consideration already ensure the geotechnical stability.

Among the additional phase-related new constraints, the 'bench-lag' was the only one added as an additional layer in the new set of blocks precedences $\mathcal{O}(b)$, securing its hard constraint definition. It defines the precedence constraint on a phase basis in number of benches (ie what is the number of benches shift (lags) between two phases during extraction). The new 'sink rate' control no longer requires the definition of direct and indirect neighbourhoods with which to calculate the scores to penalise deviations, now being defined as a counter (delta of active Z coordinates) of maximum number of benches being extracted per phase and period. The 'max active phases' constraint defines the maximum number of phases that can simultaneously be extracted within a single period t.

It is important to highlight that this adapted scheduler operates on phase-bench logic for vertical precedencies while maintaining block selectivity and incorporating uncertainty from simulations. It is being named as a phase-bench stochastic scheduler (PBS) but essentially departs from most of the conventional phase-bench models, which follow aggregation of blocks into panels and homogenous mining assumptions.

CASE STUDY

This section describes and examines the application of the method herein described at a large iron ore mine complex. First, the user is referred to a brief description of the mine complex modelling and the materials used (simulations). Next, the conventional mine plan is presented and submitted to a Risk Analysis approach (Godoy and Dimitrakopoulos, 2011) to quantify the risk embedded in the existing LOM production schedule, allowing future fair comparison against DBS and PBS stochastic optimisations that incorporate the simulations. After that, the following subsections analyse the results of the proposed framework described in the Method section.

Using these simulations, annual forecasts of cash flows, ore tonnages, metal tonnages, grades, concentrates are generated for each scenario. Risk profiles are then presented based on 10th, 50th, and 90th percentiles (P_{10}, P_{50}, and P_{90}), which represent non-exceedance probabilities within an 80 per cent confidence interval. Such an approach provides a means of understanding the risk of the mine schedule associated with any production target.

Mine complex overview

The mining operation consists of a conventional open pit bench mining method whose the main operational goal is to fulfill the mill at its capacity. The existing processing plant includes two phases that concentrate hematite and magnetite through gravity and magnetic separation circuits. Figure 2 illustrates the modelling of material flow.

Iron (Fe) and sulfur (S) grades are key elements of interest, since the Fe grade defines the recoveries and a mill feed quality target (minimum 26 per cent Fe) while S grade is a matter of great concern regarding the quality of material feeding the mill (maximum of 80 ppm is allowed). According to the technological processing of materials, some lithologies are pre-defined as waste and the mineralised ones are available for economic evaluation and subsequent destination policy definition.

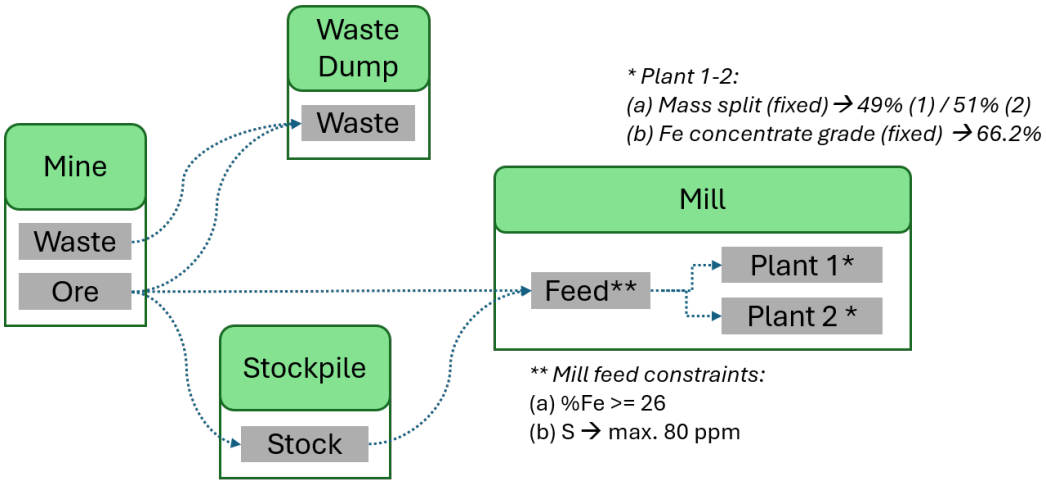


FIG 2 – Mine complex material flow modelling.

The value chain model and the mass flow assume that at any destination the materials are homogenised. Two types of material are feeding the 'Mill' destination – incoming material from the mine or from the stockpile – being blended to produce the mill 'Feed'. This homogeneous material is routed to two plants (a mass split 49 per cent / 51 per cent is predefined between Plant 1/Plant 2). Plant recoveries are a function of %Fe and the contaminant %MgO feed head grades. It is assumed a fixed value for the concentrate Fe grade of 66.2 per cent, therefore the yearly recovery values are the attributes that show the direct impact on Fe-MgO grade uncertainty in the concentrate.

Conditional simulations

Gaussian-based simulation techniques were applied to generate Fe and S grade simulations. The grade values for other elements (such as MgO) were obtained from the estimated block model, using Ordinary Kriging (OK) or inverse distance weighting to a power of three (ID3) methodologies. A total of 20 simulations (combined realisations of Fe-S) were selected based on previous studies (Albor and Dimitrakopoulos, 2009) that demonstrated to be a sufficient number to provide stable life-of-mine optimisation results.

Conventional mine plan and risk analysis

The Risk Analysis procedure was able to quantify the impact of the simulations on the conventional mine plan production forecasts. Figure 3 highlights the differences for Fe (a) and S (b) mill head grades.

For the Fe mill grade the forecast derived from deterministic optimisation closely aligns with the P_{50} values obtained through risk analysis. However, this deterministic forecast does not adequately capture uncertainty or support risk quantification, especially regarding the high risk of falling below the feed grade's minimum threshold of 26 per cent, which is significant from year 7 to 13 (exceeding 30 per cent as summarised by Table 1). This misrepresentation of Fe grade variability has a direct impact on calculated grade recoveries, concentrate tonnages and on the yearly cash flows (as will be summarised in the final subsection).

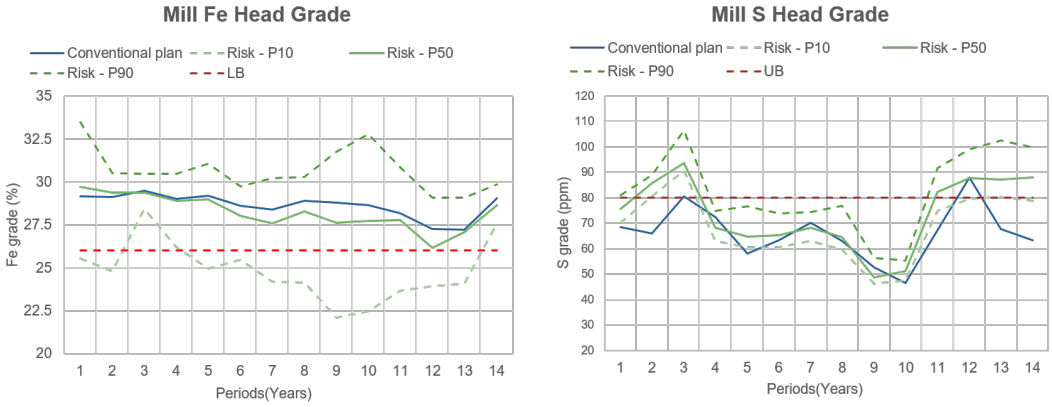


FIG 3 – Conventional and risk analysis mill forecast comparison: (a) Fe grade; and (b) S grade.

TABLE 1

Fe mill target deviation probabilities thresholds by period.

Period	1	2	3	4	5	6	7	8	9	10	11	12	13	14
% probability Fe<26%	15	10	0	5	15	10	30	15	30	35	35	45	40	0

For the sulfur (S) mill head grade – an important contaminant that requires careful control at the processing plant – the analysis indicates a high risk of exceeding the preferred upper bound limit in the first four years and last five years of production. The risk analysis highlights that the deterministic schedule's estimated grade does not adequately capture the uncertainty or quantify the risks associated with S grade variations feeding the mill.

Stochastic DBS optimisation and phase design

The stochastic DBS optimisation generated a semi-operational mine sequence illustrated in Figure 4a by period-coloured filled blocks for a typical mine cross-section. It is possible to verify that the DBS sequence was successful when producing an aggregated block sequence that rationally respects 3D components required for the global optimisation, such as sink rake (eight benches per period) and mining width, despite some scattered blocks and local deviations. The main production target (Fe mill grade) was met as exhibited in Figure 5a, despite a 1-year LOM reduction.

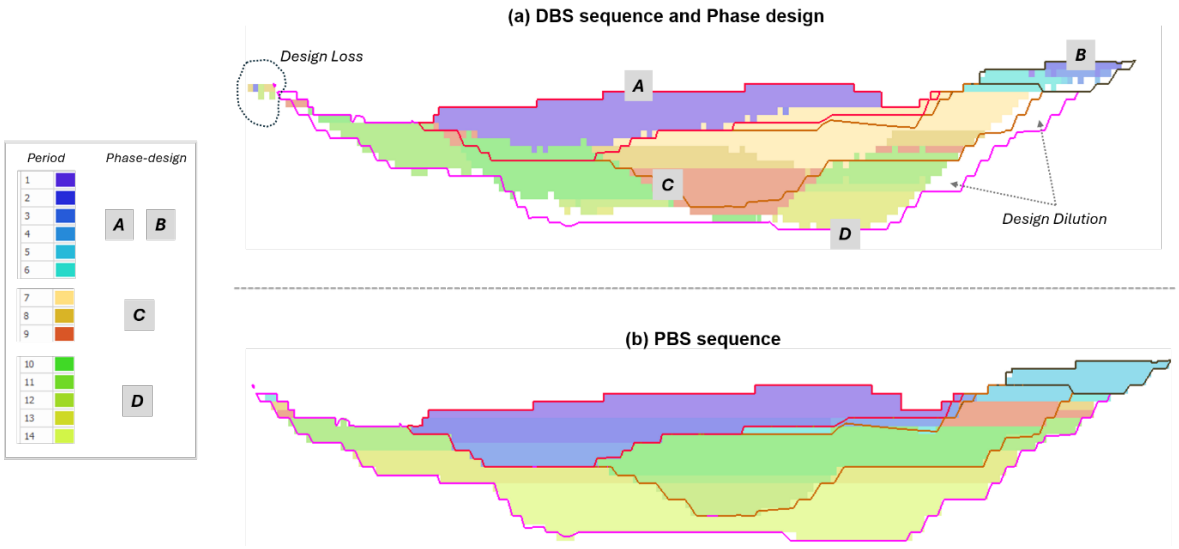
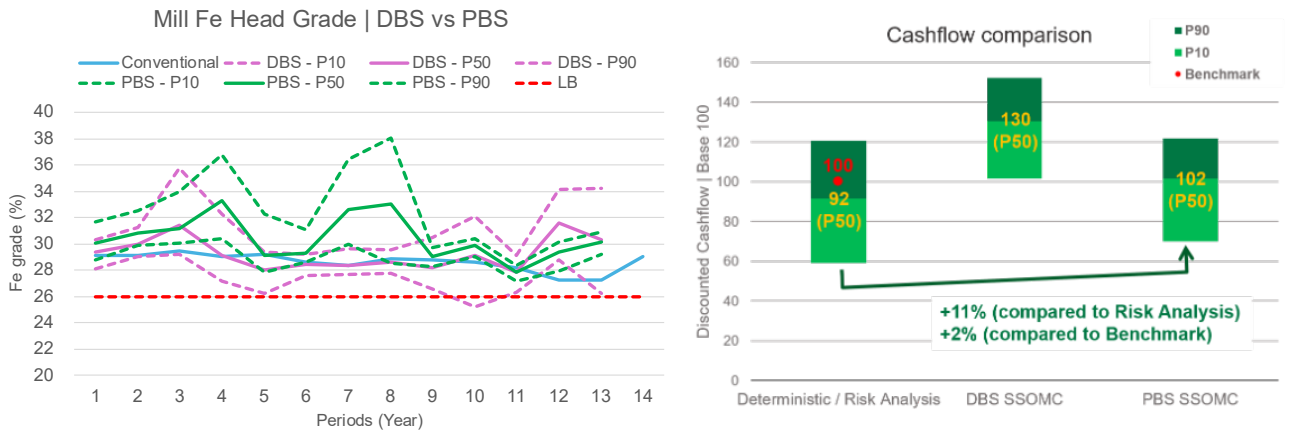


FIG 4 – Typical mine cross-section with phase-design and (a) DBS; and (b) PBS sequences.



(a) **(b)**

FIG 5 – Stochastic DBS and PBS: (a) Fe mill target profiles; and (b) NPV comparison.

In order to refine the DBS stochastic schedule with operational constraints, a block-grouping approach defined the phases regions for detailed mine design. This block aggregation was guided by spatial affinity and periods proximity analysis. In Figure 4a the solids boundaries over the cross-section illustrate some of the phases designed and its spatial adherence to the DBS grouped periods of extraction. Periods 1 to 6 mostly define phases A and B (separated due to spatial disconnection), while phase C fits in general to periods 7 to 9 and phase D broadly accommodates period 10 onwards. It is worth mentioning that the mine design process unavoidably leads to loss and dilution, also highlighted in Figure 4a, since it contemplates refined operational constraints such as ramp access and local minimum mining width.

Stochastic phase-bench optimisation

The phases designed in the previous step were used to establish the vertical precedences between phases-benches and a strict sink rate control was set (eight benches per period, as for the DBS). The Figure 4b illustrates the sequence of extraction produced by the stochastic phase-bench optimisation application.

By comparing the DBS and PBS sequences on Figure 4 it can be seen that for most of the periods there is a consistent match between the areas that both stochastic optimiser is selecting to mine, in both cases always taking into account the NPV maximisation and risk-management from simulations to keep the targets under control, as illustrated in the Fe mill head grade profile comparison in Figure 5a. Furthermore, the PBS stochastic sequence stands out for remarkably more acceptable mine shapes, increasing the operational confidence in its execution.

Figure 5b summarises the NPV for all cases studied above. The initial expected project value (conventional mine plan, deterministic and operational) was set as the baseline (100 per cent) and the evaluation of this very same plan throughout the uncertainty and variability of the simulations (Risk Analysis) demonstrates a probable (P_{50}) decrease by -8 per cent. When allowing a SIP optimiser to generate a new schedule in block-scale and managing the risk of not achieving production targets (Stochastic DBS optimisation) it was verified a maximum probable (P_{50}) NPV increase around +30 per cent, which was later on confirmed as a potential increase (P_{50}) of +2 per cent when a phase-bench stochastic optimiser refined the results. If the stochastic PBS is analysed against the Risk Analysis, then the probable (P_{50}) value created is around +11 per cent.

CONCLUSIONS

The results highlight critical gaps in the deterministic model, particularly its inability to account for grade uncertainty. The application of simulated models allowed to quantify the risk embedded in the conventional mine plan as a probable -8 per cent decrease in NPV and a 30 per cent probability of failing to meet the mill Fe grade target in the second half of project LOM.

The proposed stochastic framework for mine optimisation, which includes a three-step approach (stochastic DBS → phase definition and design → stochastic PBS) demonstrates ability to maximise

NPV (+2 per cent from the expected deterministic benchmark or +11 per cent compared to the Risk Analysis) while managing the risk (all periods above the mill Fe grade target) and respecting critical operational parameters (ensured by phase-design and PBS mineable controls), without sacrificing the benefits of block granularity.

These advantages result in a more reliable long-term plan that ensures the operational practicality required for large-scale mining complexes, representing a significant step towards gaining industry confidence in advanced strategic planning technologies.

Future extensions of the methodology might include an automated schedule-based phase generation to remove any subjective component in the process, such as a Breadth First Search algorithm (BFS) adapted approach suggested by Farmer and Dimitrakopoulos (2018). Another topic of interest for research is the improvement of the mineability constraints for the DBS solver, potentially reducing even more the 'gap' of value creation and 'operationalisation' demonstrated in the results.

REFERENCES

Albor, C and Dimitrakopoulos, R, 2009. Stochastic mine design optimisation based on simulated annealing: pit limits, production schedules, multiple orebody scenarios and sensitivity analysis, *Mining Technology*, 118(2):79–90. https://doi.org/10.1179/037178409X12541250836860

Aras, C, Dagdelen, K and Johnson, T B, 2019. Generating pushbacks using direct block mine production scheduling algorithm, in *Mining Goes Digital*, pp 426–436. https://doi.org/10.1201/9780429320774-49

Benndorf, J and Dimitrakopoulos, R, 2013. Stochastic long-term production scheduling of iron ore deposits: integrating joint multi-element geological uncertainty, *Journal of Mining Science*, 49(1):68–81. https://doi.org/10.1134/S1062739149010097

Bley, A, Boland, N, Fricke, C and Froyland, G, 2010. A strengthened formulation and cutting planes for the open pit mine production scheduling problem, *Computers and Operations Research*, 37(9):1641–1647. https://doi.org/10.1016/j.cor.2009.12.008

Caccetta, L and Hill, S, 2003. An application of branch and cut to open pit mine scheduling, *Journal of Global Optimization*, 27(2–3):349–365. https://doi.org/10.1023/A:1024835022186

Dimitrakopoulos, R and Lamghari, A, 2022. Simultaneous stochastic optimization of mining complexes – mineral value chains: an overview of concepts, examples and comparisons, *International Journal of Mining, Reclamation and Environment*, 36(6):443–460. https://doi.org/10.1080/17480930.2022.2065730

Dimitrakopoulos, R and Ramazan, S, 2004. Uncertainty-based production scheduling in open pit mining, *Transactions of the Society for Mining, Metallurgy and Exploration*, 316:106–112.

Dimitrakopoulos, R, 2011. Stochastic optimization for strategic mine planning: a decade of developments, *Journal of Mining Science*, 47:138–150. https://doi.org/10.1134/S1062739147020018

Dimitrakopoulos, R, Farrelly, C T and Godoy, M, 2002. Moving forward from traditional optimization: grade uncertainty and risk effects in open-pit design, *Mining Technology*, 111(1):82–88. https://doi.org/10.1179/mnt.2002.111.1.82

Farmer, I and Dimitrakopoulos, R, 2018. Schedule-based pushback design within the stochastic optimisation framework, *International Journal of Mining, Reclamation and Environment*, 32(5):327–340. https://doi.org/10.1080/17480930.2017.1289606

Gershon, M, 1987. Heuristic approaches for mine planning and production scheduling, *International Journal of Mining and Geological Engineering*, 5(1):1–13. https://doi.org/10.1007/BF01553529

Godoy, M and Dimitrakopoulos, R, 2011. A risk quantification framework for strategic mine planning: method and application, *Journal of Mining Science*, 47:235–246. https://doi.org/10.1134/S1062739147020109

Goodfellow, R and Dimitrakopoulos, R, 2016. Global optimization of open pit mining complexes with uncertainty, *Applied Soft Computing*, 40:292–304.

Goodfellow, R and Dimitrakopoulos, R, 2017. Simultaneous stochastic optimization of mining complexes and mineral value chains, *Mathematical Geosciences*, 49:341–360. https://doi.org/10.1007/s11004-017-9680-3

Hustrulid, W A, Kuchta, M and Martin, R, 2013. *Open Pit Mine Planning and Design* (CRC Press – Taylor and Francis Group: London).

Jelvez, E, Morales, N, Nancel-Penard, P, Peypouquet, J and Reyes, P, 2016. Aggregation heuristic for the open-pit block scheduling problem, *European Journal of Operational Research*, 249(3):1169–1177. https://doi.org/10.1016/j.ejor.2015.10.044

Johnson, T B, 1968. Optimum open pit mine production scheduling, PhD thesis (Operations Research Department, University of California: Berkeley).

Johnson, T B, 1969. Optimum production scheduling, in *Proceedings 8th International Symposium on Computers and Operations Research in the Mineral Industry – A Decade of Digital Computing in the Mineral Industry* (ed: A Weiss), pp 539–562 (AIME: New York).

Khalokakaie, R, Dowd, P A and Fowell, R J, 2000. Lerchs-Grossmann algorithm with variable slope angles, *Mining Technology*, 109(2):77–85.

Lerchs, H and Grossmann, I F, 1965. Optimum design of open pit mines, *Canadian Institute of Mining Bulletin*, 58:47–54.

Letelier, O R, Espinoza, D, Goycoolea, M, Moreno, E and Muñoz, G, 2020. Production scheduling for strategic open pit mine planning: a mixed-integer programming approach, *Operations Research*, 68(5):1425–1444. https://doi.org/10.1287/opre.2019.1965

Monkhouse, P H L and Yeates, G, 2007. Beyond naive optimisation, in *Orebody Modelling and Strategic Mine Planning*, Spectrum Series 14, pp 3–18 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Montiel, L and Dimitrakopoulos, R, 2015. Optimizing mining complexes with multiple processing and transportation alternatives: an uncertainty-based approach, *European Journal of Operational Research*, 247(1):166–178.

Pimentel, B S, Mateus, G R and Almeida, F A, 2010. Mathematical models for optimizing the global mining supply chain, in *Intelligent Systems in Operations*: *Methods, Models and Applications in the Supply Chain* (ed: B Nag), pp 133–163 (IGI Global). http://dx.doi.org/10.4018/978-1-61520-605-6.ch008

Ramazan, S and Dimitrakopoulos, R, 2013. Production scheduling with uncertain supply: a new solution to the open pit mining problem, *Optimization and Engineering*, 14(2):361–380. http://dx.doi.org/10.1007/s11081-012-9186-2

Saliba, Z and Dimitrakopoulos, R, 2019. Simultaneous stochastic optimization of an open pit gold mining complex with supply and market uncertainty, *Mining Technology*, 128(4):216–229. https://doi.org/10.1080/25726668.2019.1626169

Incorporating geological, equipment, and geotechnical uncertainty into mine planning

I Minniakhmetov[1] and Z Fletcher[2]

1. Principal Geoscientist, BHP, Perth WA 6010. Email: ilnur.minniakhmetov@bhp.com
2. Principal EMP, BHP, Perth WA 6010. Email: zory.fletcher@bhp.com

ABSTRACT

Mine plans often under-represent three dominant operational risks: variability in the *in situ* resource, uncertainty in equipment performance, and geotechnical hazards. This paper presents a practical framework that incorporates all three into a single planning environment and quantifies their impacts in decision-ready terms. Multiple conditional simulation realisations, equipment performance distributions, and simulated geotechnical failures are combined and scheduled under identical constraints as the base plan. Cross-scenario analysis summarises outcomes as $P_{10}/P_{50}/P_{90}$ envelopes for key metrics including tonnes, qualities, rehandling, and schedule slippage. A case study on a modified iron ore data set demonstrates a super-additive effect: modelling any single source alone produces modest risk shifts, but their interaction widens envelopes materially, increases rehandling likelihood, and introduces meaningful tails in timing and quality distributions. Joint treatment is necessary to reveal compounding risks invisible in single-factor stress tests. The framework informs strategic decisions—stockpile footprint allocation, fleet sizing, pit access sequencing—that sit beyond the two-year budget planning horizon, enabling leaders to set commitments at explicit confidence levels and define trigger points for operational adjustments.

INTRODUCTION

Mine plans are typically built on deterministic resource models and fixed equipment parameters. This obscures three dominant operational risks: variability in the *in situ* resource, uncertainty in equipment performance, and geotechnical hazards. When these risks materialise, operations face unplanned rehandling, crusher throughput losses, quality breaches, and fleet constraint violations—consequences invisible when commitments are set. The purpose of incorporating uncertainty is not to generate hundreds of plans for execution, but to stress-test the robustness of the current plan, identify risks that could prevent its implementation, and define buffers to protect bottlenecks and the critical path.

Previous work (Minniakhmetov and Serjeantson, 2025) demonstrated a methodology for consuming geological uncertainty into planning by emulating the full cycle—conditional simulation, grade control sampling, model updates, and rescheduling—across 100 realisations for a historically mined iron ore operation comprising multiple orebodies feeding crushers and ore handling plants.

This paper extends the framework by combining geological, equipment, and geotechnical uncertainty in a single planning environment and demonstrating their super-additive interaction. Critically, the mitigations revealed—additional stockpile footprint near active pits, fleet procurement, accelerated pit access—require decisions beyond the two-year budget planning horizon, in the strategic planning space with multi-year lead times. These studies must therefore be conducted proactively to inform strategic decisions before constraints materialise. The methodology applies equally to tactical and strategic planning horizons.

METHODOLOGY

Geological uncertainty

Conditional simulations of Fe, Al_2O_3, SiO_2, P, and loss on ignition (LOI) are generated using hierarchical truncated pluri-Gaussian simulation for categorical variables and turning bands simulation for continuous grades (Williams and Minniakhmetov, 2025). Ten realisations are propagated through the Alastri Tactical Scheduler (ATS) reserving process to produce scenario-specific reserve files that feed directly into scheduling. Each realisation represents a plausible *in situ*

reality; the ensemble tests whether the plan remains feasible across the range of geological outcomes.

Equipment uncertainty

Historical loader utilisation and availability (UA) data are fitted with statistical distributions (typically Gaussian, approximately ±8 per cent relative perturbation). A custom importer within ATS perturbs loader UA values per period, generating ten equipment scenarios that test whether the plan can absorb capacity shortfalls without breaching constraints.

Geotechnical uncertainty

Geotechnical risk is incorporated by simulating pit-wall failure events that remove access to planned mining areas for defined durations, testing spatial resilience and the ability to maintain production from alternative sources. In this work, the geotechnical failure scenario is included to demonstrate the technical feasibility of incorporating such events into the probabilistic scheduling framework. The specific failure geometry and duration used are illustrative; in practice, these inputs would be informed by structural geology and geotechnical engineering assessments of slope stability risk for each pit.

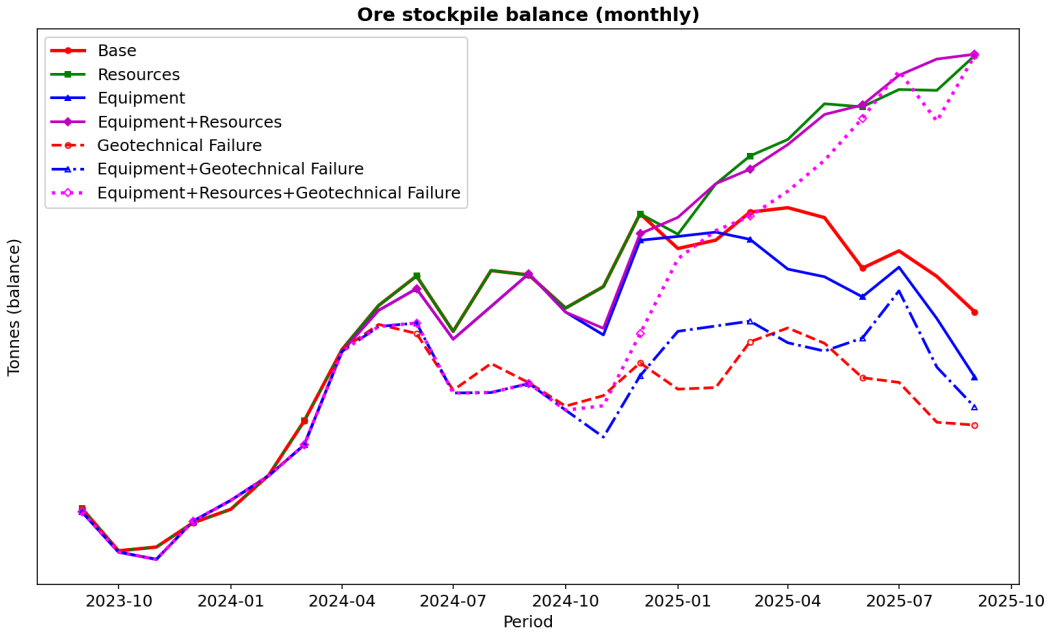
Combined scenarios

All three uncertainty sources are combined within ATS. Each scenario is scheduled using identical constraints and objectives as the base plan. Batch scheduling produces detailed schedule reports, analysed as $P_{10}/P_{50}/P_{90}$ envelopes for crusher tonnes, grades, rehandling, and truck hours. The goal is to quantify where the plan is at risk and what buffers—stockpile capacity, fleet margin, schedule float—are needed to protect it.

RESULTS

The case study applies the framework to a 24 month plan for a historically mined iron ore hub comprising multiple orebodies, stockpiles, crushers, and ore handling plants. The underlying geological and operational data are real; however, selected parameters have been modified to illustrate potential failure modes, including blending-grade shortfall scenarios. All failures and constraint breaches presented are simulated outcomes generated by the framework, not historical operational events. Three variability sources were tested independently and in combination.

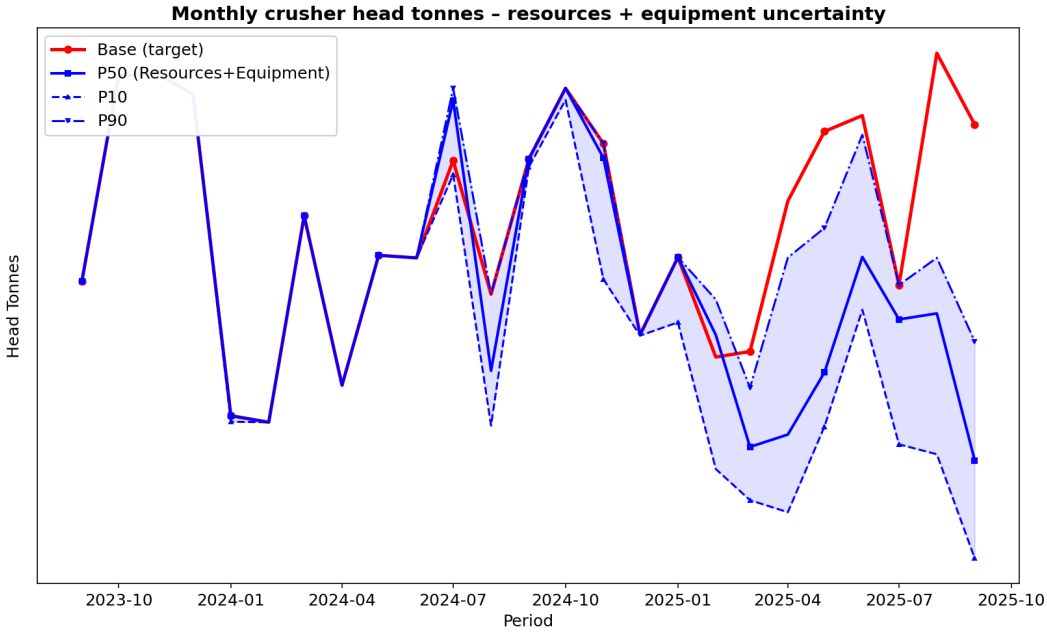
Figure 1 shows the ore stockpile balance under different uncertainty combinations. Solid lines represent scenarios without geotechnical failure; dashed lines represent scenarios that include a simulated pit-wall failure. All dashed-line scenarios exhibit lower ore stockpile tonnages, threatening strategic buffer objectives. While the risk of breaching the minimum stockpile threshold within this two-year window is low, the reduced ending balance sets the starting position for the subsequent planning cycle—which would be at material risk if future plans assume stockpile levels approximately twice those shown under uncertainty scenarios.



Note: Absolute tonnage values removed due to commercial sensitivity.

FIG 1 – Ore stockpile balance under geological, equipment, and geotechnical uncertainty scenarios. Dashed lines (geotechnical failure scenarios) show reduced buffer levels. Absolute values removed due to commercial sensitivity.

Figure 2 shows monthly crusher head tonnes under combined resource and equipment uncertainty. The base plan (red) represents the committed target. In Year 2, the P_{50} falls below target due to insufficient blending-grade material—a shortage driven by the smoothing inherent in kriging that overstates blending-grade tonnages relative to conditional simulations.



Note: Absolute tonnage values removed due to commercial sensitivity.

FIG 2 – Monthly crusher head tonnes under resource and equipment uncertainty. P_{50} falls below target in Year 2 due to blending-grade shortage. Absolute values removed due to commercial sensitivity.

When all three sources are applied together, the combined effect is super-additive. Variability-driven rehandling increases truck demand beyond fleet capacity, amplifying the risk of not delivering

productive movement—a critical path constraint. These compounding interactions are invisible when each source is tested in isolation. The plan appears feasible under any single perturbation but is vulnerable to simultaneous ordinary variability across multiple dimensions.

STRATEGIC IMPLICATIONS

The mitigations required—proximal stockpile footprint (physical ground area near active pits for additional stockpile construction), fleet procurement, new pit access—sit beyond the two-year budget planning horizon, in the strategic planning space with multi-year lead times. This framework provides quantified risk evidence to support those decisions proactively. It enables planners to:

- identify which buffers protect the critical path and size them against quantified risk.
- set commitments at an explicit confidence level (eg P_{70}–P_{80}) with trigger points for adjustments.
- rank uncertainty sources by impact on plan robustness to prioritise mitigation investment.
- communicate risk envelopes within corporate assurance and budgeting processes.

The workflow is transparent, repeatable, and contained within commercially available software, offering a consistent probabilistic lens from tactical through strategic planning horizons.

ACKNOWLEDGEMENTS

The author acknowledges the support of the Resource Engineering, Geoscience, and Mine Planning teams, the Resource Centre of Excellence leadership at BHP, and Micromine for technical collaboration on ATS.

REFERENCES

Minniakhmetov, I and Serjeantson, J, 2025. Consuming in situ uncertainty into mine planning processes – from reserve calculation to dynamic decision-making, in Proceedings of the Mineral Resource Estimation Conference 2025 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Williams, C and Minniakhmetov, I, 2025. In-situ uncertainty quantification and implications for risk mitigation – a case study from an iron ore operation, in Proceedings of the Mineral Resource Estimation Conference 2025 (The Australasian Institute of Mining and Metallurgy: Melbourne).

From risk assessment to stochastic optimisation – a systematic approach to de-risk an operating gold mine production schedule

A Rispoli[1], L C Andrade[2] and D Alegre[3]

1. Senior Director, Mining and Integrated Planning, IAMGOLD, Toronto ON, Canada. Email: adrienne_rispoli@iamgold.com
2. Mine Planning Consultant, KPI Mining Solutions, Montreal QC, Canada. Email: laura.andrade@kpimining.com
3. Mine Planning Consultant, KPI Mining Solutions, Montreal QC, Canada. Email: douglas.alegre@kpimining.com

INTRODUCTION

Traditional mine planning relies on a single estimated (kriged) orebody model to generate deterministic production schedules. While computationally convenient, this approach smooths grade variability, leading to production forecasts that are unlikely to be achieved in reality (Ravenscroft, 1992). As gold prices reach historic highs and deposits become increasingly complex, a more rigorous stochastic approach is required to quantify geological uncertainty and optimise long-term value (Leite and Dimitrakopoulos, 2014).

This study conducted in 2025 is structured in three stages applied to Essakane Gold Mine:

1. Risk assessment of the current production year to validate against actual performance.
2. Risk assessment across the life-of-mine (LOM) to capture the impact of uncertainty on long-term outcomes.
3. Stochastic optimisation to maximise project value while managing risk.

A set of stochastic simulations of Au grades is used throughout to represent geological uncertainty.

The Essakane Mine was selected for this study because of its historical complexity in grade predictability. Reconciliation outcomes have been highly variable, with results ranging from significantly positive to negative variances between modelled grades and mill production (F3 reconciliation factor).

GEOSTATISTICAL SIMULATIONS

A set of orebody realisations was generated using geostatistical simulation conditioned to Essakane's drill hole data. Comparison with the kriged block model revealed a pronounced smoothing effect: simulated models reproduce the global statistics of the drill hole data set more accurately. Where grade-control data incorporating production blastholes and supplementary reverse circulation (RC) drilling were available, simulated models showed a markedly closer match to grade control than the estimated model. This demonstrates that simulations provide a more realistic representation of grade variability and, critically, quantify the range of geological uncertainty for mine planning decisions.

RISK ANALYSIS

The risk analysis method applies the existing deterministic production schedule, extraction sequence and ore/waste volumes, to the set of simulated orebody models, producing stochastic production forecasts expressed as P_{10}, P_{50}, and P_{90} risk profiles.

Current year (2025)

Applied to 2025 in the study, the risk analysis found that the deterministic model overestimated contained gold above cut-off by an average of 32 per cent. Originally defined ore volumes incorporate substantial low-grade material based on simulated models, suggesting higher dilution than forecast (Figure 1). These findings are corroborated by operational results: Q2 2025 actuals were below target due to a temporary grade under performance associated with phase sequence (IAMGOLD Corporation, 2025). The full year 2025 F3 results confirmed fewer tonnes, lower grade,

and fewer ounces, thus a negative reconciliation performance consistent with the simulation-identified overestimation.

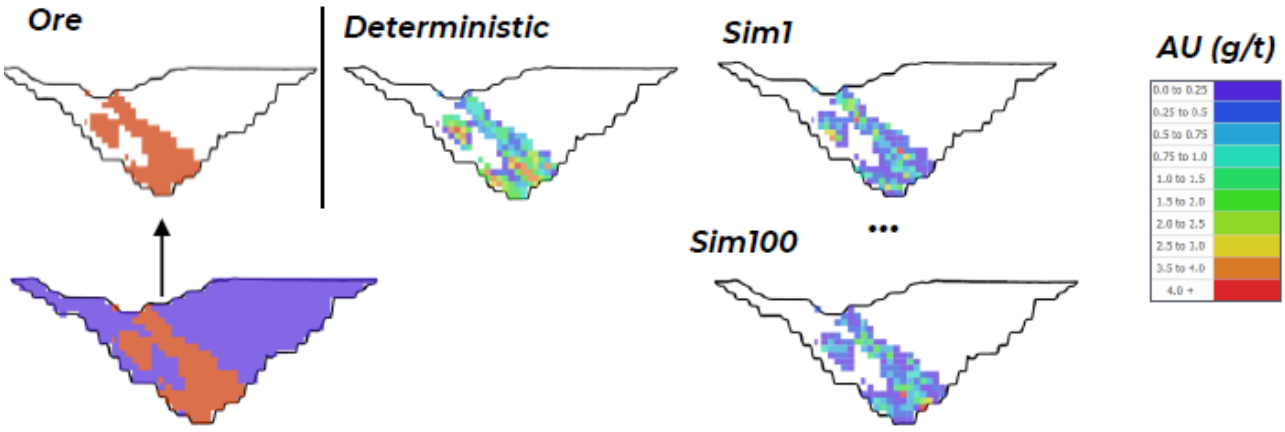


FIG 1 – Dilution in deterministic ore volumes.

Life-of-mine plan (2026–2029)

Extending the risk analysis to the full LOM plan revealed persistent overestimation of recovered gold ounces and head grades across all planned years. In the study, P_{10} to P_{90} risk profiles with the simulation consistently fall below the deterministic forecast, meaning the deterministic plan's net present value (NPV) substantially overstates the most likely outcome. Critically, the analysis in the study identified a risk of negative discounted cash flow in 2028 when geological uncertainty is appropriately accounted for (Figure 2). Total costs show minimal variability across simulations, confirming that economic risk is almost entirely a function of gold grade uncertainty on the revenue side. Identifying the areas of geological uncertainty can support further drilling program targets.

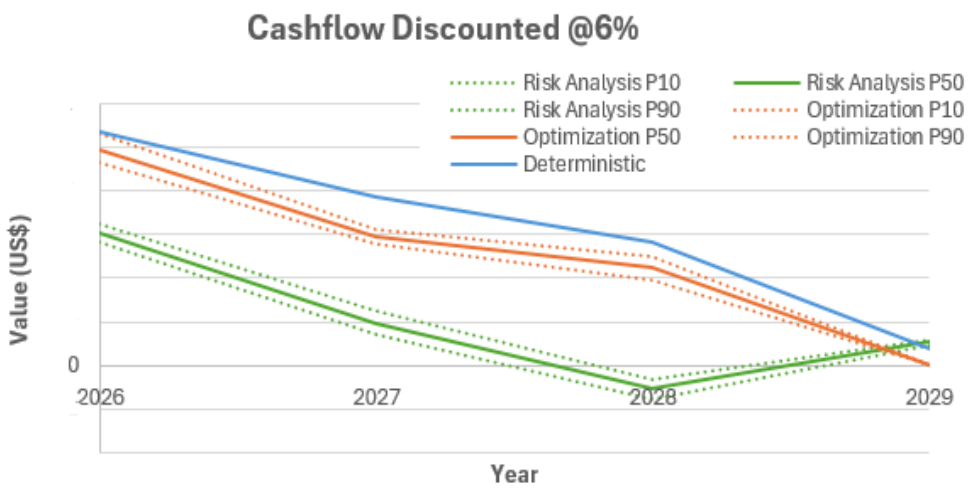


FIG 2 – Risk profile of the deterministic LOM plan and stochastic LOM plan.

STOCHASTIC OPTIMISATION

Conventional mine planning optimises sequentially and independently: pit limits, phase design, extraction sequence, cut-off grades, stockpiling, and processing destinations are resolved in separate steps with separate tools. Stochastic optimisation replaces this fragmented workflow with a simultaneous optimisation and considers all orebody simulations concurrently. Rather than optimising against a single smooth model, the stochastic optimiser identifies schedules that are resilient across the full range of geological scenarios, directly producing probabilistic production forecasts and risk-managed outcomes (Goodfellow and Dimitrakopoulos 2016, 2017).

Base case results

The stochastic base case optimisation produced a materially different extraction sequence compared to the deterministic plan. The stochastic schedule prioritises mining areas closer to surface and with higher grade certainty, reducing both geological and economic exposure. The most significant outcome was the elimination of the previously identified risk of negative discounted cash flow in 2028, while simultaneously achieving an increase in NPV relative to the risk-assessed deterministic forecast as displayed in Figure 2. The optimised cut-off grade policy differs per annum and per material type, reflecting the optimiser's ability to dynamically adapt to grade uncertainty rather than applying a fixed, deterministic threshold.

Growth scenario and gold price sensitivity

A growth scenario was evaluated by running the stochastic optimiser at gold prices of US$2500, US$3000, and US$3500/oz, reflecting the current elevated price environment. Higher gold prices enable a lower cut-off grade policy, unlocking additional material as economic ore and increasing reserve potential through 2032. The sensitivity analysis in terms of the NPV demonstrates that the opportunity cost at elevated prices significantly outweighs the downside risk; even if realised prices fall below the most optimistic assumption, the production schedule remains financially robust. The growth scenario, after accounting for associated capital expenditure, delivers a higher average NPV than the base case, validating the strategic case for reserve expansion at current prices.

CONCLUSIONS

This study demonstrates that deterministic mine planning risks consistently overestimating recoverable gold at Essakane with the previous block model, and that the magnitude of this overestimation was quantifiable through stochastic simulation. The predicted overestimation trend in Au ounces identified through risk analysis for 2025 was confirmed by Essakane's reconciliation results. This key finding enables the operation to strategically target future diamond drilling programs in areas of geological uncertainty and to refine domaining methodologies in subsequent resource model updates. These improvements are expected to reduce in the mine plan local grade variability within the orebody and enhance overall grade predictability in future block model revisions.

Stochastic optimisation resolves the core limitation of conventional planning by simultaneously optimising all mine planning decisions across multiple geological scenarios. The resulting schedules are more conservative where uncertainty is high and more aggressive where grade confidence is strong, confirming the risk management capability of the approach. The growth scenario suggests that the current gold price environment presents a strategic opportunity to expand the reserve base without increasing financial risk.

ACKNOWLEDGEMENTS

IAMGOLD approved the publication of the results presented in this paper. The authors thank the Essakane and IAMGOLD Corporate Technical Service teams, KPI Mining Solutions Mine Planning team, and Mine Mind Geostatistics team for their collaboration.

REFERENCES

Goodfellow, R and Dimitrakopoulos, R, 2016. Global optimization of open pit mining complexes with uncertainty, *Applied Soft Computing*, 40:292–304. https://doi.org/10.1016/j.asoc.2015.11.038

Goodfellow, R and Dimitrakopoulos, R, 2017. Simultaneous stochastic optimization of mining complexes and mineral value chains, *Mathematical Geosciences*, 49(3):341–360. https://doi.org/10.1007/s11004-017-9680

IAMGOLD Corporation, 2025, August 7. IAMGOLD reports second quarter 2025 results. Available from: <https://www.iamgold.com/english/investors/news-releases/news-releases-details/2025/iamgold-reports-second-quarter-2025-results/>

Leite, A and Dimitrakopoulos, R, 2014. Stochastic optimization of mine production scheduling with uncertain ore/metal/waste supply, *Mining Science* and Technology, 24(6):755–762. https://doi.org/10.1016/j.ijmst.2014.10.004

Ravenscroft, P J, 1992. Risk analysis for mine scheduling by conditional simulation, *Transactions of the Institution of Mining and Metallurgy (Section A),* 101:A104–A108.

Strategic mine planning – value, governance and decision-making

Strategic value driven planning through integrated multi-mine optimisation

F Damasceno[1], J Sega[2] and P Silva[3]

1. Mine Planning Specialist, Vale, Nova Lima, MG 34006-068, Brazil. Email: felipe.damasceno@vale.com
2. Technical Manager, Vale, Nova Lima, MG 34006-068, Brazil. Email: juliano.sega@vale.com
3. Technical Manager, Vale, Nova Lima, MG 34006-068, Brazil. Email: pedro.silva@vale.com

INTRODUCTION

Strategic mine planning in iron ore has traditionally been structured around local decisions, often solving each asset or production system with limited visibility of downstream interactions. Integrated mine planning has already improved this picture by representing mine, plant, waste and product constraints in a single optimisation environment. This shifts the planning problem towards a portfolio-level decision context, where trade-offs must be made between project options, capital allocation choices and the response of the system to changes in prices, premiums and technical conditions. This work presents a practical evolution from integrated mine planning to global optimisation, in which multiple iron ore operations are represented in a unified strategic model designed to support value-driven decisions across the broader business rather than within individual assets.

METHODOLOGY

The proposed framework represents a multi-mine value chain in a single optimisation model spanning more than 25 mines and over 30 processing facilities, integrated through shared logistics systems and product streams. The problem is formulated as a mixed-integer linear program (MILP) that maximises the discounted economic value of the system. Continuous variables describe material flows, production levels and blending, while binary variables represent discrete strategic choices, such as whether a capital project is activated and when it starts. Capital projects are therefore modelled as value-driven triggers, activated only when they improve the global objective. The formulation captures production, processing, quality, infrastructure and project start-up constraints across hundreds of origin-destination flows and structures. Figure 1 summarises the framework's main drivers, integrated value-chain model and strategic outputs.

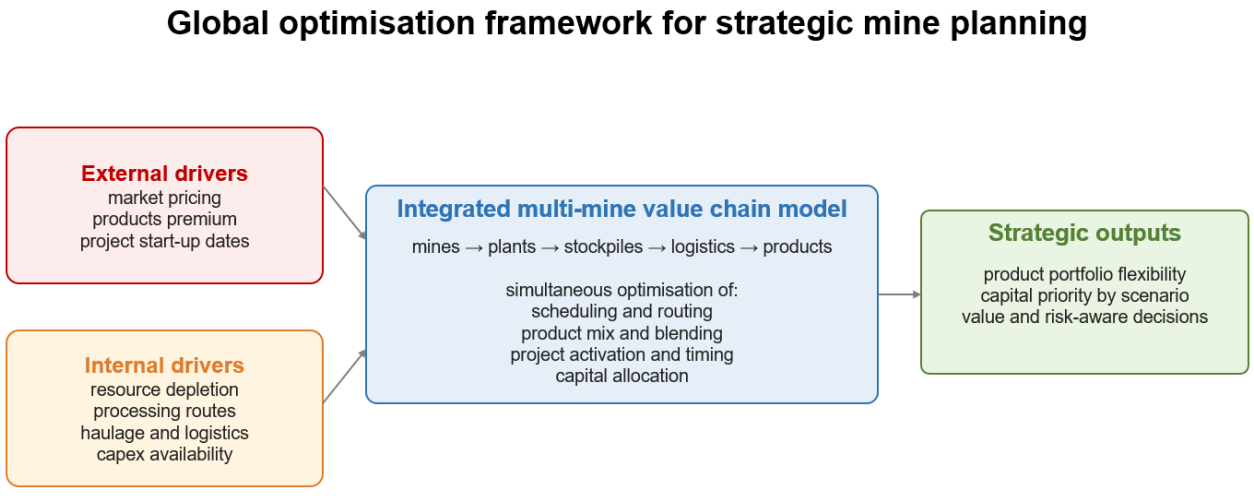


FIG 1 – Conceptual representation of the global optimisation framework used to link uncertainty, multi-mine flows and strategic decisions.

Because the model combines a large number of binary decisions with a fine-grained flow network over a 30-year horizon on an annual basis, it belongs to a class of problems that cannot be solved reliably by heuristics or spreadsheet-based methods. The framework is solved with Gurobi, a state-of-the-art commercial MILP solver that consistently ranks among the fastest available for large-

scale problems. Gurobi applies a branch-and-cut strategy combined with advanced presolve, cutting planes and parallel search, and can run multiple algorithms concurrently so that the fastest method prevails. This provides a guaranteed optimal solution, or a solution with a measurable optimality gap when stopped early, which is essential for a model of this scale and complexity.

The framework is intended for scenario-based decision support rather than one-off schedule generation. Rather than optimising under a single set of assumptions, it is used to learn the value of choices across many combinations of market and technical conditions, including changes in product pricing, premium outlooks, project timing scenarios, processing configurations and capital competition. It also captures uncertainty associated with asset availability, particularly when access to future mining areas depends on project timing and readiness, which has become more relevant as sustainability requirements and decarbonisation increase the complexity and timelines of project execution. High-performance virtual computing environments make the timely evaluation of these scenarios operationally viable.

CASE STUDY APPLICATION

The framework was applied to a business case focused on portfolio flexibility and value maximisation under uncertainty. A broad set of product-flexibility and project alternatives was screened and an initial subset was taken forward for integrated scenario evaluation. These alternatives included route changes, processing flow sheet configurations and project timing options, assessed based on their contribution to overall value generation and capital efficiency. Rather than disclosing detailed asset-level assumptions, the case study is presented as an illustration of how a global optimisation framework can translate uncertainty into strategic choices. Scenario sets were constructed to compare different market outlooks and operating conditions, with particular attention to the implications of weaker premium environments and to the effects of variations in project timing on product mix and mine scheduling.

RESULTS AND DISCUSSION

The case study showed that global optimisation leads to a different decision logic relative to conventional planning. Under a conventional approach, projects and product routes tend to be assessed largely on local merit, whereas the global model re-prioritises the same set of alternatives according to their contribution to system-wide value, changing which projects are activated and when. The most important shift was not simply in total production, but in the composition and timing of value delivery.

A central output of the approach is the distinction between robust and scenario-dependent alternatives. Robustness is assessed by how consistently an alternative is selected across the scenario set: alternatives activated under most combinations of assumptions are classified as robust, while those activated only under specific premises are classified as scenario-dependent. This measure moves the decision from a single optimal plan towards an understanding of which choices create value reliably and which depend on favourable conditions. In practice, alternatives focused on premium products proved more sensitive to weaker premium environments, while lower-capital flexibility options remained selected across a broader range of scenarios and were therefore more robust.

This is particularly relevant in integrated iron ore portfolios where premium products are dependent on the development of specific mining areas and on the interaction between mine availability, downstream flexibility and market conditions. In that context, global optimisation acts as a decision-support framework that reveals how variations in market conditions and operational constraints can simultaneously propagate through product strategy, project timing and capital allocation. By evaluating alternatives inside a common optimisation environment, the planning process moves away from isolated what-if studies and towards a more coherent comparison of competing choices, exposing hidden interactions between assets and clarifying where capital is genuinely adding value.

CONCLUSIONS

Global optimisation represents a meaningful evolution of integrated mine planning. Its main contribution is not only a broader model scope, but a stronger decision framework for situations where value depends on the interaction between market conditions, product positioning, project and operational uncertainties and capital choices across multiple mines. By formulating the problem as a large-scale MILP and solving it with a state-of-the-art solver, the approach can distinguish robust alternatives from scenario-dependent ones, identify when premium-focused pathways lose competitiveness and show how product mix and project timing should adapt to changing assumptions. For organisations managing large, interconnected mining systems, this provides a practical basis for more disciplined strategic planning under uncertainty.

Value generation in mining – does mineral resource management still matter?

P Knights[1], H Burgers[2] and C Aji[3]

1. Professor and Discipline Leader for Mining, School of Mechanical and Mining Engineering, The University of Queensland, St Lucia Qld 4072. Email: p.knights@uq.edu.au
2. Associate Professor and Head of the Strategy and Entrepreneurship Discipline, School of Business, The University of Queensland, St Lucia Qld 4072. Email: h.burgers@uq.edu.au
3. Graduate of Advanced Finance and Economics Program, The University of Queensland, St Lucia Qld 4072. Email: c.aji@student.uq.edu.au

ABSTRACT

This paper revisits a 2009 study which analysed the share and reported mineral reserve histories of a select group of 14 mining companies over the period from 2000 to 2008. With the exception of one or two cases, the study concluded that total shareholder returns (TSR) generally followed total reserve increment (TRI). In this revisited study, the share and reported mineral reserve histories for 15 global mining companies were analysed over the period 2010 to 2024. The results indicate that total shareholder return does not necessarily depend on mineral reserve growth. With the exception of First Quantum, TSR performance of companies can be classified into three groups; the first category has total shareholder return commensurate with total reserve increment (this includes Glencore. South32, Gold Fields, Rio Tinto and Antofagasta Minerals); a second category ('mid-fielders') consists of Anglo American, BHP, Freeport McMohRan and Teck Resources; and a third, underperforming, category consists of Vale S.A. and the four large gold producers, Newmont, Newcrest, Barrick and Kinross Gold. Compared with the 2009 study, a larger proportion of companies have total shareholder returns that have underperformed their respective total reserves increment. The causal factors behind his require more research but are hypothesised to include financial variables such as debt taken on for acquisitions or perceived liabilities. This is specifically the case when companies borrow at the height of the commodity price cycle, suffer cost blow outs in developing new projects and then commence production in the low point of the cycle. The timing of investments in resource management, project development, operations and market management relative to respective commodity price cycles plays an important role in total shareholder return.

INTRODUCTION

In 2009, Camus, Knights and Tapia presented a conference paper that challenged the notion that value creation in mining mainly concerns optimisation of production and costs. Their paper contended that the capacity of a company to effectively manage upstream processes – resource related activities whose purpose is to continually replenish prospective mineral resources – is as, if not more, important than the ability to effectively construct mines and produce at low cost with zero harm. To support this they calculated total shareholder return (TSR) and compared this with Total Reserves Increment (TRI) for a number of prominent international mining companies over the period 2000–2008. This period coincided with a complete economic cycle, being the first half of the so-called super cycle of resource investment driven by the resurgence of the Chinese economy. The global financial crisis affected commodity prices in 2008/2009, however recovery was swift and the super cycle returned from 2009 until late 2012. With the exception of two or three cases, total shareholder return in general followed total reserves increment.

The purpose of this paper is to revisit and update the 2009 study. For a selected group of 15 global mining companies share price, production and minerals reserve data was sourced for the period 2009–2024. This period includes the slow recovery following the 2012/2013 commodity price slump, the imperative to optimise production from capacity investments made during the super cycle leading to increased adoption of digitalisation and automation technologies, coal divestments and the shift towards clean power and electric vehicles, the invasion of Ukraine by Russia and the COVID pandemic and subsequent recovery. Total shareholder return and total reserve increments were calculated using the same approach as for the 2009 study, using an average copper price from 2024 to calculate copper-equivalent (CuEq) production and reserves.

PRIOR ART

The following paragraphs, drawn from Camus, Knights and Tapia (2009), contend that the mining business value chain can be conveniently grouped in the four activities as depicted in Figure 1. These are the primary activities defined in the value chain framework proposed by Porter (1985). This model also considers some support activities overarching the value chain, such as general management, technology management and human resources management.



FIG 1 – Mining business value chain.

Splitting the mining business this way has become widespread after Porter's (1985) seminal contribution. However, measuring the value-added contribution of each activity along the value chain to then set the company's course on the right activities and organise the business accordingly remains a challenge.

The upstream resource-related activities are described as mineral resource management. This function covers the exploration, planning and evaluation activities whose respective objectives are to discover mineral resources and transform them into mineral reserves. The outcome is a business plan that defines the methods and rate at which the orebody is mined and processed (life-of-mine production plan), along with the on-site and off-site infrastructure needed to deliver the mineral products to market. This is critical as Baker and Giacomo (1998) estimate that 75 per cent of mining projects fail due to not understanding geological orebody properties!

The downstream processes are the industrial-type activities that build and execute the business plan. These begin with the project management function accountable for building the development component of the plan, basically the engineering and construction works. After project commissioning, the operations management function takes responsibility for and executes the ongoing productive component of the plan. The outcome, in this case, is delivery of production, which typically includes mining and processing activities, including logistics and transportation. At the end of the value chain is the marketing and commercial management function whose aim is to realise revenue from the mineral products.

In mining, value creation initiatives more often focus on downstream activities, particularly project and operations management. Instead, the proposition here favours mineral resource management, as it is claimed to be the most value-accretive activity. As value contribution throughout the value chain cannot be directly assessed using public domain information, this proposition is supported by inference from analysis of company financial performance.

The hypothesis is that mining companies that excel in TSR over an entire economic cycle are those that also excel in increasing their reserves and production, hereafter referred to as TRI.

METHODOLOGY

TSR is an index that measures the overall company financial performance over time. It takes into account the share price variation and dividends paid to shareholders in a period. Hence, it provides the total return obtained by shareholders during that period. Dividends can be simply added to the share price difference although the most accepted assumption is their reinvestment in additional shares.

A key feature of this index is that while its absolute value is contingent to market conditions, companies' relative position within a peer group reflects the market perception of both individual and overall performance. This characteristic is useful to isolate the effect of non-controllable market variables (eg commodity and consumable prices) on the index and thus focus attention on the controllable variables that also affect TSR. It is calculated as:

$$TSR = \frac{P_i - P_{i-1} + \sum_i Div}{\mathrm{P}_{i-1}}$$

where:

P_i is the share price at the end of period i

P_{i-1} is the share price at the end of period i-1

Div are the dividends paid in period i

TRI is an index that measures variation in mineral reserves and production over time. TRI is calculated as:

$$TRI = \frac{R_i - R_{i-1} + \sum_i Prod}{\mathrm{R}_{i-1}}$$

where:

R_i attributable recoverable reserves at the end of period i

R_{i-1} attributable recoverable reserves at the end of period i-1

Prod attributable production during the period i

Mineral reserves are the economically and socially acceptable mineable part of mineral resources. As such, they represent future production that companies expect to extract from their mines. Its economic and technical viability is established by the planning and evaluation function. This involves certain market assumptions together with the capital and operating costs of the downstream activities spanning from project development, through operations, to marketing.

In the model under analysis – where TSR and TRI are the dependent and independent variables, respectively – the counterpart of the company share price is mineral reserves. For operating mining companies, dividends are frequently the matching part of production. This correspondence in model variables is in fact what gives economic sense to the term ‘deposit’, which is usually used in mining parlance to refer to an orebody.

In order to confirm this hypothesis, this study compares TSR and TRI for a selected group of 15 mining companies over the period 2009–2024. The mining companies included in this study broadly represent the whole mining industry as they cover a wide spectrum; from global, diversified mining groups, through mid-size, more focused players, to large and mid-tier international gold producers. Table 1 lists all the companies included in this research study along with their stock exchange trading information.

TABLE 1

Mining companies included in the study.

Company	Incorporated in	Listed on	Ticker symbol	Currency
Anglo American	South Africa/UK	LSE	AAL	GBP
Antofagasta PLC	Chile/UK	LSE	ANTO	GBP
Barrick Gold	Canada	NYSE	ABX	USD
BHP Ltd	Australia/UK	ASX	BHP	AUD
First Quantum Minerals	Canada	TSE	FM	CAD
Freeport McMoRan	USA	NYSE	FCX	USD
Glencore	Switzerland	LSE	GLEN	GBP
Gold Fields Ltd	South Africa	NYSE	GFI	USD
Kinross Gold Corp	Canada	NYSE	KGC	USD
Newmont Mining Corp	USA	NYSE	NEM	USD
Newcrest Mining Ltd	Australia	ASX	NCM	AUD
Rio Tinto Ltd	UK/Australia	ASX	RIO	AUD
South 32 Ltd	Australia	ASX	S32	AUD
Teck Resources	Canada	TSE	TECK.B	CAD
Vale S.A.	Brazil	NYSE	VALE	USD

DATA COLLECTION AND ANALYSIS

Copper-equivalent (CuEq) reserve and production series were developed for major mining companies from 2009 to 2024. Reserve and production quantities were taken from annual reports, and commodity price inputs were sourced from the price decks used in regulatory disclosures. All metals were first valued using the relevant reporting prices, which allowed the different components of each company's reserve base to be expressed in a common monetary framework. The combined value of all contained metals was then converted into an equivalent quantity of copper, producing a single CuEq measure for each company-year observation. This enabled direct comparison across firms with different commodity mixes by translating multi-metal reserves into a consistent unit.

A fixed price deck was applied across the entire time series to ensure that the CuEq estimates reflect underlying physical changes rather than shifts in commodity pricing. Using constant prices avoids the double counting that occurs when rising prices inflate CuEq reserves independently of geological or operational movements. It also prevents reserve reductions caused by mining depletion from being masked by contemporaneous increases in metal prices. Applying one price basis provides a stable foundation for tracking reserve additions, depletions and reclassifications across time and across companies.

This framework aims to generate a consistent and comparable data set that supports analysis of long-term reserve trends within the global mining industry. The resulting CuEq data set provides a basis for analysing long-run reserve dynamics, enabling direct comparison of reserve scale, concentration, and materiality across the industry over time.

Data collection difficulties

Compiling a consistent reserve and production data set across multiple companies and years presented several challenges. Reserve disclosures vary in structure, level of detail and reporting conventions, particularly with respect to metallurgical recoveries, attributable ownership, and the treatment of multi-metal orebodies. Companies also differ in how they classify and report reserve changes, which required harmonisation to ensure comparability across the data set.

Another difficulty arises from the interaction between reserve quantities and the price assumptions used in regulatory reporting. Reserves are disclosed in physical units, yet their valuation is inherently tied to the commodity prices applied by each company at the time of reporting. Changes in these price decks can create apparent movements in reserves even when the physical quantities are unchanged, complicating efforts to isolate genuine geological or operational changes. Ensuring that copper-equivalent values reflected real shifts in the reserve base rather than price-driven effects required careful standardisation of inputs and consistent treatment of price assumptions across companies and years.

Data processing complexities

Extending the copper-equivalent series back to 2009 required addressing several methodological constraints, the most significant being the treatment of commodity prices across the historical period. Two options were available: reconstruct the series using the price deck applied in each reporting year or apply a single, fixed price deck across the entire time span. A constant price basis was adopted to avoid embedding year-specific price volatility into the copper-equivalent calculations and to ensure that observed movements reflected physical reserve changes rather than fluctuations in metal markets. This mitigated the risk that rising commodity prices would give the appearance of reserve growth even when underlying geological or operational factors resulted in depletion or reclassification.

Additionally, updating the analysis in later years required substituting the previous price deck with the new reporting prices, which resulted in all copper-equivalent values from 2009 onward adjusting accordingly despite unchanged physical reserve quantities. This sensitivity is inherent to valuation-based conversions and reflects the dependence of copper-equivalent figures on the selected pricing basis. Corporate events also required standardised treatment. Mergers and acquisitions were incorporated by aligning reserve and production data with the reporting entity for each year so that only the reserves attributable to the company at that time were included. Share buybacks were excluded from the copper-equivalent calculations because they do not affect the physical reserve base, although they were relevant to the assessment of total shareholder return.

Choice of baseline

The baseline year in our analysis was selected as 2009. This was in part because it was natural progression from the previous study which included the 200–2008 period, but also the Global Financial Crisis of 2008/2009 caused a brief reset in share and commodity prices for many of companies. The period 2009–2024 includes at least two commodity price cycles (one of seven years from 2009–2016 and another shorter cycle of four years from 2016–2020, cut short due to COVID). The initial baseline for the analysis is therefore chosen to be the start of the 2009 cycle.

RESULTS

The aggregate results for Total Shareholder Return versus Total Reserves Increment for the 15 global mining companies included in the study are shown in Figure 2. A line with unit gradient extending from the origin has been included as a reference. We have listed dual listed companies according to one Exchange. As dual listings are interchangeable across countries, there is little meaningful difference, as any difference would be immediately exploited by arbitrage traders. Note that Rio Tinto (and First Quantum) are positioned well above the 1:1 reference line. Further investigation of this showed that the baseline year coincided with Rio Tinto's somewhat contentious purchase of Alcan in 2009. This acquisition was largely debt financed. As a result, Rio Tinto share values dropped from around $110 to $30 at the start of 2009. The share price was back up at around $75 by the end of the year, however the implication is that the annual share price forms a very low

baseline for the TSR measurement. This led us to investigate the effect of changing the baseline year.

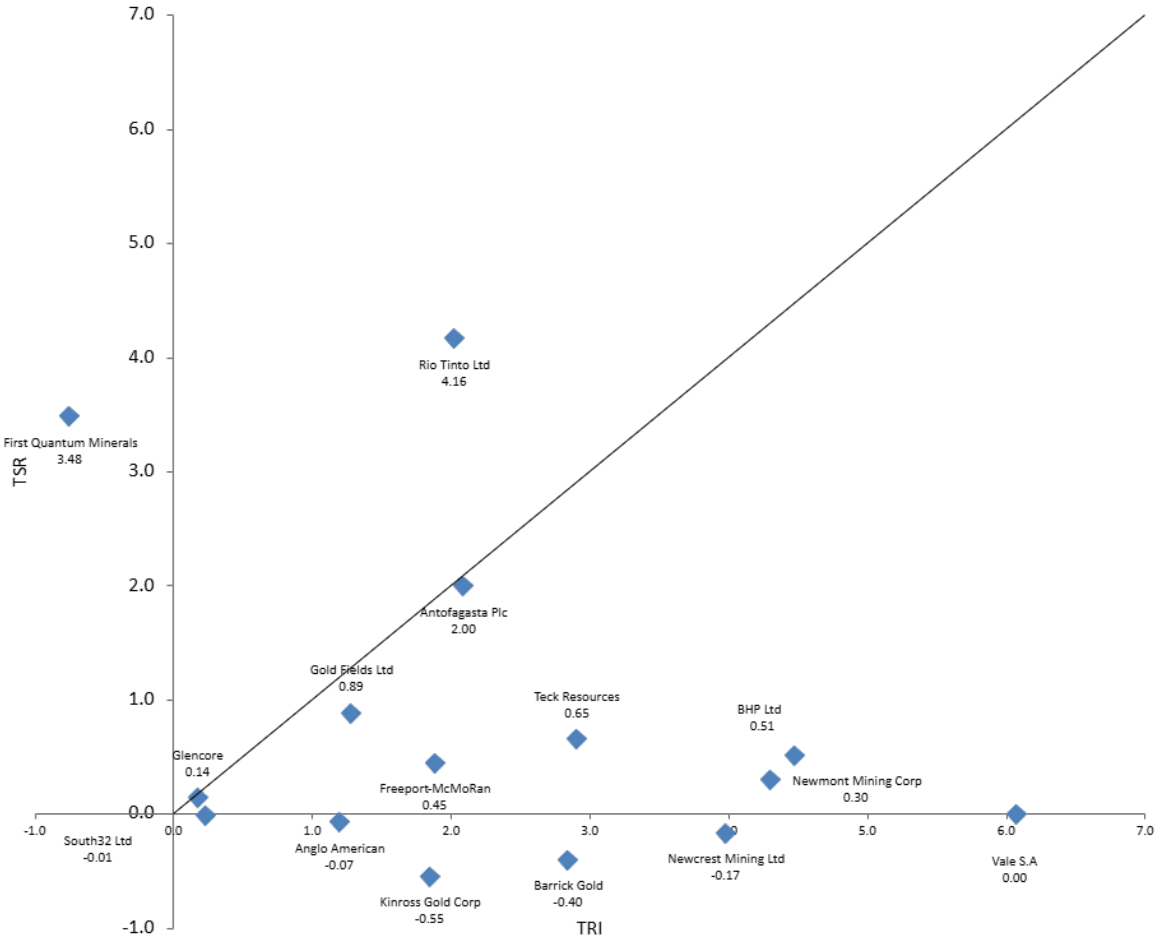


FIG 2 – Total shareholder return versus total reserves increment for the period 2009–2024.

Revised baseline

Figure 3 shows the effect of selecting a different start date by reducing the window of analysis to 2010–2024. Rio Tinto's performance is now highly correlated with its total reserve increment, similar to that of BHP and Anglo American. With the exception of First Quantum, The TSR performance of companies can be classified into three groups.

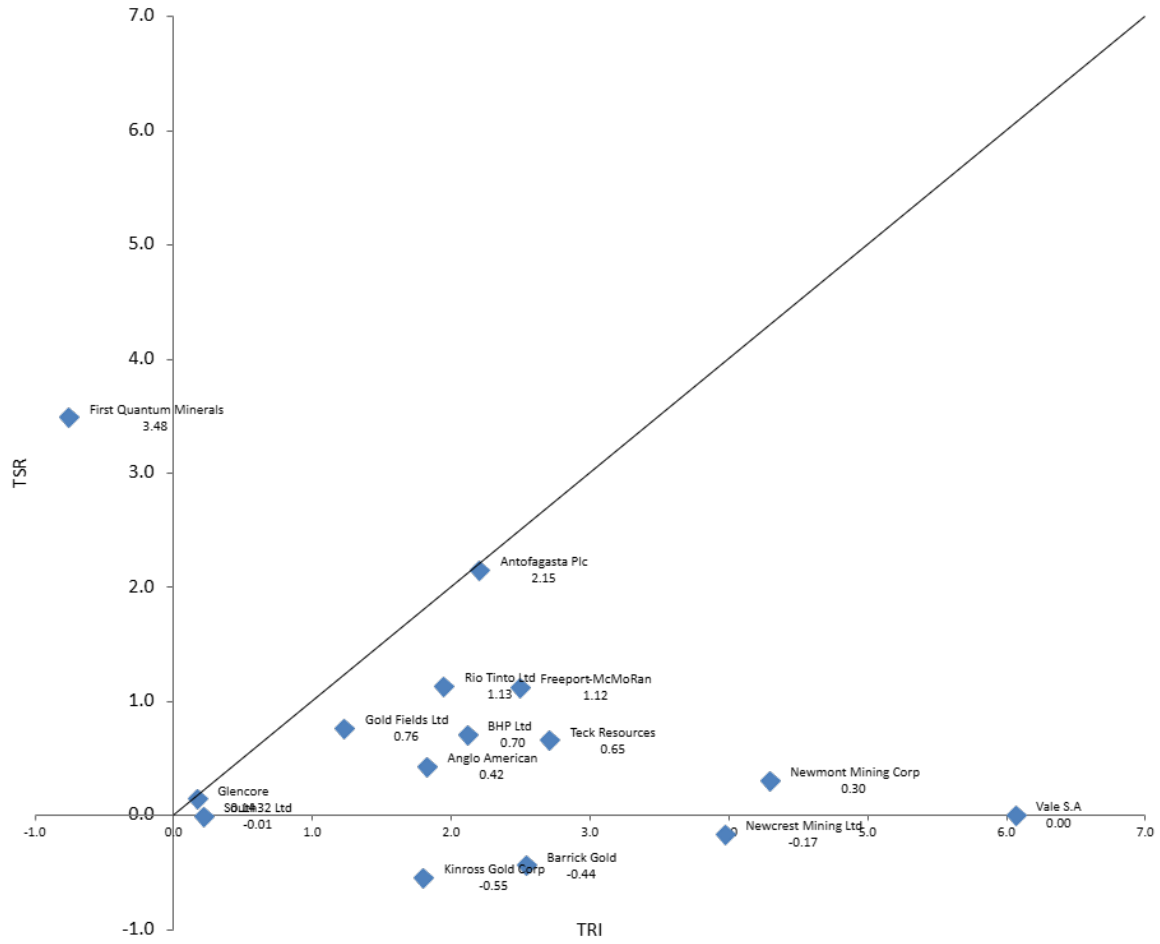


FIG 3 – Total shareholder return versus total reserves increment for the period 2010–2024.

The first category is where total shareholder return is commensurate with total reserve increment. This includes Glencore. South32, Gold Fields, Rio Tinto and Antofagasta Minerals. The second category (the so-called 'mid-fielders') consists of Anglo American, BHP, Freeport McMohRan and Teck Resources. The third category, the underperformers, consists of Vale S.A. and the four large gold producers, Newmont, Newcrest, Barrick and Kinross Gold.

In the first category, Antofagasta Minerals is a large copper miner, Gold Fields is a dominant gold producer and Glencore and South 32 are diversified mineral producers having significant legacy coal assets. Glencore is predominantly a trading company. Its 2024 revenues were USD230 billion. However not all of this is primary minerals production, as Glencore is one of the world largest recyclers of end-of-life electronics, and a major recycler of complex secondary copper, nickel, cobalt, gold, silver, platinum and palladium bearing feeds. As coal producers, their share prices may be weighed down due to heightened financial risk of incurring stranded assets as the world moves to phase out thermal coal power stations. Rio Tinto has benefited from its stronger position in iron ore and a more aggressive growth strategy in commodities such as copper and aluminium. Rio Tinto divested its coal interests in 2017 and has subsequently acquired significant lithium interests. It can be argued that the market premium that Rio Tinto enjoys is as a result of positioned itself as a future-facing company by embracing materials required for the energy and AI technology transitions. Google Gemini states:

> *"Rio Tinto has benefited from a focus on its iron ore and expanding copper businesses, which have performed strongly, especially with a robust expansion strategy in Mongolia and Guinea".*

In the second category, Google's Gemini AI explained the marginal underperformance of Anglo American and BHP as follows:

"Anglo American has underperformed Rio Tinto due to factors like operational challenges, including the lingering effects of the pandemic on South African and South American assets, a less advantageous commodity mix, and the recent disruption from the failed BHP takeover bid".

Meanwhile:

"BHP has underperformed Rio Tinto from 2009-2024 due to several factors, including its exposure to underperforming nickel and coal assets, while Rio Tinto has benefited from its stronger position in iron ore and a more aggressive growth strategy in commodities like copper and aluminium. BHP's financial performance has also been affected by a major dam collapse settlement, and its focus on cost reduction and strategic assets has resulted in a more cautious approach to shareholder returns compared to Rio Tinto's higher dividends over the period".

First Quantum Minerals appears to be an anomaly given its higher total shareholder returns and negative total reserves increment. Google Gemini states:

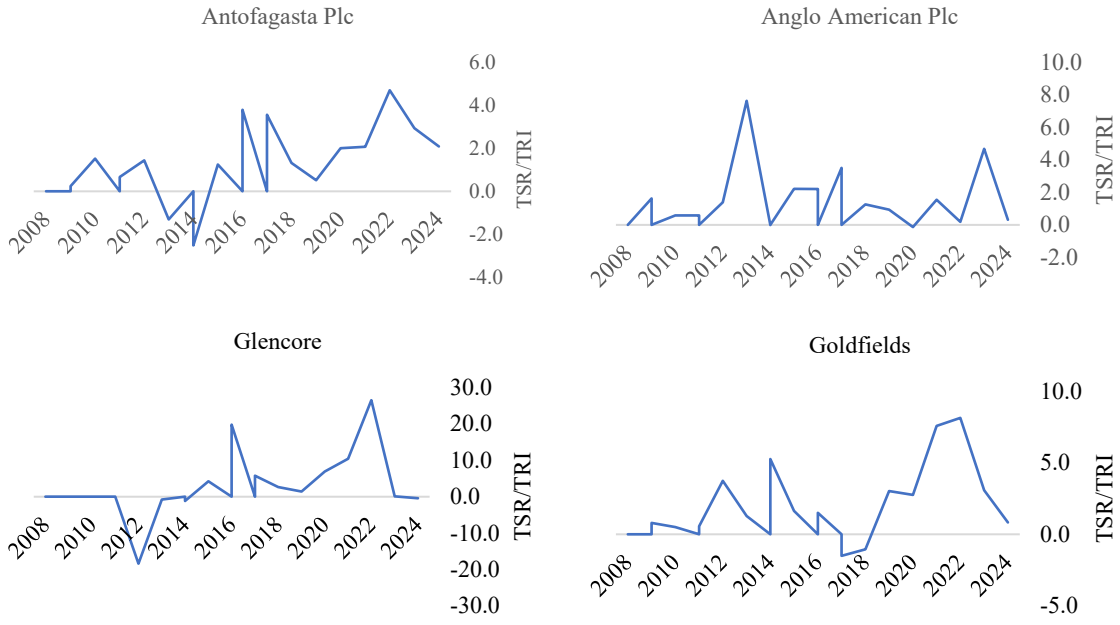
"The company's performance was a mix of significant expansion and development, such as the Kansanshi S3 expansion in Zambia, and periods of struggle, like the Cobre Panama shutdown and its impact on profitability and market value. While the company implemented strategic financial measures to stabilize its balance sheet and had strong performance from its Zambian operations during certain periods, the closure of its largest asset caused a significant stock drop."

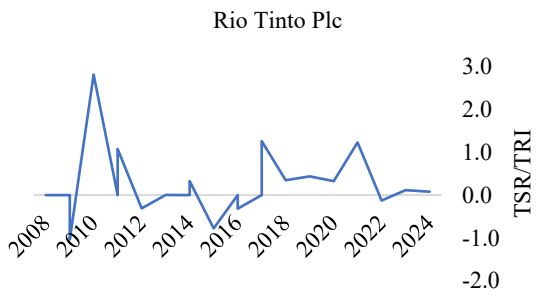
It is somewhat surprising to see the large gold miners, Barrick, Newmont and Newcrest in the underperforming group. Google's Gemini AI suggested the following:

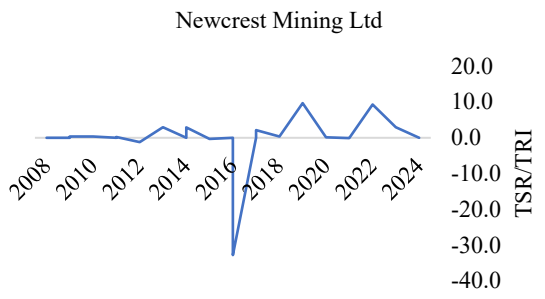
"Gold companies underperformed from 2009 to 2024 due to a combination of rising production costs, poor capital discipline, and operational challenges, which squeezed profit margins despite the rising price of gold. Companies borrowed heavily and pursued projects that became unprofitable when gold prices fell, leading to write-downs and investor scepticism. Additionally, gold miners are subject to operational risks and reserve depletion that physical gold does not face".

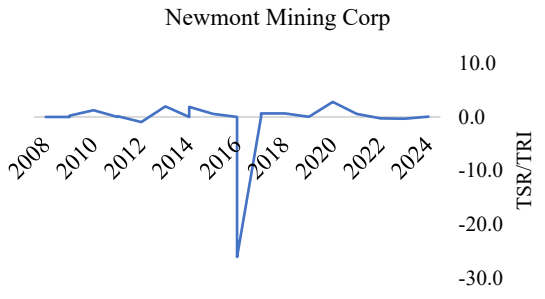
Longitudinal analysis

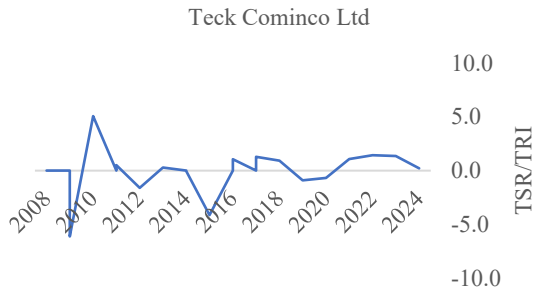
In order to understand these trends in more detail, longitudinal plots of the TSR/TRI rations were prepared for each of the mining companies in the study. These plots are reproduced in Figure 4.

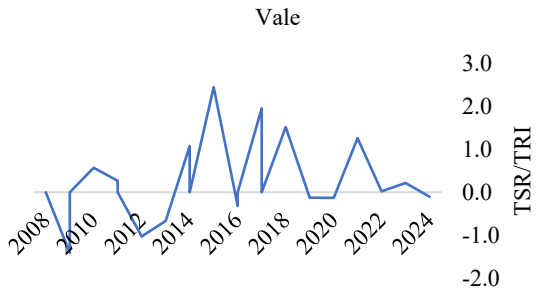


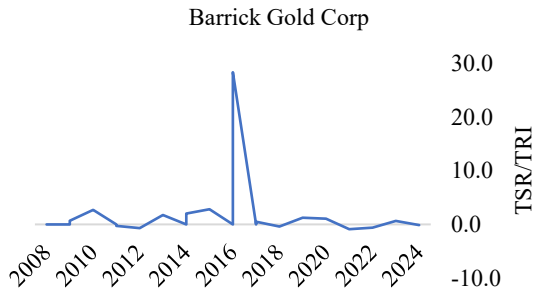

Rio Tinto Plc
3.0
2.0
1.0
0.0
-1.0
-2.0
TSR/TRI
2008
2010
2012
2014
2016
2018
2020
2022
2024

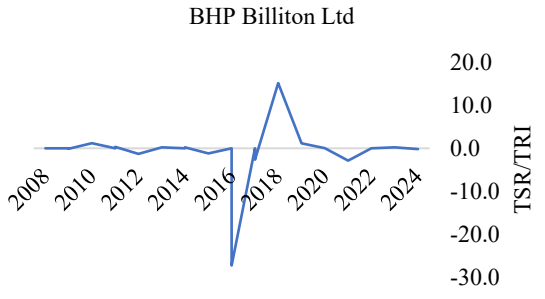

Newcrest Mining Ltd
20.0
10.0
0.0
-10.0
-20.0
-30.0
-40.0
TSR/TRI
2008
2010
2012
2014
2016
2018
2020
2022
2024

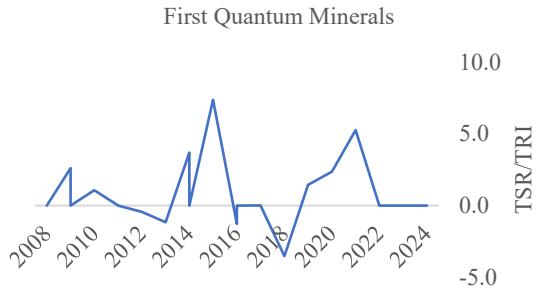

Newmont Mining Corp
10.0
0.0
-10.0
-20.0
-30.0
TSR/TRI
2008
2010
2012
2014
2016
2018
2020
2022
2024

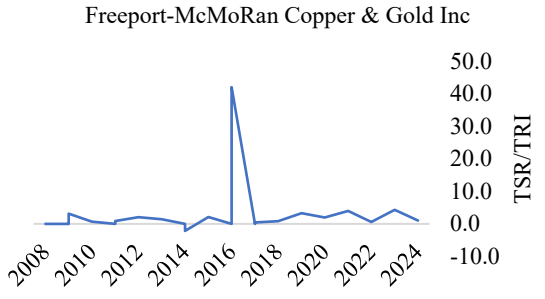

Teck Cominco Ltd
10.0
5.0
0.0
-5.0
-10.0
TSR/TRI
2008
2010
2012
2014
2016
2018
2020
2022
2024

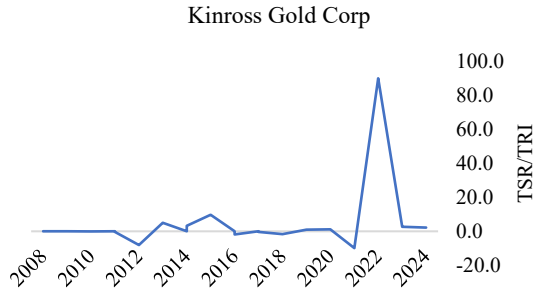

Vale
3.0
2.0
1.0
0.0
-1.0
-2.0
TSR/TRI
2008
2010
2012
2014
2016
2018
2020
2022
2024


Barrick Gold Corp
30.0
20.0
10.0
0.0
-10.0
TSR/TRI
2008
2010
2012
2014
2016
2018
2020
2022
2024


BHP Billiton Ltd
20.0
10.0
0.0
-10.0
-20.0
-30.0
TSR/TRI
2008
2010
2012
2014
2016
2018
2020
2022
2024


First Quantum Minerals
10.0
5.0
0.0
-5.0
TSR/TRI
2008
2010
2012
2014
2016
2018
2020
2022
2024


Freeport-McMoRan Copper & Gold Inc
50.0
40.0
30.0
20.0
10.0
0.0
-10.0
TSR/TRI
2008
2010
2012
2014
2016
2018
2020
2022
2024


Kinross Gold Corp
100.0
80.0
60.0
40.0
20.0
0.0
-20.0
TSR/TRI
2008
2010
2012
2014
2016
2018
2020
2022
2024

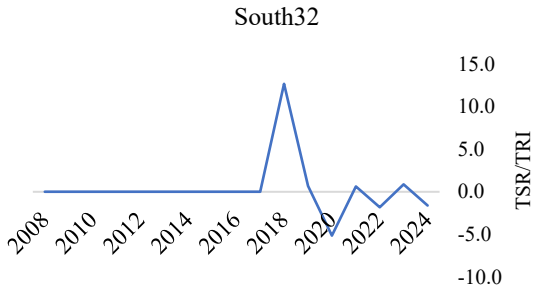


FIG 4 – Longitudinal plots of TSR/TRI ratios for the period 2009–2024.

It is immediately apparent that in 2016 a spike in the ratio occurs for 6 of the 15 companies. Co-pilot advanced the following response as to why this might be the case:

> *"Mining share prices in 2015 were primarily driven by the steep decline in commodity prices, especially iron ore and gold; weakening Chinese demand and stock-market turmoil; stronger USD and falling oil prices; financial strain from losses and high leverage; and companies' reactive cost management, capex cuts, and asset rationalization. The S&P GSCI dropped ~24% year-to-date, with mining indices like GDX and XME plunging ~22% and ~51% respectively. Iron ore prices fell ~44% year-over-year (Q1 to Q3), weighed down by oversupply from Australia/Brazil and Chinese demand weakness. Gold prices struggled, reaching six-year lows by December, driven by stronger US dollar, rate hike fears, and lackluster demand from China and India."*

Our conjecture is that this spike reflects a recovery from the 2015 response.

Rio Tinto has experienced a number of sudden increases in TSR/TRI ratio, in 2009 (being the rebound from the GFC), 2016, 2017 and 2020. It is postulated that the 2017 bounce is due to a spike in iron ore prices plus the divestment of its Australian coal assets. The 2020 bounce is most probably related to the COVID pandemic that severely curtailed iron ore production in Brazil causing a massive increase in iron ore prices.

DISCUSSION

Over the period of the study the NASDAQ Composite index has delivered a return of in excess of 840 per cent (see <https://www.macrotrends.net/1320/nasdaq-historical-chart>). This is indicative of the extraordinary rise of technology and clean energy stocks. A technology trend recent report from McKinsey and Company (Yee *et al,* 2025) identified 13 frontier technology trends, including: Agentic AI; artificial intelligence, application specific semiconductors, advanced connectivity, cloud and edge computing, immersive reality technologies, digital trust and cybersecurity, quantum technologies, robotics, mobility, bioengineering, space technologies and energy and sustainability technologies. In 2024 these technology industries collectively attracted an estimated USD772.1 billion in equity investment. In the same year, S&P Global (2024) estimate the investment in corporate mining exploration to be $12.8 billion, representing just 1.65 per cent of the investment equity in technology stocks!

For much of the recent decade, equity markets have been largely dominated by technology and sustainable energy stocks. Mining projects, often associated with long, uncertain lead times to develop new ore reserves, have had to fight for a share of equity finance in the current market. We contend that the category 1 cohort of mining companies identified in this study (eg Rio Tinto) have a future-facing portfolio of minerals that are strongly aligned with the growth of technology industries. Their portfolios are dominated by critical minerals that are in demand and have generated large cash flows in recent years (namely copper and iron ore).

Compared with the 2009 study, a larger proportion of companies now lie in the category where total shareholder returns have underperformed total reserves increment. Further research is required in order to identify and analyse the causal factors behind this trend. Some reasons may lie in financial

variables such as debt required for acquisitions or perceived liabilities associated with these companies. This is specifically the case when companies borrow at the height of the commodity price cycle, suffer cost blow outs in developing new projects and then commence production in the low point of the cycle. Returning to Figure 1, it is apparent that the timing of investments in resource management, project development, operations and market management relative to respective commodity price cycles plays an important role in total shareholder return.

Some of the underperforming group of companies (including Vale, Newmont, Newcrest, Barrick and Kinross Gold) have assets in high-risk countries such as South Africa, Mali, PNG and Pakistan. Barrick is facing large environmental legacies associated with the Pascua Lama mine located on the border of Argentina and Chile. Anglo American has coal assets in Australia and, through its ownership of de Beers, is exposed to potentially stranded assets due to the market disruption from synthetic diamonds. Anglo American has significant copper assets in South America, however these have not lived up to capacity in recent years. BHP is developing the Jansen potash project in Saskatchewan with the potential for this to become a new business unit orientated at the agricultural rather than technology sector. It is also subject to ongoing legal actions in the case of the Samarco and Brumadinho tailings dam failures.

CONCLUSION

Returning to the central hypothesis of this study, namely does shareholder value for mining companies follow mineral reserves growth, it is clear that the answer is ‘not necessarily’. This study compared the total shareholder returns for 15 global mining companies with their reported total reserves increments. With the exception of First Quantum, TSR performance of companies can be classified into three groups; the first category has total shareholder return commensurate with total reserve increment (this includes Glencore. South32, Gold Fields, Rio Tinto and Antofagasta Minerals); a second category (‘mid-fielders’) consists of Anglo American, BHP, Freeport McMohRan and Teck Resources; and a third, underperforming, category consists of Vale S.A. and the four large gold producers, Newmont, Newcrest, Barrick and Kinross Gold.

Compared with the 2009 study, a larger proportion of companies have total shareholder returns that have underperformed their respective total reserves increment. The causal factors behind his require more research but are hypothesised to include financial variables such as debt taken on for acquisitions or perceived liabilities. This is specifically the case when companies borrow at the height of the commodity price cycle, suffer cost blow outs in developing new projects and then commence production in the low point of the cycle. The timing of investments in resource management, project development, operations and market management relative to respective commodity price cycles plays an important role in total shareholder return.

ACKNOWLEDGEMENT

The authors would like to acknowledge the financial support of Rio Tinto for this study. In particular, we express our gratitude to Ian Rees, Chris Tabb, Nigel James and Ben Myall for many useful conversations during the study.

REFERENCES

Baker, C K and Giacomo, S M, 1998. Resources and reserves: Their uses and abuses by the equity markets, Ore Reserves and Finance: A joint Seminar between the AusIMM and ASX, Sydney.

Camus, J, Knights, P and Tapia, S, 2009. Value Generation in Mining – A New Model, in Proceedings of the Australian Mining Technology Conference, Brisbane, Australia.

Porter, M, 1985. *Competitive Advantage: Creating and Sustaining Superior Performance* (Simon and Schuster: New York).

S&P Global, 2024. World Exploration Trends 2024 [online]. Available from: <https://www.spglobal.com/market-intelligence/en/news-insights/research/world-exploration-trends-2024> [Accessed: Nov 2025].

Yee, Y, Chui, M, Roberts, R and Smit, S, 2025. McKinsey Technology Trends Outlook 2025, 5th edition, July 2025 [online]. Available from: <https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/the-top-trends-in-tech-2025> [Accessed: Nov 2025].

Closing the gap between strategic and tactical mine planning

E Ponce[1], A Hannett[2] and A Flett[3]

1. Strategic Mine Planning Superintendent, Agnico Eagle Mines, Cochrane ON P0L 1C0, Canada. Email: erick.ponce@agnicoeagle.com
2. Assistant General Manager, Technical Services Manager, Agnico Eagle Mines, Cochrane ON P0L 1C0, Canada. Email: arthur.hannett@agnicoeagle.com
3. Short Term Mine Planning Superintendent, Agnico Eagle Mines, Cochrane ON P0L 1C0, Canada. Email: anton.flett@agnicoeagle.com

ABSTRACT

Strategic mine plans are commonly developed to maximise project value; however, operational constraints often emerge as plans progress toward execution. This can create a disconnect between long-term planning assumptions and achievable operational outcomes. At Detour Lake Mine, an operational gap assessment was undertaken to evaluate phase interaction, vertical advance rates, phase minimum mining width, truck congestion, shovel area, haulage access and geotechnical constraints. The assessment was conducted jointly by strategic planning, mine operations and short-term planning teams and was used to test the practicality of existing phase designs and mining sequences. Results identified opportunities to simplify the mine layout, improve mining widths, reduce equipment congestion and enhance operational flexibility. The assessment framework was subsequently incorporated into mine design and scheduling decisions, reducing the number of mining phases from eight to six and improving access to ore. The work demonstrates how integrating operational considerations into strategic planning can improve mine plan robustness, executability and long-term value realisation.

INTRODUCTION

Detour Lake Mine (Figure 1) is a large-scale open pit gold operation located in north-eastern Ontario, Canada, within the Abitibi Greenstone Belt. The mine is a bulk-tonnage, low-grade system exploited using conventional truck-and-shovel mining, where value is primarily driven by high throughput, operational efficiency, and long mine life.



FIG 1 – Detour Lake Mine 2024, Agnico Eagle.

Mineralisation is associated with the Sunday Lake Deformation Zone, a major regional structure hosting disseminated gold within quartz-carbonate vein stockworks and sulfide minerals.

The process plant operates at over 28 Mt/a using a conventional flow sheet consisting of crushing, grinding, gravity recovery, and carbon-in-pulp (CIP) leaching.

PHASE AND SCHEDULE RISK ANALYSIS

While strategic mine plans are designed to maximise project value, operational constraints often emerge as plans transition from Life-of-Mine (LoM) schedules to medium-term plans, short-term plans, and annual budgets. These constraints can include permitting requirements, environmental commitments, equipment productivity, workforce availability, haulage limitations, and geotechnical considerations, potentially creating a gap between planned and achievable outcomes.

To quantify this gap, an operational assessment was conducted at Detour Lake Mine to evaluate the practicality of the existing phase designs and mining sequence under operating conditions. The assessment focused on phase interaction, vertical advance rates, ore exposure, haulage access, truck congestion, shovel operating areas, mining front availability, and geotechnical constraints.

The analysis provided a framework for evaluating mining sequence efficiency, phase sizing, and operational practicality, allowing operational constraints and improvement opportunities to be incorporated directly into pit design and long-term scheduling decisions. The assessment demonstrated that phase design and the interaction between successive mining phases are significant drivers of mine plan robustness, operational flexibility, and value realisation.

At the time of the assessment, the operation mined approximately 125 Mt/a using a fleet of four hydraulic shovels, three electric rope shovels, and 51 haul trucks. While this configuration provided flexibility in ore access, it also introduced challenges related to equipment congestion, restricted mining widths, phase interaction, and operational complexity.

Based on workshops conducted between Mine Operations, Mine Technical Services, and Strategic Mine Planning, a set of operational risk criteria was developed to assess phase design and mining sequence performance, as shown in Table 1.

TABLE 1

Risk analysis.

Risk	Phase width	Vertical advance	Truck congestion	Shovel area	Phase interaction
High	80–110 m	> 6	Over 95 Mt	< 10 ha	No geotechnical berm between mining phases
Medium	110–140 m	5–6	75–95 Mt	10–12 ha	Geotechnical berm between mining phases
Low	>140 m	< 5	< 75 Mt	> 12 ha	Geotechnical berm and increase ramp width between mining phases

These parameters were used to redesign the mining phases and overall pit configuration to improve operational efficiency and reduce operational risk through the mine planning process.

STRATEGIC PHASE ANALYSIS

These criteria were subsequently used to redesign the mining phases and overall pit configuration, improving operational efficiency and reducing operational risk within the mine planning process.

The methodology included the assessment of equipment productivity, haulage and loading estimates, application of geotechnical criteria for slope design, reserve modelling, and the development of strategic mine schedules. Optimisation analyses were conducted using Direct Block Scheduling (DBS) to define and confirm the mine sequence direction. Additionally, multiple phases design scenarios (Sets A–Q) were developed. Their impacts were evaluated based on the risk

analysis, productivity achievement and indicative net present value (NPV) over 5-year, 10-year, and life-of-mine horizons.

The assessment focused on Phase 5, which contains the largest proportion of the mine's gold inventory over the next ten years of operation. As mining of this phase commenced in 2024, expanding its footprint was no longer considered a practical option. Instead, the evaluation focused on optimising the phase configuration to improve operational safety, increase equipment productivity, and enhance overall mining efficiency.

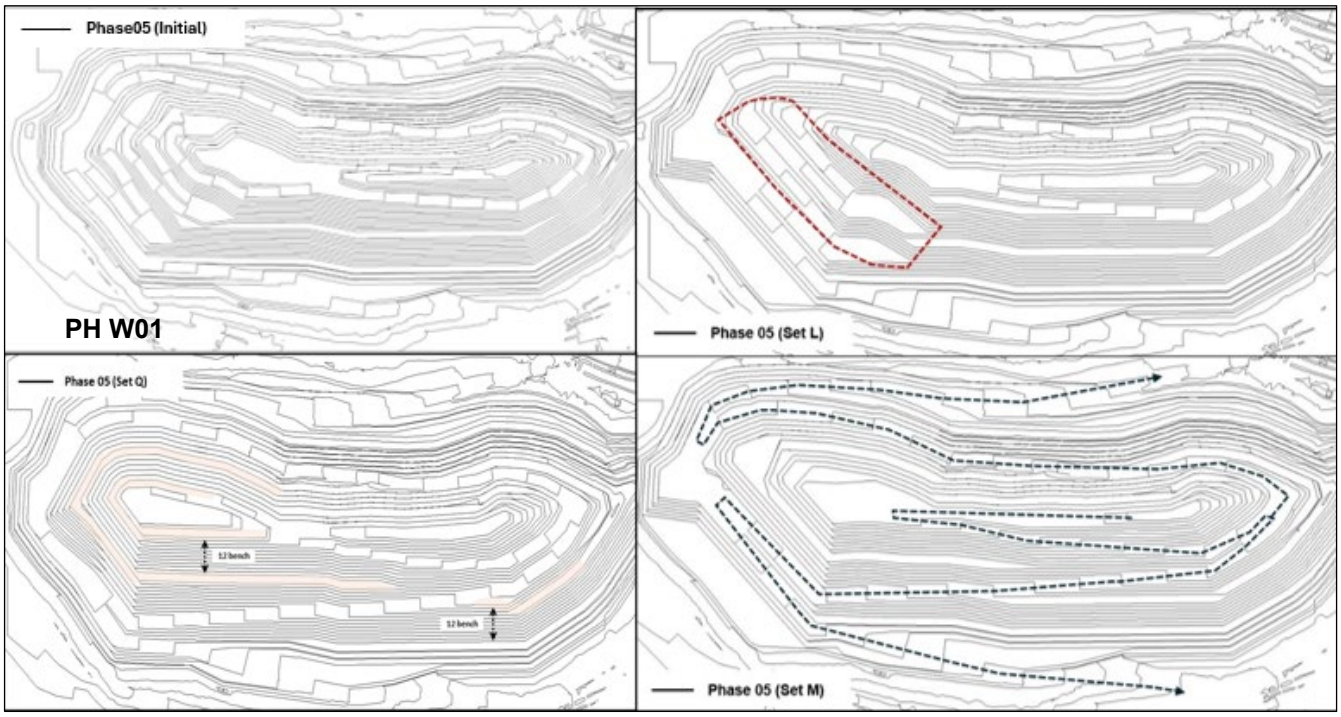


FIG 2 – Options phase 05 designs.

The original design of Phase 5 included a switchback haulage system along the west wall. Two main alternatives were evaluated:

1. Reducing the mining footprint in the south-west sector of Phase 5 while maintaining the switchback configuration along the west wall (Set L).
2. Introducing a southern access ramp while simultaneously reducing the mining extent in the south-west sector (set M).

The operational assessment showed that maintaining the switchback configuration preserved spillage risks and increased operational exposure due to reliance on a single access ramp. In contrast, the addition of a southern ramp reduced access-related risks and improved operational flexibility. Furthermore, this alternative enabled faster phase development, allowing earlier access to ore and improved production scheduling.

This evaluation led to a broader discussion regarding the minimum phase width required for 72-tonne electric rope shovels to operate efficiently while loading on both sides. Considering the space requirements for all supporting equipment, a minimum mining width of approximately 140 m was estimated to achieve optimal productivity.

Based on these considerations, the set Q design was developed, eliminating the switchback configuration and incorporating a southern ramp. This design improves operational efficiency, reduces risk, and aligns with equipment productivity requirements.

A similar analysis was conducted for the subsequent phases.

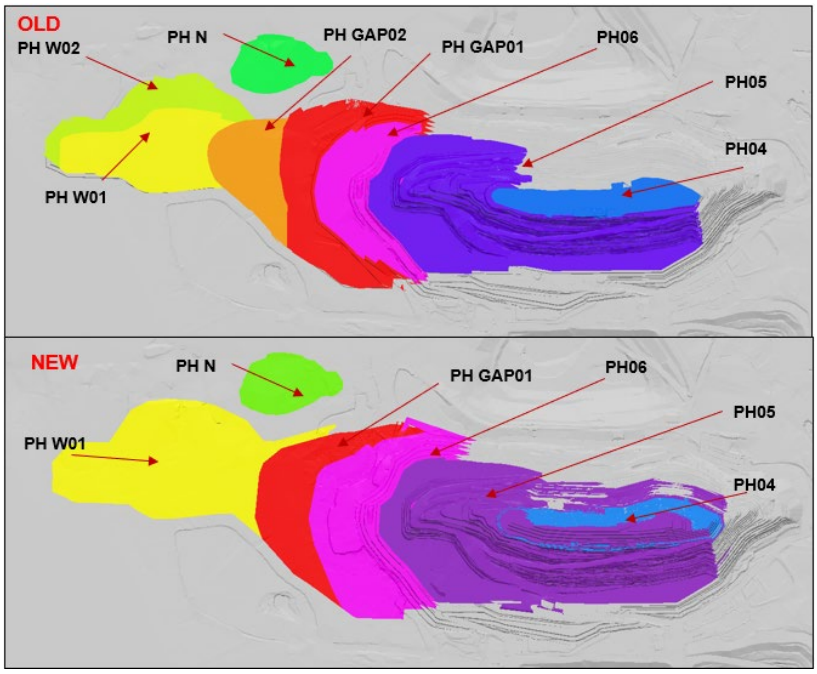


FIG 3 – Phases redesign.

CONCLUSIONS

The operational gap assessment conducted at Detour Lake Mine demonstrated that incorporating operational considerations into strategic mine planning can significantly improve the practicality and robustness of long-term mine plans. Application of this methodology reduced the number of mining phases from eight to six, resulting in wider mining fronts, fewer switchbacks, improved shovel productivity, reduced haulage constraints, and increased operational flexibility.

Mine planning process robustness

The planning framework developed for Detour Lake successfully integrated geotechnical, operational, and economic considerations into the strategic planning process. This approach improved alignment between long-term planning objectives and operational execution requirements while enhancing overall mine plan robustness.

Phase and sequence risk analysis

The risk assessment framework provided a practical method for evaluating phase design and mining sequence performance. Parameters such as phase width, vertical advance rate, truck congestion, shovel operating area, and phase interaction were used to identify operational risks and optimise pit configuration. The resulting phase redesign improved both operational efficiency and mining practical.

Operational flexibility

The study demonstrated that shovel productivity, equipment allocation, and haulage access can have a significant impact on project value, particularly during the first five to ten years of operation. Improvements such as increased ramp width and enhanced access arrangements in areas such as Set Q contributed to greater operational flexibility and reduced execution risk.

Balanced planning approach

The results highlight the importance of balancing value maximisation with operational and geotechnical considerations. Although more aggressive mining strategies may improve short-term value through accelerated ore exposure, they can also increase operational complexity and risk. Conversely, overly conservative approaches may reduce risk but limit economic performance. The most effective mine plans are those that achieve an appropriate balance between safety, operability, flexibility, and value generation.

Beyond the pit shell – a framework for strategic mine planning to challenge constraints and support value-driven decisions

B Price[1]

1. Senior Engineer, Strategic Production Planning, Rio Tinto Copper, Salt Lake City Utah 84009, USA. Email: branwenprice@riotinto.com

ABSTRACT

Strategic mine planning increasingly relies on advanced optimisation tools, yet the strategic value delivered by these software packages is often shaped more by their underlying constraints and assumptions than by the optimisation engines themselves. Planners explore a wide range of scenarios surrounding mining and processing assumptions, but often assume constraints linked to tailings, or infrastructure are non-negotiable. As a result, potentially high-value constraints may remain unchallenged. It can be difficult to identify which constraints and assumptions materially influence strategic value of proposed development pathways without having to evaluate an overwhelming number of scenarios. This paper shifts strategic mine planning from the optimisation of material movement schedules to a process that informs strategic operational decisions. It introduces a structured workflow that makes constraint assumptions transparent and testable.

The proposed framework integrates three steps into the front end of the strategic planning cycle. First, constraint triage classifies planning inputs as hard (non-negotiable), soft (negotiable with cost or time), or assumed (historical practice or internal conventions). This classification helps to clarify which constraints legitimately apply to a given site or asset, and which may represent opportunities to unlock value or reduce risk. Second, a focused set of bookend scenarios is developed to bound the feasible planning space. This matrix includes a conservative case in which key constraints remain fixed, and a targeted unlocked case in which a selected subset of high impact assumptions is relaxed within realistic technical and operational limits. Rather than optimising dozens of scenarios, this approach focuses on the analysis of the key value drivers and quantifying value and risk associated with each constraint set. Third, standardised communication of outcomes that is focused on informing strategic decisions. This communication emphasises consistent metrics between schedules, such as net present value (NPV) uplift relative to capital spend, sensitivity to price, and risk associated with execution delays or operational performance. A structured approach improves communication between technical planners, study teams, operations, and senior leaders, ensuring that optimisation insights quantify value and risk in a meaningful way.

The contribution of this work is not a new optimisation algorithm, but a versatile framework that strengthens the link between integrated scheduling and the quality of strategic decisions. By quantifying the value and risk associated with adjusting specific constraints, planners can move away from selecting a single ‘optimal’ schedule and instead support more timely, transparent, and defensible investment decisions under uncertainty.

INTRODUCTION

Strategic mine planning integrates fundamental aspects of the entire mining value chain from geology through operation to market. Advanced optimisation tools provide solutions to integrated production system including mining sequence, cut-off grade, processing options and capital spend (King, 2014). These optimisers rely entirely on input assumptions, that include system capacity constraints and financial parameters, provided by the asset that planners can choose to adhere to or not. Often operational assumptions such as equipment availability and productivities are flexed within a mine plan, but constraints linked to infrastructure, water allocation, and tailings capacity remain fixed. Parameters that were used in previous planning cycles are maintained without questioning if value could be unlocked by opening one, or a combination, of these constraints. Planners evaluate schedules that compare mining and processing options without addressing the potentially larger value levers such as water or tailings capacity, or a deferral of expansion throughput in favour of metal recovery. The result is a planning space that may evaluate dozens of

scenarios while being too narrow to unlock potential value. As a result, the optimisation insights might never reach the capital decision they were built to inform.

Strategic mine planning is most useful when treated as a decision-support activity rather than a routine modelling exercise. Literature in strategic planning has focused on two streams: integrated optimisation and stochastic modelling. While both have improved software, tools, and modelling approaches; neither operationalises the choice of the constraint set itself. That gap is the focus of this paper.

The following framework develops constraint framing from a recognised problem to a deployable practice in mine planning. Most importantly, this practice can be applied to various time horizons at a site-based level, not only at a strategic level. The framework makes constraint assumptions explicit before optimisation, focuses scenario work on bookends that bound value and risk, and translates schedule outputs into a common language shared by study teams, operations and executives centred around informing strategic decisions.

BACKGROUND

Decision analysis is rooted in properly framing the problem. Defining the boundaries and the purpose is key to ensure the analysis answers the necessary questions. Decision analysis identifies information and alternatives as the most impactful decision levers (Howard and Abbas, 2015). Strategic mine planning literature by contrast often focuses on one of two streams: integrated optimisation or uncertainty. Decision levers, integrated optimisation and uncertainty, all improve the output of an optimiser bound by a constraint set.

Mine planning has historically pushed hard on the information lever: orebody data, equipment availability trends, metal recovery curves, and schedule logic. However, treating mining sequence, cut-off grade and processing destination as sequential decisions can leave significant value on the table (King, 2014). Integrated solvers close that gap by optimising the decisions together. Yet alternatives as a decision lever are not often prioritised, or even considered, in planning cycles. Creativity is pushed aside in favour of providing the most optimal technical or financial plan in the shortest amount of time.

Stochastic modelling has focused on addressing uncertainty and risk associated with strategic mine planning. This method can integrate uncertainty in variables, such as geological information and metal prices, directly into the schedule (Dimitrakopoulos, 2018). Recent work extends the thinking to circular dependencies between capacity, cut-off grade, and sequencing (Balci, Cetin and Erkayaoglu, 2025). Likewise, practitioner case studies at complex mines and multi-product operations show that value is unlocked only when the constraint set is deliberately revisited rather than inherited (Wharton, 2009; Coutts, 2009). This framework aims to improve value-chain optimisation with circular dependencies by treating cut-off grade as a policy driven value lever rather than a fixed input.

Integrated optimisation and modelling uncertainty both aim to improve planning by determining higher value, or lower risk scenarios. However, neither addresses the quality of the decision process around the optimisation. Which constraints are accepted as binding? Why? Is the optimiser being asked the right question? Mine planners know that the quality of the inputs shapes the quality of the output, but often accept current constraints as non-negotiable based on previous plans and internal site conventions. Mine planners need a tool to enable them to challenge some of these constraints without being overwhelmed by an exhaustive scenario matrix.

A FRAMEWORK FOR CONSTRAINT-AWARE STRATEGIC PLANNING

The framework adds three steps to the strategic planning cycle: constraint triage, bookend scenarios, and communication centred around decisions. Each step produces a standardised data set or input specification that technical and leadership audiences review before continuing the planning process. The constrain triage informs the key levers evaluated by the bookend scenarios. After reflecting on initial results, a fourth step can be added to revisit these constraints and then bookend scenarios to refocus on the most promising value drivers based on the outcomes (Figure 1). This helps to

maintain the purpose of strategic planning on quantifying value and risk, and improve engagement throughout the process.

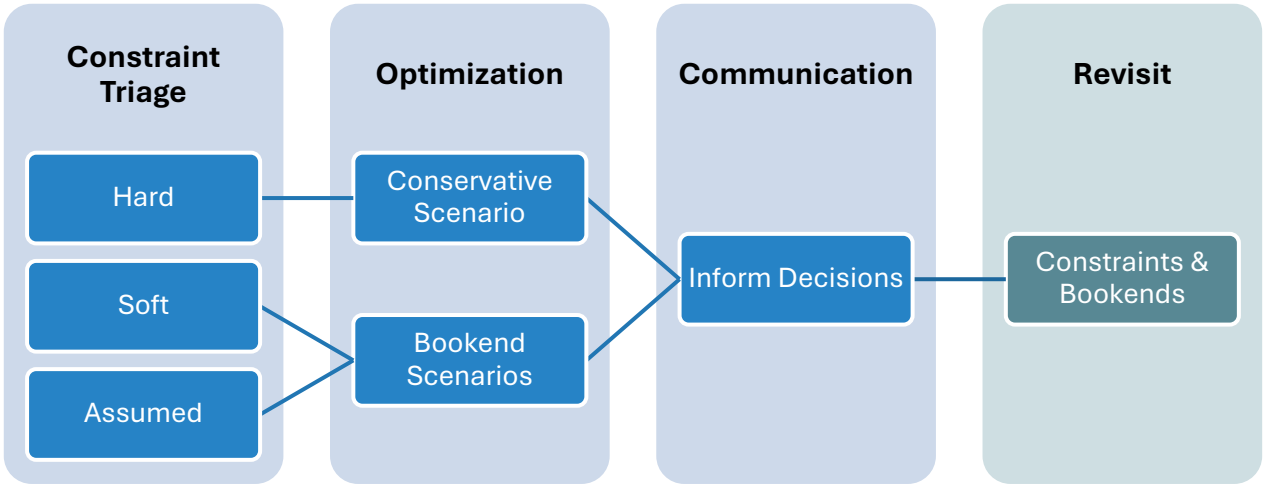


FIG 1 – Framework for constraint-aware strategic planning.

Constraint triage

Constraint triage sorts every planning input into one of three categories: hard, soft, and assumed constraints. The focus is determining which category a constraint falls into given the planning time horizon. Most planning teams find a small number of constraints carry a disproportionate share of unrealised value. Identifying these constraints helps determine which assumptions warrant an unlocked scenario and can provide insightful optimisations.

Hard constraints are non-negotiable on the relevant planning horizon. Such as geotechnical slope criteria, regulatory limits on concentrate production or water discharge, statutory closure obligations, and fundamental physical limits on comminution or recovery. A hard constraint has a clear basis of estimate, an identifiable owner, and no realistic pathway to relax the constraint inside the planning time horizon. It is important to remember that a constraint may be non-negotiable within a mid-range mine plan but is an assumed constraint for a strategic mine plan. For example, recovery and comminution parameters are a hard constraint for a budget planning cycle, but a strategic evaluation should determine if this constraint is in fact non-negotiable, or if it is a value lever that can enable an net present value (NPV) uplift with capital project.

Soft constraints are negotiable with cost, time or capital. Examples include processing throughput, fleet size and equipment availability, or in-pit conveyance systems. A soft constraint can be relaxed with an explicit capital expenditure. Whereas, assumed constraints are carried forward from prior studies, internal conventions or historical practices. They are often difficult to defend on a first principles basis and easy to overlook because they are embedded within the site as a non-negotiable practice. These range from legacy cut-off grades and fixed strip ratios, to tailings or port facility capacities, to recovery curves built on outdated metallurgical test work.

For example, if a site has maintained a 0.20 per cent CuEq cut-off grade consistently in their annual plans for the last ten years, it may be an assumed constraint. Operational expenses and commodity prices significantly impact cut-off grade. Therefore, under the current economic conditions 0.18 per cent CuEq, or even 0.15 per cent, may be the economic cut-off grade. Additionally, this lower grade material may not have accurate recovery data since it has been classified as waste for so long. Cut-off grade and recovery will be classified as an assumed constraint for material that falls in the range of 0.20–0.15 per cent CuEq, and a couple of scenarios can test the value of reducing the cut-off grade. This potential value can help bolster the argument of capital expenditure on metallurgical testing to determine accurate recovery data associated with the lower grade ore.

Compiling these categorised constraints into a register can be an effective tool to help streamline the process in the future. Soft and assumed constraints are ranked within the register by estimated influence on value. The ranking uses a simplistic approach to determine potential value, such as a basic grade-tonnage curve or benchmarking from a similar asset. The point is to identify a small

number of constraints that have the most potential for the highest value uplift and distinguishing between what the asset *must* accept from what it has *chosen* to accept.

Challenging these constraints often creates tension. The goal is to help structure the dialogue between different departments, such as maintenance, processing, environment or technical, to determine which key-value lever can unlock the most opportunity. Each of these stakeholders brings a different perspective to the mine plan, and discussing the constraint registry and priorities for analysis together can help drive a more engaged planning process. A constraint workshop to triage the various constraints can be incredibly valuable. This space encourages planners to work with operational owners to determine a credible list of soft and assumed constraints. Challenging assumed constraints requires buy-in from operational owners not just mine planners. This results in a constraint register that the asset will honour and value-levers that can unlock potential. A potential failure mode of this framework is an over-classification of constraints as hard. Defaulting to categorising every infrastructure or environmental input as a hard constraint could result in missed opportunities to unlock value. However, a constraints workshop can alleviate this by providing a structured space for engaging with departments and people who have different perspectives and priorities. Allowing the planning team to focus on key-value constraints that matter to the stakeholders across the asset.

Bookend scenarios

Bookend scenarios replace the extensive scenario list that is used in strategic planning with a focused matrix that targets challenging and evaluating the key constraints found during the constraint triage. This provides the bounds of the feasible planning parameter space. Moreover, it provides guidance for later detailed analysis to expedite the initial stages of the planning process.

The conservative scenario maintains the soft and assumed constraints being evaluated at their current assumptions. It is used as a starting point to compare against unlocked scenarios. The conservative scenario must be defensible on its own merits. It is the schedule the asset delivers if no constraints are challenged and the plan remains unchanged.

The unlocked scenarios relax a subset of key-value levers that are expected to provide the most opportunity based on the constraint triage. Each assumption is changed one-by-one, and the schedule is re-optimised in steps to ensure value can be linked directly to one constraint. The new assumptions are based on engineering judgement, industry benchmarking, or a new technology. For example, tailings capacity has been identified as a soft constraint but can unlock significant value with an expansion. The bookend scenarios will not simply double the tailings configuration regardless of the significant value that would unlock. Instead, the capacity assumption is relaxed to the next credible engineered alternative, such as a new lift with a given capex that is similar to the previous design. Additionally, there may be an opportunity to have a scenario that evaluates different stacking such as a process improvement like coarse particle flotation, or re-mining tailings to open capacity in the current facility. Another example is mill throughput as an assumed constraint. Though the size of the grinding circuit may be fixed, a bookend scenario could help evaluate a new technology to help debottleneck the system. Mill throughput should be increased to a reasonable rate with a credible cost estimate based on that technology. Likewise, if the mill has seen significant unplanned downtime, and this was flagged, a scenario should evaluate the risk to value, if that trend of lower throughput continues. This analysis can help support a refurbishment or capital expenditure on new mill equipment. Limiting the number of unlocked scenarios allows planners to focus on the key-value drivers and analyse the risk associated with each new constraint assumption.

Once initial scenario outcomes have been analysed, intermediate scenarios can be added where the value uplift between the conservative and unlocked scenario is large enough to warrant further resolution. These intermediate scenarios revisit the constraints and should focus on which part of the constraint is driving the bulk of the uplift. For example, a concentrator capital project increases throughput and recovery because of three smaller projects. Each of these projects can be decoupled to determine which aspect is the largest value lever. In hard ore, it may be the comminution improvement that delivers the most value; whereas the significant capital cost associated with a recovery improvement to the flotation circuit may not be driving the majority of the NPV uplift.

Therefore, a discussion around capital discipline and project delivery can be informed by these intermediate scenarios to ensure the project with the highest potential value is prioritised.

The goal of the bookend scenarios is to answer two questions: how much value is currently constrained by unchallenged assumptions? And which specific constraints contribute the largest share of that value? Focusing optimisation on the key-value levers can quantify value and risk in a timely and transparent way to help inform strategic decisions.

Communication to inform decisions

Clear communication of optimisation outcomes, value-levers, and risk is a critical aspect of mine planners' role. This step can be challenging, or poorly received by stakeholders, if there are intense pressures, time constraints, or if the outcomes highlight a significant risk to production guidance. However, being able to communicate results in a clear and consistent manner targeted to respective stakeholders ensures the important findings are received correctly. Reporting every scenario in the same format and showing the NPV uplift, or decrease, relative to the conservative scenario can help keep the conversation structured around the trade-offs rather than individual scenario results. To communicate risk, metrics can be reported as ranges rather than single values wherever the underlying constraint carries residual uncertainty. For example, an unlocked scenario that evaluates recovery uplift for a concentrator project in three options can be communicated as a range to centre the discussion around value opportunity and risk.

The waterfall graph in Figure 2 evaluates the bookend cases against the conservative case to engage leaders with the key-value levers rather than a large table of metrics. This structures the outputs on which metrics are used to drive decision-making, and how the trade-off between maintaining an assumed constraint or relaxing it can unlock opportunity.

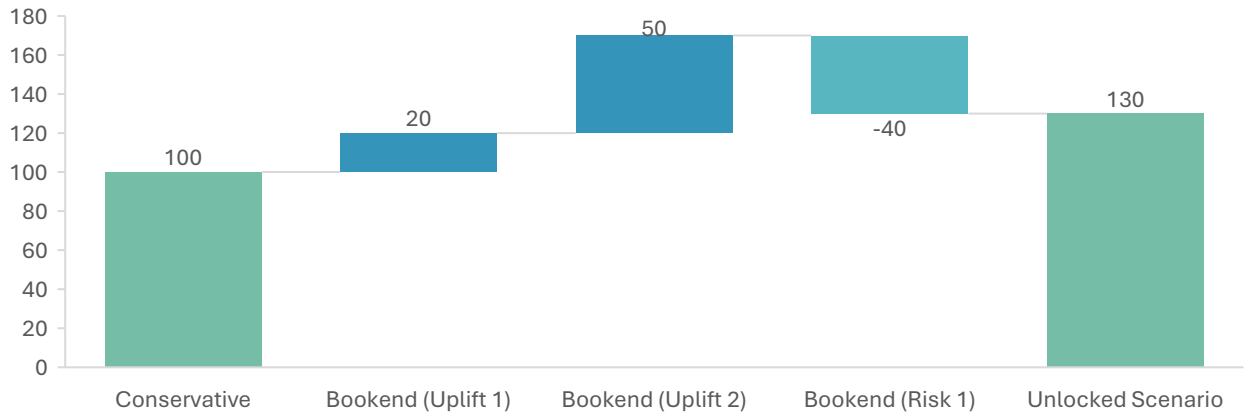


FIG 2 – Illustrative waterfall graph.

The communication of results is most powerful as a standing format. When every study or scenario presents NPV uplift, capital intensity, and execution risk in the same structure, leaders can compare opportunities across operations on a like-for-like basis. Planners can build a portfolio view to showcase which soft and assumed constraints have the greatest potential to unlock value for the business. This can help guide studies, capital discipline, and operational improvements. In instances where scenarios are highlighting risk to production, a common format can help shorten the review cycle, because the conversation can focus on what step changes need to be made to get back on track.

DISCUSSION

Scenarios must be defensible, not aspirational. Unlocked scenarios must stay inside realistic technical and operational limits, supported by a credible engineering pathway and an honest cost or time to relax the constraint. Without that discipline, the framework risks producing optimistic schedules that lack the basis for capital approval and erode planner credibility. A common example in mine planning would be a discrepancy between maintenance planning and production planning

inputs. This often causes tension between the two departments and can result in either unrealistic targets or plans that are too conservative to achieve production guidance. Either of which erodes trust in the mine plan and reduces collaboration. If the planned assumption for haul truck fleet mechanical availability is 88 per cent, but due to operational challenges the average mechanical availability has been 67 per cent. This can dramatically change the number of haul truck hours planned within a schedule. However, if the maintenance department is reporting in the annual planning cycle that mechanical availability targets will be met; there is often no avenue to identify the risk to the mine plan. By utilising bookend scenarios, mine planners can determine a reasonable value tied to achieving the mechanical availability target. The discussion can shift from unrealistic targets, or blame, to a risk centred approach. Value and risk can also help drive an operational improvement project, if there is significant value to unlock by outperforming the mechanical availability target. Bookend scenarios can guide the operation by emphasising which metrics matter most. Furthermore, if the bookend scenarios show that the bottleneck is loading capacity, not haul truck hours, then the conversation surrounding inputs can focus on unlocking shovel hours instead. The key takeaway is that by tying value directly to a specific assumption parameter, the operation can drive the true value levers of the business instead of being caught up in assumptions that are not constraining the system.

The idea that a system's value is governed by a small number of constraints is not new. Goldratt made this argument in *The Goal* (Goldratt and Cox, 1984), where he introduced the Theory of Constraints and reasoned that the strength of a chain is set by its weakest link. He also stated that an hour lost at a bottleneck is an hour lost for the whole system. This supports the central argument of this paper. It is the constraints and assumptions given to an optimiser, rather than the optimiser itself, that shapes the strategic value of a plan. Goldratt also stressed that an organisation must first define its goal and then choose measures that reflect it. His observation that people behave according to how they are measured maps onto the emphasis this framework places on consistent, decision-focused metrics such as NPV uplift relative to capital spend. Just as local efficiency measures can move a plant away from its goal, a mine plan optimised against inherited or assumed constraints can produce an efficient schedule that may never unlock the largest value levers.

Goldratt's five focusing steps also provide a useful contrast to the workflow proposed here. His steps are to identify the constraint, decide how to exploit it, subordinate everything else to that constraint, add capacity to the constraint, and then repeat the process with the new constraint that appears. Constraint triage performs the identify step by separating what genuinely binds the system from what has been accepted as binding. The bookend scenarios exploit and elevate the selected constraints by quantifying the value and the capital required to relax them within credible limits. His fifth step is the most relevant to this work. Once a constraint is broken, the analysis must return to the first step, because the binding constraint will have moved elsewhere in the system. This is the same behaviour seen in the haul-truck example above, where relaxing truck availability can expose loading capacity as the true bottleneck and shift the focus to unlocking shovel hours. Goldratt paired this fifth step with a warning about inertia, which is the tendency to carry forward an assumption after the conditions that justified it have changed. That warning describes the failure mode this framework is designed to address. Assumed constraints, legacy cut-off grades, and inherited capacities are the inertia that quietly caps value. By making these assumptions explicit and testable, and by revisiting them each planning cycle rather than defaulting to historical practice, the framework applies Goldratt's process of ongoing improvement to strategic mine planning.

Stochastic methods evaluate uncertainty within a given constraint set, whereas this framework focuses on understanding, testing, and refining the constraints themselves. The two approaches are complementary and together support more informed strategic decisions. A bookend matrix optimised stochastically delivers curated alternatives with a consideration of geological and price risk. In practice, that combination is where the next generation of strategic planning work sits—well-framed alternatives, optimised stochastically, reported in a concise manner so that leaders can decide between alternatives with confidence.

CONCLUSIONS

Strategic mine planning delivers the most value when treated as decision support, rather than a routine modelling output. With this mindset, as the plan moves forward into detailed technical work,

it helps ensure the analysis remains focused on its strategic intent, rather than becoming lost in the details. The literature on integrated-optimisation and uncertainty has improved strategic planning by developing tools and methods that can determine higher value and lower risk scenarios inside a given constraint set. However, decision analysis shows that alternatives are often the largest unused levers. The framework presented here operationalises that lever for strategic mine planning. It is a structured workflow that can be applied equally to greenfield, brownfield, and life-of-asset reviews, as well as annual planning cycles. Constraint triage exposes what can be challenged and what is genuinely fixed. Bookend scenarios bind the feasible planning space without an excessive scenario list. Communication to inform decisions can shift the conversation from a selection between schedules to a deliberate choice on which constraints to challenge, and at what cost. As a result, strategic planning can support more timely, transparent, and defensible investment decisions under uncertainty.

ACKNOWLEDGEMENTS

The author would like to acknowledge the many teams involved in strategic planning across Rio Tinto, whose work and collaboration underpin the ideas reflected in this paper.

Finally, I acknowledge Rio Tinto for permission to publish this paper.

REFERENCES

Balci, M, Cetin, E and Erkayaoglu, M, 2025. Mine planning optimisation under circular reference loop characteristics, *Mining, Metallurgy and Exploration*, 42:1–18.

Coutts, B, 2009. Integrated strategic planning at EKATI Diamond Mine, in *Proceedings of the Orebody Modelling and Strategic Mine Planning Conference 2009*, 8 p (The Australasian Institute of Mining and Metallurgy: Melbourne).

Dimitrakopoulos, R, 2018. Stochastic mine planning: methods, examples and value, in *Advances in Applied Strategic Mine Planning*, pp 101–128 (Springer: Cham).

Goldratt, E M and Cox, J, 1984. *The Goal: A Process of Ongoing Improvement* (North River Press: Great Barrington).

Howard, R A and Abbas, A E, 2015. *Foundations of Decision Analysis* (Pearson: Harlow).

King, B, 2014. Integrated strategy optimisation for complex operations, in *Proceedings of the Orebody Modelling and Strategic Mine Planning Conference 2014*, 10 p (The Australasian Institute of Mining and Metallurgy: Melbourne).

Wharton, C, 2009. Coal mining in the Hunter Valley – a strategic mine planning case study, in *Proceedings of the Orebody Modelling and Strategic Mine Planning Conference 2009*, 8 p (The Australasian Institute of Mining and Metallurgy: Melbourne).

Mind the gap – aligning strategic vision with technical evidence in strategic mine planning

A Suchak[1]

1. MAusIMM, Principal Mining Engineer, BHP, Perth WA 6155. Email: amit.suchak@bhp.com

ABSTRACT

In most mining companies, strategic planning and opportunity assessment processes begin shaping the direction for future growth and development well before any formal study work is undertaken. These activities draw on portfolio drivers, geological understanding, technical indicators and corporate strategy to define focus areas. As strategic teams explore these possibilities, they develop an internal understanding of potential pathways well in advance of the level of technical evidence required to support their appropriate representation in external communication. This creates a structural gap between early internal direction setting and ability to communicate scenarios externally in a manner that reflects their technical maturity. This paper examines this planning gap and introduces the Strategic-Evidence Alignment (SEA) Framework as a practical approach.

This gap has become more noticeable with increasing complexity in strategic mine planning. It is now standard practice to incorporate system level optimisation, technology transitions, decarbonisation requirements, supply chain considerations and cross functional interactions that extend beyond traditional engineering scopes. While these approaches increase internal insight, resulting scenarios often remain below the confidence thresholds required for consistent external communication. Study phases (Scoping, Pre-feasibility, Feasibility) remain key engineering milestones while project toll gates define funding and progression. However, neither was designed to judge strategic strength or communication suitability of early-stage scenarios. This leaves organisations without a structured way to distinguish between what is strategically valuable internally and how it should be communicated externally. Observations from recent industry disclosures illustrate this behaviour where scenarios required clarification or refinement to align communication with evidence maturity.

This paper formalises this challenge as the Strategic–Evidence Gap and applies the SEA Framework, built around Strategic Readiness, Evidence Readiness, and Communication Discipline. Together, these elements align strategic intent with supporting evidence, improving clarity and discipline in communication of strategic mine plans.

INTRODUCTION

Strategic mine planning has evolved from a largely technical exercise into a cross functional decision process, particularly in complex planning environments that influence business direction, asset positioning, investment timing and enterprise risk. In many organisations, opportunity identification is initiated through portfolio reviews, life-of-mine updates, changes in geological understanding, equipment strategy, infrastructure constraints, evolving market conditions, and decarbonisation imperatives. These inputs often drive strategic thinking well before associated concepts are mature enough to support formal reporting or external disclosure.

As a result, internally coherent scenarios frequently emerge at an early stage and are sufficiently robust to influence long-term direction, even though their technical basis (eg geological confidence, engineering definition, and cost assumptions) continues to evolve. These scenarios are developed to support internal decision-making and value optimisation, rather than primarily for external communication.

This creates a persistent gap, referred to as the Strategic–Evidence Gap, which presents a practical dilemma for organisations. Strategically important scenarios may be withheld from external communication where supporting evidence is insufficient, resulting in under-representation of long-term value and optionality. Conversely, when early-stage scenarios are communicated prematurely, they risk being interpreted as misleading or insufficiently substantiated. This dynamic is observed across the industry as reflected in disclosures from both early-stage developers and established

producers (Larvotto Resources, 2024; QMines Limited, 2024; Patagonia Lithium Limited, 2025a, 2025b; Pursuit Minerals Limited, 2025a, 2025b; Resolute Mining Limited, 2026).

The Strategic-Evidence Gap introduces both governance and communication risk, as strategic decisions and external narratives may evolve ahead of the technical evidence required to support them. The challenge is therefore not simply one of timing, but of appropriately representing scenario maturity – what can be communicated, how clearly it can be expressed, and how this aligns with the robustness of supporting evidence.

This gap affects both external communication and internal alignment, as scenarios are interpreted across technical and non-technical functions at different assumed levels of maturity and certainty. It also highlights the limitation of current study-stage and reporting practices, which focus on technical maturity and reporting credibility but do not explicitly address how early-stage strategic scenarios should be evaluated and represented prior to formal technical validation.

This paper addresses that limitation by introducing the Strategic–Evidence Alignment (SEA) Framework, which provides a structured approach to aligning strategic intent, evidence maturity, and communication suitability across decision-making and external communication. The framework is primarily intended for strategic mine planners responsible for shaping and evaluating long-term development pathways, while enabling alignment between technical teams, governance processes, and stakeholder communication functions.

The SEA Framework introduced in this paper is informed by, and builds upon, established industry classification systems (JORC, 2012) for study maturity and reporting rigour.

THE PLANNING GAP AND STRUCTURAL DRIVERS

The planning gap referred to in this paper as the Strategic–Evidence Gap arises because strategic insight (ie identification of preferred development pathways) develops faster than technical certainty. When a planning team models a future pathway, it may already be evident that a particular pit sequence, underground access concept, processing route, or infrastructure configuration offers superior long-term value. However, the supporting technical work may still be at an early stage of maturity.

This relationship between the breadth of internally developed strategic scenarios and the subset that can be communicated at a given stage is illustrated conceptually in Figure 1.

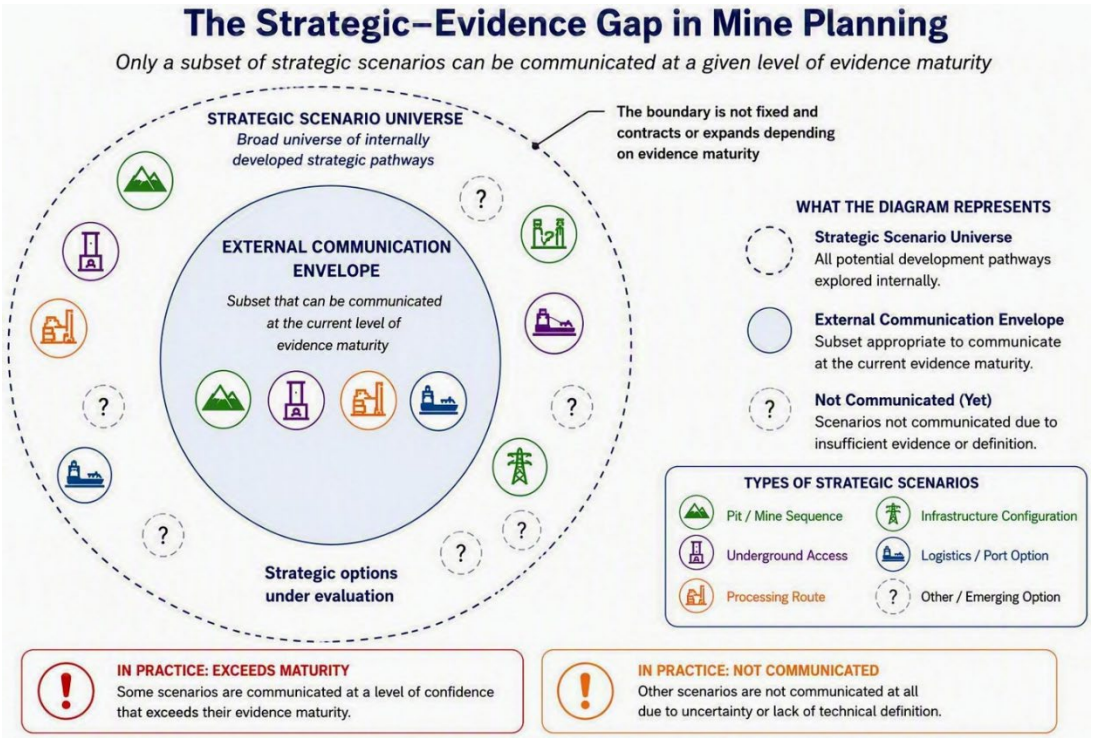


FIG 1 – Strategic scenarios and communication envelope.

As shown in Figure 1, strategic mine planning generates a broad universe of potential development pathways, many of which are directionally credible and inform long-term value creation. The external communication envelope represents a subset of this scenario universe, where the level and specificity of what is communicated should reflect the maturity of supporting technical evidence. However, in practice, this relationship is not consistently maintained. Scenarios are often communicated at a level of confidence that exceeds their evidence maturity while others are not communicated at all due to uncertainty or lack of technical definition.

The difference between the broader scenario universe and the communicable subset defines the practical expression of the planning gap. This is not a limitation of planning capability, but a structural outcome of how strategic insight develops ahead of technical validation.

Typical areas where technical maturity lags strategic clarity include (JORC, 2012):

- geological confidence and resource definition
- geotechnical and hydrogeological parameters
- metallurgical performance and recovery assumptions
- infrastructure feasibility and integration
- capital and operating cost certainty.

This dynamic becomes more pronounced in complex mining systems. A change in mine sequencing may alter energy demand, emissions profile, haulage requirements, water balance, and supply chain exposure. Similarly, a technology transition may unlock additional value while introducing operational and implementation risks.

As these interdependencies increase, maturity becomes multi-dimensional, with different components of the system progressing at different levels of confidence rather than along a single sequential pathway. Modern mine planning is therefore characterised by strong interaction between geological, mining, processing, infrastructure, environmental, and commercial domains. Strategic scenarios emerge as integrated system-level constructs rather than isolated technical ideas.

As scenarios evolve, their components do not mature uniformly. Some elements may achieve relatively high confidence early, while others remain conceptual. This uneven maturity introduces escalation risk, where confidence in the overall strategic direction increases ahead of the maturity of the underlying technical inputs.

Decisions are increasingly made using partially mature systems, which is where misinterpretation begins. Stakeholders may not distinguish between what is demonstrated, what is assumed, and what remains under investigation. Without explicit separation, early-stage scenarios can be perceived as established outcomes.

This interplay between system interdependence, continuous development, and uneven maturity is a primary mechanism underlying the Strategic–Evidence Gap.

LIMITATIONS OF EXISTING PLANNING AND COMMUNICATION FRAMEWORKS

Study phases

Study phases provide a structured approach to developing engineering definition. Scoping studies test options, pre-feasibility studies refine alternatives, and feasibility studies confirm preferred pathways.

However, these stages are designed to increase technical certainty and enable comparison of scenarios, but do not assess strategic relevance, portfolio alignment or timing sensitivity of a scenario. As a result, a concept may be strategically important despite being at an early stage of engineering maturity, while a technically advanced option may no longer align with broader business priorities.

In practice, this often leads to informal reliance on strategic judgement outside the formal study framework, particularly in early-stage planning where engineering definition has not yet caught up with strategic intent. This becomes particularly relevant when these early-stage strategic positions are subsequently considered for external communication, often driven by the need to demonstrate future growth potential or reflect strategic direction, with the extent and level of detail governed by organisational risk appetite.

Project toll gates

Project toll gates provide governance over capital allocation and project progression. They ensure that investment decisions are justified, risk-managed, and appropriately controlled. However, they primarily address readiness for funding rather than readiness for communication. A project may progress through a gate based on its investment case, while still requiring careful qualification in how it is described externally.

This distinction is critical, as funding approval does not necessarily imply that the supporting assumptions are sufficiently mature for full external representation. This creates a disconnect between investment decision and how project maturity is communicated externally.

Communication of strategic mine plans to stakeholders

While study phases and toll gates define internal development and governance processes, communication of strategic mine plans represents the external interface of these decisions.

Effective communication of strategic mine plans requires clearly defined assumptions, appropriate levels of supporting evidence, and transparent articulation of uncertainties. These elements are essential for maintaining credibility with a range of stakeholders including investors, partners, regulators and internal decision-makers.

However, early-stage strategic scenarios are not developed to meet these communication expectations as a primary objective. Even when internally robust, they may lack sufficient supporting evidence to justify external communication at the same level of detail or confidence.

If not carefully managed, this can lead to a situation where preliminary scenarios are communicated in a manner that implies a level of certainty not yet supported by technical evidence. This increases the risk of misinterpretation, unintended expectations and perceived commitments or loss of credibility with stakeholders.

Critical observation

Together, these structures provide strong engineering definition, governance, and reporting controls. However, they do not explicitly distinguish between:

- strategic robustness
- evidence sufficiency (ie the degree to which technical support is sufficient to underpin the level and specificity of external communication)
- communication suitability.

As a result, the transition between internal planning and external communication is not formally governed. This reinforces the challenge identified in stakeholder communication, where internally developed scenarios may be communicated without clear alignment to the maturity of supporting evidence. In practice, this is not a binary assessment, but one of proportional alignment where the level of communication should reflect the maturity of the underlying technical evidence.

This limitation explains the persistence of the Strategic–Evidence Gap identified earlier.

Evidence–communication misalignment in practice – lessons from recent industry disclosures

The limitations outlined above are not only conceptual but are observable in practice. Recent public disclosures, including Australian Securities Exchange (ASX) announcements from 2025 and 2026,

demonstrate a consistent pattern of misalignment between communication and technical maturity, consistent with regulatory guidance on the disclosure of forward-looking statements (ASIC, 2024). These disclosures are drawn from publicly available information. The purpose of these examples is not to evaluate individual disclosures, but to illustrate recurring structural patterns in how strategic scenarios are communicated relative to evidence maturity. While these disclosures provide a clear reference point, the behaviour is not jurisdiction-specific or commodity-specific. The same pattern is observed across the industry where strategic scenarios are developed and communicated ahead of full technical validation.

Three representative cases from recent ASX disclosures illustrate this pattern:

Case 1 – Maiden Resource with Embedded Production Target (Greenfield Project): A company included a 15-year mine life and 10 000 t per annum (t/a) production rate in its January 2025 (Patagonia Lithium Limited, 2025a) announcement of a maiden Inferred Mineral Resources reported under the JORC Code (JORC, 2012). A subsequent retraction announcement stated that these figures were ‘conceptual in nature’ and not compliant with reporting requirements (Patagonia Lithium Limited, 2025b). The underlying Mineral Resource was not retracted; however, the production target language was not supported by the required technical basis.

Case 2 – Throughput Scenarios Interpreted as Production Targets (Development Project Update): A company issued a project update in August 2025 referencing multiple throughput scenarios (5000 t/a, 12 500 t/a, 17 500 t/a) in forward-looking development framing (Pursuit Minerals Limited, 2025a). A subsequent retraction announcement clarified these figures were ‘aspirational’ and based on ‘conceptual planning work,’ not compliant production targets under ASX Listing Rule 5.16 (ASX, 2019) (Pursuit Minerals Limited, 2025b). The strategic development pathway remained unchanged, but the level of communication exceeded evidence maturity.

Case 3 – Scoping Study with 100 per cent Inferred Resources (Established Operating Miner): A company released a scoping study announcement in May 2026 including production targets and financial assessments (Resolute Mining Limited, 2026). The following day, the company retracted the announcement, acknowledging production targets were based on 100 per cent Inferred Resources with no reasonable basis under ASX Listing Rule 5.16 (ASX, 2019). This case is particularly notable as it involved an established operating miner rather than a junior explorer, demonstrating the pattern extends beyond early-stage companies.

A consistent feature across these cases is that the underlying strategic concept was not invalidated. The scenarios remained relevant to long-term planning and in many instances continued to guide subsequent technical work. What changed was the recognition that the level of external communication exceeded the maturity of the supporting evidence.

The common thread is the introduction of quantitative detail (production rates, mine life, throughput) before evidence maturity supported it. This confirms the pattern is structural to strategic mine planning practice, not limited to specific company types, project stages, or commodity sectors.

These examples provide observable instances of how the absence of a formal mechanism to align strategic intent, evidence maturity and communication leads to regulatory intervention or disclosure refinement. The way in which this misalignment progressively develops is described in the following section.

Escalation mechanism

Interpretation often escalates faster than evidence maturity. The cases outlined above show that misalignment between evidence and communication often develops as a gradual progression in how scenarios are expressed over time. This progression reflects how planning outputs are progressively refined and communicated, rather than discrete transitions between defined stages.

In practice, scenarios are first developed for internal planning as conditional constructs, supported by partial and evolving technical inputs. As these scenarios gain internal traction, elements of them begin to enter external communication, often to provide transparency around future opportunities or strategic direction.

During this transition, the level of detail included in communication tends to increase incrementally. This progression can be described as:

Options → Expectations → Perceived Commitments

Options represent internally developed scenarios used to explore possibilities. At this stage, assumptions are evolving, and technical support may be uneven across different components of the scenario.

Expectations emerge when these scenarios are communicated externally with additional definition—such as indicative production ranges, development sequencing, or timing assumptions.

Perceived commitments arise when stakeholders interpret this increased specificity as representing a defined or endorsed development pathway.

A critical feature of this progression is that the perceived level of evidentiary support increases with the level of communication detail, regardless of whether the underlying technical maturity has changed.

This is consistent with the examples in the section above where incremental additions of detail particularly quantitative or time-bound information materially influenced how scenarios were interpreted. In these cases, the introduction of production targets or development parameters conveyed a level of certainty that was not yet supported by sufficient technical definition.

Importantly, this escalation is rarely intentional. It occurs as a natural consequence of:

- the need to provide clarity when communicating complex scenarios.
- the selective progression of certain technical components ahead of others
- the use of specific numeric parameters to support strategic narratives.

As a result, a scenario that is internally understood as conditional can be interpreted externally as a commitment once presented with sufficient specificity.

This divergence between internal language and external interpretation is summarised in Table 1, which shows how increasing evidentiary maturity can be perceived differently by internal and external stakeholders.

TABLE 1

Escalation mechanism in mine planning communication.

Evidence maturity	Internal language	Typical external interpretation
Low	Concept, option, scenario	Early-stage opportunity (limited confidence)
Moderate	Indicative pathway, ranges, preliminary parameters	Expected outcome (increasing confidence implied)
High (but incomplete)	Defined pathway, emerging target (eg production rates, timing, capacity)	Interpreted as a committed development pathway or production outcome

This divergence reflects a mismatch between internal awareness of evidence maturity and external perception of certainty.

The escalation mechanism explains why the cases described above required subsequent clarification or withdrawal. In those cases, the introduction of quantitative parameters such as production rates, mine life, or throughput corresponds to a shift from ‘expectations’ to ‘perceived commitments.’ Although the underlying evidence maturity remained unchanged, the level of specificity conveyed a degree of definition that led external stakeholders to interpret the scenarios

as established outcomes. This demonstrates how communication, rather than technical advancement, can drive escalation in perceived certainty.

In the absence of a structured approach to aligning the level of communication with the maturity of technical evidence, this progression can occur without explicit control. As a result, scenarios may be unintentionally presented in ways that imply a higher degree of definition or certainty than is supported.

STRATEGIC–EVIDENCE ALIGNMENT (SEA) FRAMEWORK AND CONCEPTUAL MODEL

Framework overview

The Strategic–Evidence Alignment (SEA) Framework addresses the limitation identified above by explicitly separating three distinct but interdependent dimensions of scenario development and communication, as summarised in Table 2.

TABLE 2

SEA Framework.

Element	Purpose	Key question	Primary application
Strategic readiness	Assess long-term value and alignment	Is this strategically robust?	Internal planning
Evidence readiness	Assess technical support	Is it sufficiently supported?	Technical evaluation
Communication discipline (alignment)	Guide representation	How should it be expressed?	Disclosure and stakeholder engagement

The framework recognises that, in practice, these dimensions are often treated as interchangeable. This can result in scenarios being communicated at a level of specificity that exceeds the maturity of the supporting evidence.

SEA addresses this by requiring explicit and separate evaluation of each dimension, particularly at the point where scenarios transition from internal planning to external communication.

Importantly, these dimensions are not sequential but operate simultaneously. Scenarios must be assessed across all three elements when transitioning from internal planning to external communication. A scenario may be:

- strategically compelling but not yet evidentially supported.
- technically advanced but not aligned with broader strategy.
- sufficiently developed for internal use but not appropriate for external communication at a high level of specificity.

The key contribution of the framework is that it introduces a structured link between Evidence Readiness and Communication Discipline, ensuring that the level and specificity of communication remain proportionate to the maturity of supporting technical evidence.

This enables decision-makers to assess not only whether a scenario is viable or fundable, but also whether it is appropriate to communicate, and at what level of detail.

The framework does not introduce additional stage gates or replace existing study processes. It operates continuously across scenarios, separating strategic value, evidence maturity, and communication so that each can be evaluated independently while remaining aligned in practice.

Framework components

Strategic Readiness evaluates whether a scenario should influence long-term planning and decision-making. It requires a clear understanding of how the scenario aligns with corporate objectives and portfolio priorities, contributes to long-term value creation, and integrates across mining, processing, infrastructure, environmental, and market considerations. It also requires explicit recognition of value sensitivity to key drivers. A strategically ready scenario may not be fully defined, but it has sufficient coherence and internal consistency to inform direction and guide further work. This ensures that scenario evaluation remains driven by value and strategic relevance, rather than being constrained prematurely by evidence maturity.

Evidence Readiness is inherently multi-dimensional, requiring consistent assessment across geological, engineering, economic, and integration parameters. It assesses the level of technical maturity and reliability underpinning a strategic scenario. It focuses on the strength, quality, and completeness of the supporting technical work, including:

- geological confidence and resource definition
- maturity of mine design and processing assumptions
- infrastructure feasibility and integration
- cost estimation reliability
- identification and treatment of key risks.

The objective is to establish a clear and transparent view of what is known, what is assumed, and what remains uncertain within the scenario. This provides a transparent view of evidence maturity, forming the basis for appropriate interpretation and communication.

Communication Discipline requires that the specificity of externally communicated parameters must not exceed the maturity of the underlying evidence supporting them. It governs how scenarios are expressed once they move beyond internal planning. It determines the appropriate level and specificity of communication based on the maturity of the underlying technical evidence.

This includes:

- distinguishing between strategic intent, conceptual development, and technically supported outcomes.
- aligning granularity of disclosed assumptions (eg production rates, timing, sequencing) with evidence maturity.
- appropriately qualifying assumptions, ranges, and uncertainties.
- avoiding language that implies certainty beyond what is supported.

The objective is not to restrict communication, but to ensure that it is proportionate to the supporting evidence, preventing unintended escalation from options to perceived commitments.

In practice, Communication Discipline is demonstrated through:

- explicit qualification of uncertainty
- controlled use of parameters specific detail
- clear separation between scenarios and defined plans.

The integrated relationship between Strategic Readiness, Evidence Readiness, and Communication Discipline is shown in Figure 2.

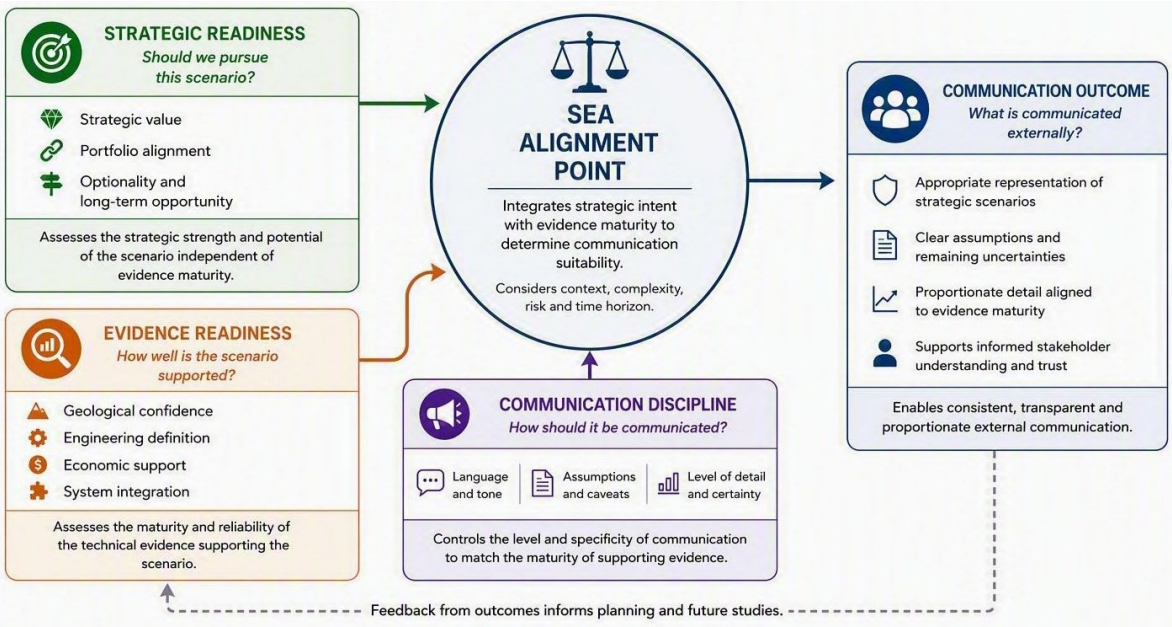


FIG 2 – SEA Framework architecture.

Figure 2 illustrates how the three SEA components operate as an integrated framework. Strategic Readiness assesses the value and alignment of a scenario, while Evidence Readiness evaluates the maturity of its technical support. Communication Discipline governs how the scenario is represented by aligning the specificity and certainty of communication with the available evidence. Their integration at the SEA Alignment Point determines an appropriate communication outcome and provides feedback for subsequent planning and technical work.

Conceptual model

The interaction between strategic development, technical maturity, and communication can be understood as a dynamic relationship. Strategic scenarios typically develop early and begin influencing decision-making, while technical evidence builds progressively through study work.

The Strategic–Evidence Gap exists where strategic clarity outpaces evidence maturity. Communication operates at this interface and must remain aligned with both.

Figure 3 illustrates this relationship, highlighting how strategic insight and technical evidence evolve at different rates, and how communication can drift ahead of technical support in the absence of explicit control. This reflects the escalation mechanism described earlier, where communication evolves from options to expectations and ultimately perceived commitments without a corresponding increase in evidence maturity.

The Evidence Readiness Threshold referred in Figure 3 as communication discipline boundary is conceptual and does not represent a fixed regulatory or study-stage boundary. In a multi-asset or multi-project organisation, individual projects may follow different trajectories of strategic insight and technical evidence development. The relationship shown in Figure 3 should therefore be interpreted as a conceptual model, with the magnitude and evolution of the Strategic–Evidence Gap varying across projects depending on their maturity, complexity, and context.

The SEA Framework provides a mechanism for managing this relationship by aligning Strategic Readiness, Evidence Readiness, and Communication Discipline to govern the transition from internal scenario development to external communication.

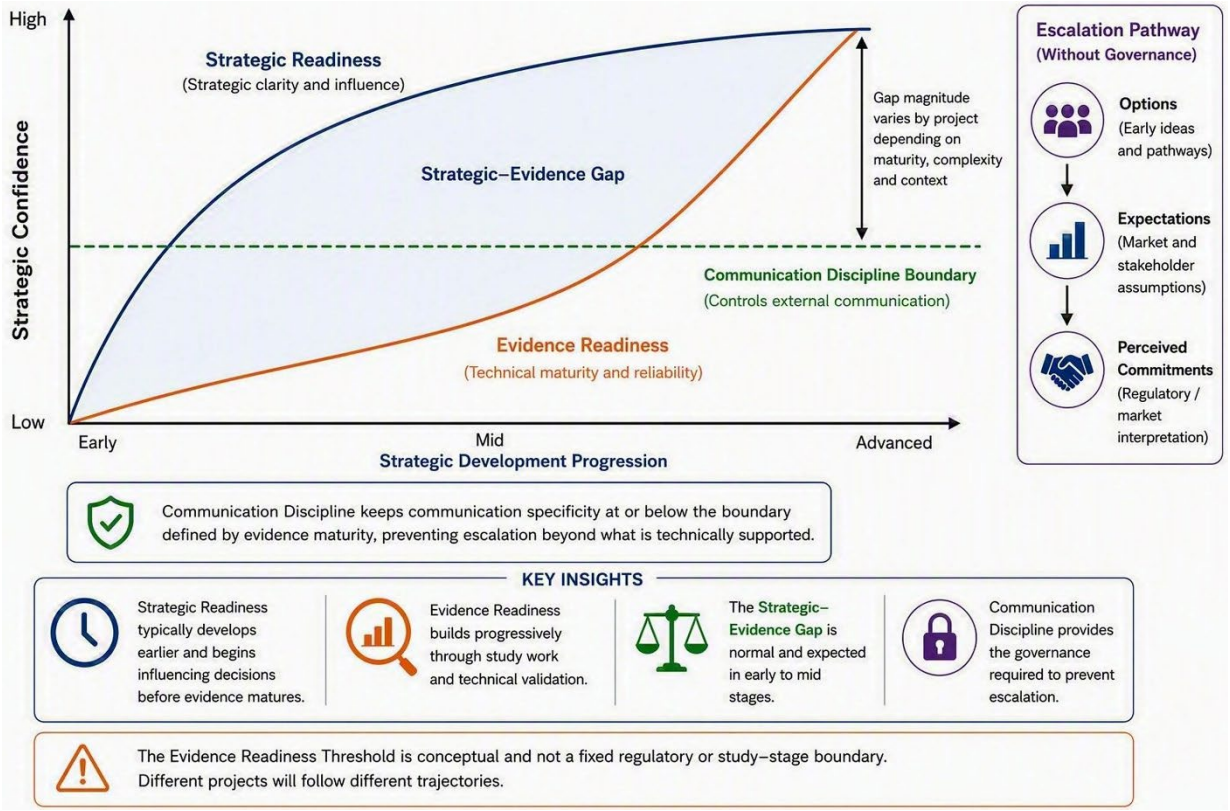


FIG 3 – Conceptual model of Strategic–Evidence Gap and communication escalation.

Practical application in strategic mine planning

Consider an operating mine identifying an adjacent opportunity with potential to extend reserve life and materially improve project economics. Early modelling indicates a potentially high-value pathway, supported by coherent strategic logic but limited technical maturity. Metallurgical test work is ongoing, geotechnical confidence is still being established, infrastructure assumptions remain conceptual, and cost estimates are sensitive to unresolved design elements.

This reflects a typical planning situation where strategic direction emerges ahead of full technical validation across all system components. In practice, this situation typically evolves along two distinct pathways: progression without explicit alignment between scenario maturity and communication, and progression under structured application of the SEA Framework.

Without the SEA Framework, such scenarios are often advanced and communicated based on their perceived value, with the maturity of supporting evidence implicitly assumed rather than explicitly assessed. As observed in practice, this frequently leads to the introduction of quantitative detail or development implications that exceed the underlying technical basis. As communication becomes more specific, the scenario is increasingly interpreted as a defined outcome, increasing the likelihood of misinterpretation and subsequent revision.

Applying the SEA Framework changes how this progression is managed. With the SEA Framework, the scenario is evaluated explicitly across three independent but aligned dimensions.

Under Strategic Readiness, the opportunity is recognised as highly compelling, with strong alignment to long-term objectives and clear potential to influence future asset direction.

Under Evidence Readiness, the scenario is identified as intermediate in maturity, with several critical elements not yet sufficiently defined to support detailed interpretation. This makes the uncertainty explicit, rather than allowing it to remain embedded within assumptions.

Under Communication Discipline, the organisation continues to develop the opportunity, while ensuring that communication reflects the current level of evidence maturity. Language remains aligned with uncertainty, avoiding implication of a defined operating state or committed development outcome.

For example, stating ‘potential to support a defined expansion at a specified production rate’ implies a level of technical certainty not supported by early-stage metallurgical or infrastructure assessment, whereas describing a ‘conceptual expansion opportunity under evaluation’ maintains alignment with evidence maturity.

Similar dynamics occur across a range of planning contexts, including expansion strategies, satellite deposit integration, and technology or infrastructure transition pathways. In each case, strategic clarity may emerge early, while key inputs such as resource confidence, modifying factors, or system integration remain unresolved. The framework enables these opportunities to continue guiding planning, while preventing communication from implying a level of definition not yet supported.

The effect on planning practice is direct. Opportunities are preserved as strategic pathways, while effort is directed toward resolving the specific uncertainties that limit evidence maturity. Metallurgical performance, geotechnical validation, and infrastructure feasibility become targeted workstreams rather than implicit assumptions embedded within a preferred scenario.

This approach ensures that organisations maintain strategic direction while systematically strengthening the underlying evidence, without allowing communication to outpace technical maturity. For decision-makers, this complements existing study phases and project governance by making explicit the relationship between scenario value, evidence maturity, and the appropriate level of communication.

The three cases presented earlier illustrate the Strategic–Evidence Gap at an industry level. Table 3 applies the SEA Framework to Case 1 as a pre-release disclosure review, showing how Strategic Readiness, Evidence Readiness and Communication Discipline could be assessed before market release. The purpose is to demonstrate how a single SEA review point can convert an emerging disclosure risk into a clear communication decision: what can be released, what must be qualified, and what should be withheld until further evidence is available.

TABLE 3

Pre-release application of SEA Framework to Case 1.

SEA component	Readiness rating	Assessment	SEA summary
Strategic readiness	High	The maiden Resource supported a credible strategic opportunity and justified continued internal planning and study work.	SEA would have allowed the Resource to be released but would have prevented fixed production rate and mine life statements because communication detail exceeded evidence maturity.
Evidence readiness	Low	The evidence supported reporting the Inferred Mineral Resource, but not a production target, mine life or development case.	
Communication discipline	Not ready	The proposed wording used quantitative production and mine life detail that could be read as a supported production target.	

In practical terms, the SEA Framework can be integrated within existing planning and governance processes at the point where scenarios transition from internal evaluation to broader communication. This includes strategic planning cycles, study stage reviews, and disclosure assurance processes. At planning cycle gates, the SEA assessment determines whether a scenario has sufficient Evidence Readiness to be elevated from internal evaluation to external reporting. At study stage reviews, it provides a structured checkpoint that aligns technical deliverables including resource classification, cost estimation, and Modifying Factor assessment with the communication intent of the sponsoring team. For disclosure assurance, the framework provides an auditable record of the evidence basis underpinning each production target or strategic scenario released to the market, directly supporting compliance with ASX Listing Rule 5.16 and ASIC guidance on forward-looking statements.

In practice, this assessment is typically led by planning functions, with input from technical specialists and governance teams, ensuring that maturity is not interpreted from a single perspective. In this way, the framework operates as an overlay to existing processes, clarifying how scenario maturity is interpreted rather than introducing additional procedural complexity.

Organisational implications and limitations

At an organisational level, the SEA Framework changes how scenarios are prioritised, progressed, and represented across the planning life cycle.

By explicitly separating strategic value from technical maturity, it enables more disciplined prioritisation of opportunities and clearer alignment between planning effort and the uncertainties that most influence project outcomes.

In contrast to existing study phases and project toll gates—which govern technical progression and investment decisions—the SEA Framework extends governance into how scenarios are interpreted and represented. By linking evidence maturity to scenario representation, it introduces a control over interpretation that is not addressed within conventional planning structures.

This has particular relevance for forward-looking narratives, where consistency between language and underlying evidence reduces the risk of misinterpretation and unintended escalation. More broadly, the framework supports consistent interpretation of scenarios across technical, strategic, and leadership functions, particularly in organisations managing multiple assets or competing development pathways.

The framework also reinforces professional discipline by explicitly distinguishing between confidence grounded in evidence and assumptions embedded within early-stage scenarios, ensuring that representations of value remain aligned with underlying maturity.

However, the framework has limitations. It does not eliminate the need for judgement, and outcomes remain dependent on the experience, discipline, and alignment of those applying it. Effective use also requires integration with existing study workflows and governance processes, rather than operating as a parallel system.

The framework does not define fixed thresholds for communication. This is intentional, as appropriate levels of representation depend on project stage, risk appetite, regulatory context, and market conditions. Consistent application therefore relies on internal alignment and clear interpretation.

In addition, the framework does not prescribe formalised methods for measuring maturity across Strategic Readiness, Evidence Readiness, and Communication Discipline. Developing consistent criteria, scoring approaches, or assessment metrics remains an area for future work. While flexibility is an inherent strength, more structured approaches to maturity assessment would support broader application.

Finally, the effectiveness of the framework depends on adoption across strategic, technical, and communication functions. Partial implementation will not address the misalignment identified in previous sections.

Future development should focus on embedding the framework within planning systems, governance processes, and decision-support tools. The primary opportunity, however, lies in consistent application. When applied effectively, the SEA Framework provides a practical means of maintaining alignment between scenario value, technical maturity, and how that value is represented.

CONCLUSIONS

Strategic mine planning increasingly operates in a context where direction emerges ahead of full technical validation. This creates a persistent gap between what is strategically compelling and what can be supported with sufficient technical confidence, resulting in either under-communication of value or the overstatement of certainty.

The Strategic–Evidence Alignment (SEA) Framework addresses this challenge by explicitly distinguishing between strategic robustness, evidence maturity, and communication suitability, and by establishing a structured link between the maturity of supporting technical evidence and the level of communication applied.

In doing so, the framework extends beyond conventional study phases and project governance, which primarily address technical progression and investment decisions. SEA introduces a complementary capability ensuring that the value of a scenario is communicated in a manner that is proportionate to its evidence maturity.

Importantly, this is not only an external consideration. Within organisations, strategic scenarios are interpreted across diverse technical and non-technical functions, each bringing different perspectives and assumptions where communication does not reflect the maturity of underlying evidence, this can lead to inconsistent interpretation of value, misaligned expectations, and inefficient effort.

The examples presented demonstrate that misalignment between evidence and communication is not theoretical, but a recurring feature of modern mine planning. The challenge for the industry is therefore not whether strategic insight can emerge early, because it clearly can, but whether it can be consistently interpreted, aligned, and communicated as it evolves.

Effective planning depends not only on identifying future opportunities, but on preserving their value through clear and proportionate communication. The SEA Framework provides a practical mechanism to achieve this by aligning scenario value, technical maturity, and how that value is expressed across the organisation and to external stakeholders.

In this sense, the quality of strategic mine planning is defined not only by what is identified, but by how clearly its value and level of certainty are understood. As planning complexity increases, the ability to control how scenarios are interpreted, not just how they are developed will become a defining capability of effective organisations.

ACKNOWLEDGEMENTS

The author acknowledges the valuable input and constructive discussions from colleagues and stakeholders who contributed to the development of the ideas presented in this work. The author also extends appreciation to BHP and its leadership team for their ongoing support in enabling a work environment that fosters technical development and critical thinking. Special thanks are extended to Marcelo Cortes for his guidance and support, and to the broader team for their insights and feedback throughout the study. The perspectives shared have significantly strengthened the conceptual framework and its practical relevance. The author also acknowledges the use of Microsoft Copilot as a language and drafting support tool in the preparation of this manuscript.

REFERENCES

Australian Securities and Investments Commission (ASIC), 2024. Mining and resources: Forward-looking statements, Australian Securities and Investments Commission. Available from: <https://asic.gov.au/regulatory-resources/takeovers/mining-and-resources-forward-looking-statements/>

Australian Securities Exchange (ASX), 2019. ASX Listing Rules – Chapter 5: Additional reporting on mining and oil and gas activities, Australian Securities Exchange. Available from: <https://www.asx.com.au/documents/rules/Chapter05.pdf>

Joint Ore Reserves Committee, The (JORC), 2012. Australasian Code for Reporting of Exploration Results, Mineral Resources and Ore Reserves (The JORC Code) [online]. Available from: <http://www.jorc.org> (The Joint Ore Reserves Committee of The Australasian Institute of Mining and Metallurgy, Australian Institute of Geoscientists and Minerals Council of Australia).

Larvotto Resources, 2024. Retraction of production targets, ASX announcement. Available from: <https://www.larvotto resources.com/wp-content/uploads/2024/05/61206337.pdf>

Patagonia Lithium Limited, 2025a. Significant Maiden JORC Inferred Mineral Resource Estimate – Lithium Resource of 3.8 million tonnes LCE, ASX announcement, 22 January 2025. Available from: <https://www.patagonialithium.com.au/2025-asx-announcements-1>

Patagonia Lithium Limited, 2025b. Production Target Retraction, ASX announcement, 24 January 2025. Available from: <https://www.listcorp.com/asx/pl3/patagonia-lithium-limited/news/production-target-retraction-3143490.html>

Pursuit Minerals Limited, 2025a. RGS Project Update, ASX announcement, 20 August 2025. Available from: <https://www.marketindex.com.au/asx/pur/announcements/rgs-project-update-6A1279061>

Pursuit Minerals Limited, 2025b. Production Target Retraction, ASX announcement, 9 September 2025. Available from: <https://www.marketindex.com.au/asx/pur/announcements/production-target-retraction-6A1282963>

QMines Limited, 2024. Production target retraction, ASX announcement. Available from: <https://data-api.marketindex.com.au/api/v1/announcements/XASX:QML:6A1270120/pdf/inline/production-target-retraction>

Resolute Mining Limited, 2026. Retraction: Scoping study for ABC Project announcement, ASX announcement. Available from: <https://www.rml.com.au/wp-content/uploads/2026/05/14-05-2026-RSG-ASX-retraction-Statement-004.pdf>

Strategic planning and optimisation

Incorporating mining direction into strategic mine planning – modelling and implementation

I Rudge[1], S Hill[2], D Place[3] and C Poblete[4]

1. R&D Software Engineering Manager, Dassault Systèmes, Brisbane Qld 4000.
   Email: iain.rudge@3ds.com
2. R&D Mine Planning and Optimisation Expert, Dassault Systèmes, Perth WA 6000.
   Email: stephen.hill@3ds.com
3. R&D Software Engineering Director, Dassault Systèmes, Brisbane Qld 4000.
   Email: david.place@3ds.com
4. GEOVIA R&D Apps Portfolio Senior Manager, Dassault Systèmes, Brisbane Qld 4000.
   Email: cristian.pobletematamala@3ds.com

ABSTRACT

The Strategic Mine Planning (SMP) problem considers the scheduling of a mineral deposit represented by a 3D block model. Material movements directly influence the net present value (NPV) of a mining operation, as the timing of ore and waste extraction determines revenue and costs. The block model is divided into bins, which may be grouped into clusters. A block can be represented by a cluster containing a single bin, while in Bench-Based Scheduling a bench may be represented by a cluster containing multiple bins. Each bin is associated with tonnes, grade, and economic value.

In practice, mining progresses according to a defined extraction direction. Existing models generally assume that all bins within a cluster are extracted at the same rate, ignoring horizontal directional precedence. This limits their ability to provide realistic guidance on extraction sequences. To address this limitation, we develop a new model in which bins within a cluster are extracted sequentially, explicitly capturing mining direction. We present the formulation and the modifications required for use with the Bienstock-Zuckerberg (BZ) algorithm along with supporting heuristics.

We investigate the TopoSort heuristic developed by Moreno, Espinoza and Goycoolea and its extension by Letelier, which provides rapid solutions to capacity-constrained scheduling problems. Based on computational observations, we develop 14 TopoSort variants, which can be quickly executed. We further combine these approaches with the Optimised-Destinations heuristic for stockpiling problems and evaluate performance on standard MineLib instances.

For more complex constrained problems, we examine the 1-Dest heuristic described by Letelier, which solves a restricted mixed-integer linear programming (MILP) problem by fixing destinations for each bin. We propose modifications that provide additional destination opportunities, improving solution quality.

Computational results on MineLib and internal benchmark data sets demonstrate improved TopoSort performance and higher NPV solutions from the modified 1-Dest heuristic, albeit with increased runtimes.

INTRODUCTION

The importance of strategic mine planning cannot be over stated. The mine plan determines how a mineral resource is extracted over time, how extracted material will be processed or removed as waste, as well providing an investment schedule. It has a direct impact on the economic value, risk, and overall viability of a mine. Poor planning can render a deposit unprofitable, while well designed plans can significantly improve returns. To support these decisions, mining engineers use specialised software that builds mathematical models of the mine. These models capture constraints such as geology, processing, extraction sequencing, blending and other economic factors. Such models allow complex real world mining problems to be analysed and solved computationally.

Optimisation techniques are central to strategic mine planning. They are used to generate schedules and allocate resources in a way that maximises economic outcomes within the bounds of constraints. The results of these models are used to make long-term planning and investment decisions that impact mine operation and profitability. Due to the size and complexity of strategic mine planning

problems, this area sees much active research and development, with improvements to modelling techniques and algorithms.

This paper extends the work of both the TopoSort and 1-Dest heuristics. TopoSort, originally introduced in Chicoisne *et al* (2012), has been extended previously in Letelier *et al* (2020). We consider variations in key steps of the heuristic and show its impact on net present value (NPV). We also present variants of the 1-Dest heuristic, building upon the work of Moreno *et al* (2017) and Letelier *et al* (2020).

A contribution of this paper is a new modelling technique to capture mining direction on a bench in strategic mine plans. This technique allows for more realistic considerations and more practical results in the planning process.

Optimisation is run on a 3D block model, used to represent the orebody, and derived from geotechnical and resource estimation work. The block model is divided into units of earth called bins, with bins able to be grouped into units called clusters. A block can be represented by a cluster containing one bin. Each bin represents a basic scheduling unit with defined attributes such as tonnes, grade and value. Clusters can be further grouped into working phases. In the case of Bench-Based Scheduling, a bench of a phase can be represented by a cluster containing a number of bins.

MINING DIRECTION

In a mining operation, the mine is often progressed following a mining direction. In current computational modelling, all bins of a cluster are removed at the same rate, which we refer to as the ‘Same Rate’ case. This modelling has computational advantages but ignores horizontal directional precedence. Further, when a cluster is too large to be removed in a single period, and is split across two time periods, the Same Rate modelling falls short. A user is given poor guidance on what sections of the cluster to remove in the first period, and which to remove later.

Better results come from specifying that bins must be removed in a sequence, which we refer to as the ‘Sequenced Bins’ case.

As can be seen in Figure 1, the Sequenced Bins case provides better guidance for the removal of bins within a cluster. Rather than removing all bins in a cluster at 50 per cent in one period, say, and the remainder in the following period, the cluster is able to be split between bins. The bins, following the sequencing are removed in their entirety, as shown in Figure 1, with the clusters split across consecutive periods.

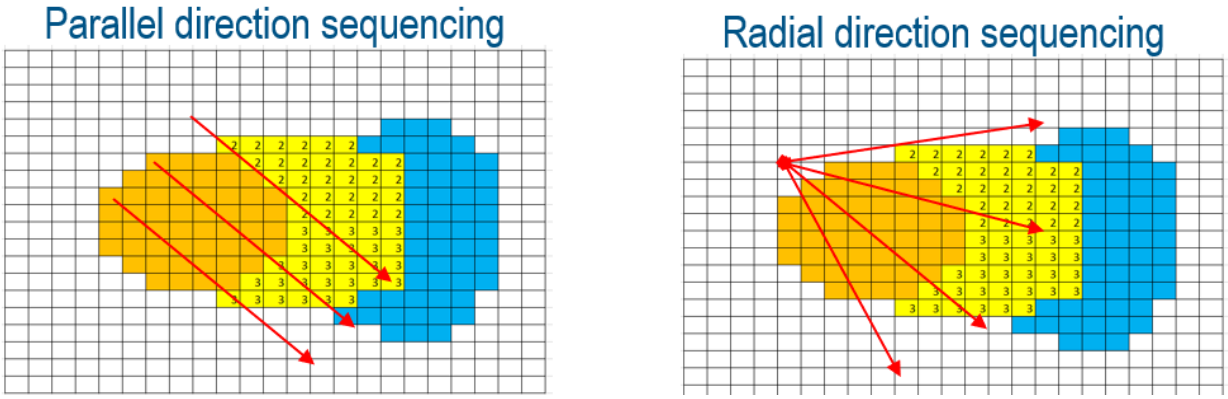


FIG 1 – Parallel and radial sequencing.

The most often used approach to modelling strategic mine planning problems so that they can be optimised using computing resources is via formulating the problem as a mixed-integer linear programming (MILP). Here, wall slopes are enforced by using precedence constraints between variables. These variables represent whether the removal of a cluster has started by a certain period, or not, providing other above clusters have been fully excavated by the period. Other variables and constraints are used to enforce different restrictions on the operations. Examples include blending constraints, capacity totals and CapEx (capital expenditure). Problems are optimised, either exactly or heuristically, to produce solutions that are used in the creation of strategic mine plans.

MILP problems contain a mixture of integer and real variables, which are typically computationally difficult to solve. This is mainly due to the size of the formulation, which can contain many millions of variables and constraints. A common approach, which is used in this paper, is to solve the easier linear programming (LP) problem (drop the constraint that some variables are integer). The LP solution is then used as a guide to solve the original problem heuristically. In the new solver from Dassault Systèmes: GEOVIA, called GMX, we solve the easier linear programming problems using our implementation of the Bienstock-Zuckerberg (BZ) algorithm (Bienstock and Zuckerberg, 2010). For an easy-to-understand description of the BZ algorithm, see Moreno *et al* (2017) and Hill *et al* (2023). We summarise the algorithm in Figure 2.

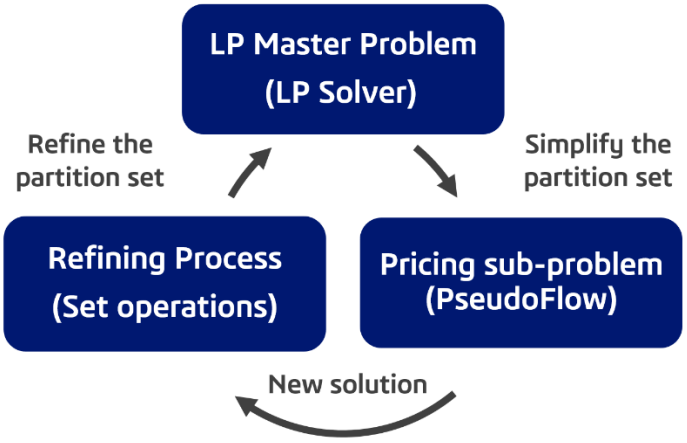


FIG 2 – A summary of the BZ algorithm.

The BZ algorithm uses partitions instead of blocks to quickly solve the LP problem. It uses a decomposition algorithm whose steps we briefly summarise:

- The LP is formulated in terms of the partition set and solved.
- The LP solution gives rise to penalties associated with the set of hard constraints, being those that are not precedence ones. These penalties are used to influence the creation of a Flow problem (which expresses the precedence constraints), solved using the PseudoFlow algorithm (Hochbaum, 2008).
- The PseudoFlow solution is used to refine the set of partitions.
- Repeat until the LP and PseudoFlow objectives converge, no new partitions are created or the dual penalties are unchanged.

For the rest of this paper, let C be the set of clusters, B is the set of bins, D is the set of possible destinations a bin can be sent to and T the set of time periods.

A significant speedup of the BZ-algorithm comes from noting, in the Same Rate case, that the PseudoFlow problem can be solved via solving an equivalent easier problem (see Moreno *et al,* 2017) to understand this ‘path contraction’ operation. For the Sequenced Bins case, we wish to also have this speedup. We provide diagrams and an explanation of our method below.

The use of Sequenced Bins modelling requires a different MILP formulation of the strategic mine planning problem than is commonly used. To best understand this change we provide the following diagram, Figure 3, which illustrates the change to the underlying graph.

The difference is the removal of the arc between the first bin of the cluster to the last one, which causes bins to be removed at a same rate. Also, we add an arc from the first bin, last destination to the last bin, last destination in the following period. This forces the cluster to be removed in at most two periods. We modify the Same Rate path contraction method to contract the graph along paths for the Sequenced Bin case.

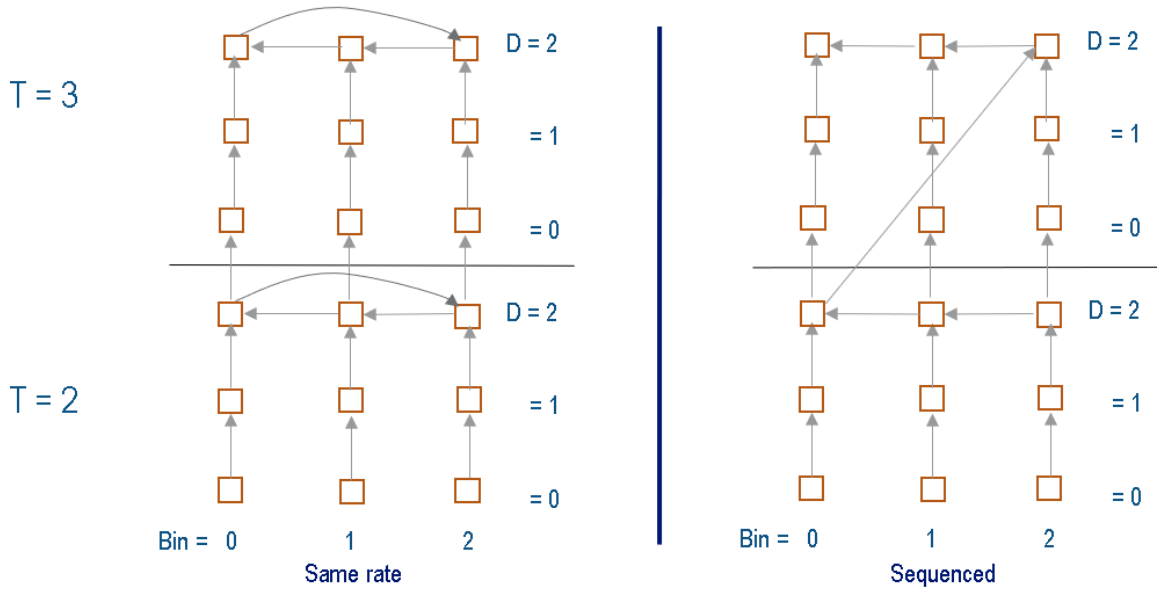


FIG 3 – Precedence graphs for both the Same Rate and Sequenced Bins cases.

In the following we detail how the Path contraction operation can be modified.

When considering the Sequenced Bins path contraction case, we note:

- The only constraints present in the Flow graph are precedence constraints.
- If it is beneficial to split the bins across two periods, rather than remove the whole set of bins in a period then this decision will be based on profitability.
- The most profitable split point can be calculated prior to solving the PseudoFlow problem.

First, for each bin and period, the most profitable destination is chosen. Following this, the costs of removing a bin in a period are added to that bin, this may include the costs from previous periods, as shown in Figure 4. For example, bin 0 has selected destination 0 as the most profitable in time period 2, and destination 1 as most profitable in the following time period. The cost for time period 3, and destination 1 must then include the previous destinations, 1 and 2.

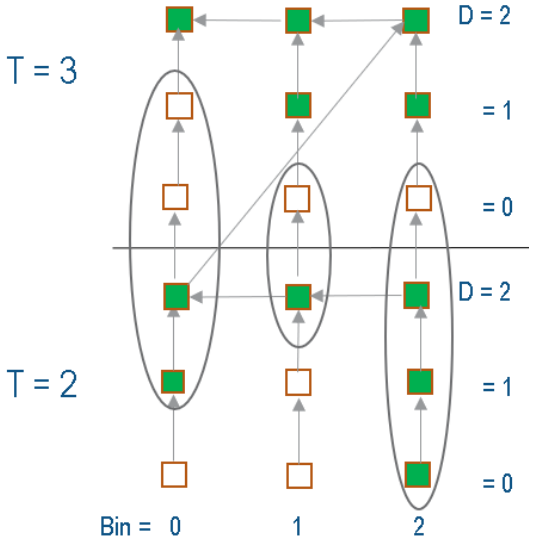


FIG 4 – Grouping destinations per bin to calculate the pseudo-costs.

Once these costs are calculated, the cluster may be split into two, based on what number of bins to remove is most profitable. By applying this method of path contraction, each cluster can be represented by two nodes per period, compared to one for the Same Rate case. Figure 5 shows the

split of a cluster into two nodes. The split is such that bins 0 and 1 are removed in period 2 and bin 2 is removed in period 3. The cluster is split and the costs attributed to each node accordingly.

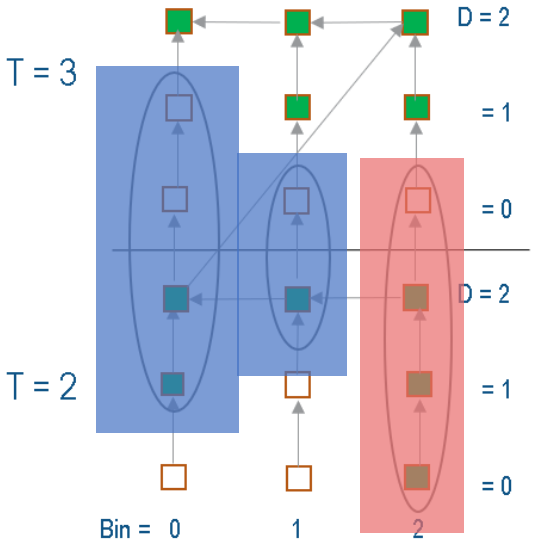


FIG 5 – Grouping a cluster into two nodes.

The path contraction method is used prior to using the PseudoFlow algorithm (Hochbaum, 2008). The number of nodes is larger than that of the Same Rate case, that is the number of nodes is of the order 2 × |C| × |T| compared to |C| × |T|. However, we note that |C| is often less than 500 for bench based scheduling problems in the literature. Additionally solving this smaller problem is preferable to solving the original graph which is of size |B| × |D| × |T| vertices, and so it much larger and more difficult to solve.

GENERATING SOLUTIONS

Once the LP has been solved, using the BZ algorithm with speedups, as implemented in GMX, we use this LP solution to generate a solution to the strategic mine planning problem. Our work follows on from the original techniques presented in (Letelier *et al,* 2020) and represents our investigations and extensions of these ideas.

TopoSort

The original TopoSort heuristic was introduced in Chicoisne *et al* (2012) and then extended in Letelier *et al* (2020). TopoSort can be used for problems that have capacity type constraints, being that all coefficients of the constraint are non-negative with positive maximum capacity. Examples of these constraints are maximum extraction and processing limits. The method can also be extended for use with problems containing stockpiles (Rezakhah, 2017).

The fundamental approach of TopoSort is to order cluster extraction by removing clusters in turn from the current surface, providing capacity constraint limits are respected. The removal of clusters from the surface ensure precedence constraints are satisfied. TopoSort is run after the linear relaxation of the problem has been solved. It uses the results from this linear relaxation as a way to weight different clusters and so guide the order of removal.

Here we present the key steps of the algorithm and refer the interested reader to the paper of Letelier *et al* (2020) for the full algorithm.

TopoSort key steps:

1. Assigning a weight, or value to all clusters. These weights are derived from the solution to the linear relaxation.
2. Create a set of all clusters that are on the surface of the pit.
3. Select the highest weight cluster from those available.

4. For this cluster, an expected time of removal comes from the weight. Assign each bin to its most profitable destination. If the bins exceed the capacity of one of the capacity constraints for this period, then split the cluster across later periods.
5. Update the remaining capacities of each affected constraint. Remove the cluster and update the surface.
6. Return to step 3 until no periods have excess capacity, or no clusters remain.

Because TopoSort effectively chooses one cluster to be removed after another, it is very fast. To contribute to the improvement of this well-known heuristic, we consider variations of the method used to weight the expected start time of clusters. We also consider alternatives to the capacity limits used to enforce constraints. Simple rule changes are applied, which are derived from observations and experimentation through our testing. Then by running several variants, which can be done quickly, we aim for a greater exploration of the solution space and improved solutions without significant increases to the overall runtime.

The variations to TopoSort that we present can be broken down into three main categories:

1. Adjustments to the weighting function.
2. Changes to how capacities are calculated.
3. Simple rule changes within the TopoSort heuristic to express ideas coming from observations during our testing.

We first describe #1 – changing the weighting function changes the order in which clusters are selected. We describe these different weightings as Regular, First, Middle Max and Zero, as described in Table 1.

TABLE 1

Summary of TopoSort weighting types.

TopoSort weighting type	Method of calculation
Regular	The weighting is the sum of each X value in the LP multiplied by the period in which it occurred, as used in Chicoisne *et al* (2012) and Letelier *et al* (2020).
First	The weighting value is equal to the first period in which the cluster was removed.
Middle	The weighting value is equal to the period between when the cluster was first removed, and when it was completely removed.
Max	The weighting value is equal to the period in which the highest amount of the cluster was removed.
Zero	The same weighting is used, as in the Regular method, however the earliest start time of the cluster is set to zero.

The weighting can be thought of as an expected start time for a cluster. When clusters are compared, the one that is expected to start earlier is selected first; in a tie the more profitable cluster is selected.

Once a cluster is selected, based on its weighting, it must be able to be excavated in a period. For this calculation, running totals of the free capacities across all periods are kept. These capacities represent how much more available resources can be used for each of the constraints in the problem. Originally, the capacity is taken from the constraint bound, with no guidance from the linear solution. In Table 2 we describe the regular capacity model, and propose three additional capacity models that consider additional information from the linear solution.

TABLE 2

Methods of calculating capacity.

Capacity type	Calculation method
Regular	The capacity used in TopoSort is the right-hand side of the capacities given in the strategic mine planning problem, as in the original version of the heuristic.
LP guided	The capacity for a given resource is based on how much of that resource is used in the Linear solution. If in a period a constraint is not binding, and a resource is not fully utilised, then the amount of the resource used becomes the new bound.
Extraction constrained	Only extraction constraints are considered, processing constraints have infinite capacity.
LP guided extraction	The capacity used for the extraction constraint is derived from the linear solution. Other constraints have infinite capacity. This is a variation of the LP guided and extraction constrained techniques.

The LP guided capacity uses the solutions from the Linear relaxation to guide TopoSort. This often provides high quality solutions, as the capacity and integer solution are structurally closer to the solutions provided by the LP solution. In the case of the extraction and LP guided extraction capacities, an additional step is required where the Optimised Destinations algorithm (Letelier *et al,* 2020) must be run. This algorithm is explored in work by Letelier *et al* (2020) and can improve the NPV of an existing solution.

We also introduce the concept of a guide capacity. A guide capacity is not a hard constraint on the problem, but rather acts to guide TopoSort to be closer to the solution of the linear relaxation. In the case of the LP derived capacity methods, the capacity can be added as a guide instead. In this case the calculations on when to use a cluster, and how much of it to be used are done comparing to the guide, but the final checks and constraints are based on the constraint bounds of the problem. This allows for deviations from the linear relaxation, while still allowing the linear solution to influence the TopoSort solution.

TopoSort is also modified via small changes to steps of the algorithm, based on our observations. The first is a method referred to as ‘Allow-Waste’. The underlying principle is that it is usually not profitable to remove waste in a period, if it does not allow removal of later ore in the same period. In this case, it is usually better to stop mining in the period, once the process constraints are at capacity. This optional method can provide better results in TopoSort, particularly in problems where ore is not always accessible throughout the time periods. We have also implemented a technique to reconsider the destinations for a bin, when assigning the cluster to be removed. Typically, in TopoSort, a bin is sent to its most profitable destination, however if that destination has no remaining capacity the cluster cannot be removed. In the case that a cluster cannot be removed in a period in its entirety; due to a binding constraint, we can change the destinations of the bins that are contributing to the binding constraint, allowing the cluster to be removed and the bins sent to alternate destinations. This alteration prioritises removing a cluster early, over sending it to a best destination. It is most practical when stockpiling is presented, as a bin may be optionally sent to a stockpile rather than a mill.

Finally, we present a new way of prioritising clusters, based on profit. We start from the bottom of the pit and repeatedly remove clusters in order of their minimum profit. A cluster can be chosen providing its removal is not dependant on the removal of a cluster below it. Our rational is that we want waste to be removed later in the schedule and ore earlier on. We use the resultant sequence as an alternative to the profit weight and term this ‘Ore First’ profit weighting.

Across our options and variants, many different combinations exist. Through testing, the most commonly dominant variants have been selected for use in GMX; with a total of 14 used. For each case, we also run the Optimised Destinations (Letelier *et al,* 2020) heuristic. The 14 variants that are

used are described in Table 3, they have been selected based on testing on our internal data sets, and the Minelib problem set (Espinoza *et al,* 2012).

TABLE 3

The 14 selected TopoSort variants.

TopoSort weighing	TopoSort capacity	Ore first	Allow waste	Rearrange destinations prior to capacity check
Regular	Regular	FALSE	FALSE	FALSE
Regular	Regular	TRUE	FALSE	FALSE
Regular	Regular	TRUE	FALSE	TRUE
Max	Regular	TRUE	FALSE	FALSE
Max	Regular	FALSE	FALSE	FALSE
Regular	Extraction	TRUE	FALSE	FALSE
Regular	LP Guided Extraction	TRUE	FALSE	FALSE
First	Regular	TRUE	TRUE	FALSE
First	LP Guided	FALSE	TRUE	FALSE
Zero	Regular	TRUE	FALSE	FALSE
Zero	Regular	TRUE	TRUE	FALSE
Max	LP Guided	FALSE	TRUE	FALSE
Max	Extraction	FALSE	FALSE	FALSE
Zero	Extraction	TRUE	FALSE	FALSE

Our tests are conducted on the Minelib (Espinoza *et al,* 2012) problems, with results presented in Table 4.

TABLE 4

Results of running 14 variants of TopoSort with optimised destinations.

Data	TopoSort solution	Varied TopoSort best solution	Optimality gap	Total heuristic runtime (seconds)
kd	409968823.6	409995546.4	0.22%	8.4
marvin	902165755.3	902165755.3	1.02%	27.9
mclaughlin	1511525329	1511720682	0.08%	1537.3
mclaughlin_limit	1323164216	1323164216	0.13%	201.7
newman1	23836695.42	23836969.35	1.94%	0.2
sm^2	1648414789	1650225843	0.13%	58.4
zuck_large	57735891.58	57739698.53	0.34%	220.0
zuck_medium	722559755	722733896.4	3.40%	49.6
zuck_small	894921144.5	894921144.5	1.17%	11.7

The Regular TopoSort of (Letelier *et al,* 2020) produces the best result for Marvin and Zuck_Small. For other problems, the use of a set of variants produces better results. Note that most of the runtime is taken up by the Optimised-Destinations time which requires the solution of a second LP.

Similarly, our methods are beneficial for bench-based scheduling problems. For internal client problems that include stockpiles, where the number of clusters is less than 500, our runtimes increase by less than a minute whereas the NPV can improve by several per cent.

1-DEST heuristic

In the prior section we detail our investigations and variation of the TopoSort heuristic which is applicable to capacity constrained problems. It can also be modified to solve stockpiling problems (Moreno *et al,* 2017).

When more complex constraints are present, where constraint coefficients are not restricted to positive values, TopoSort is not suitable. For capacity constrained problems, the removal of a cluster in an earlier period does not affect the feasibility of removing another cluster in a later period. When complex constraints are present, such as blending, the removal of a cluster can affect the feasibility of removing another cluster in a later period. For example, if too much high-grade ore is removed early on then there may not be enough in later periods to meet the blend. For this and other reasons, more sophisticated methods are necessary to find feasible and optimal solutions to these problems.

Recall that a Strategic Mine Planning MILP formulation can have many millions of variables and constraints. Typically, these problems cannot be solved directly using commercial packages such as Gurobi or CPLEX.

To tackle problems, an approach is to solve the problem using sliding time windows. Here the problem is formulated and solved over successive time period ranges. In each case, the subproblem is smaller than the original full problem. The process removes blocks from consideration that are in each successive formulation solution. Often problems are solved on a period-by-period basis. An obvious shortfall is that the solution approach is myopic, in that decisions made in earlier periods are made without consideration for what is to come. In Caccetta and Hill (2003) it was shown that for a gold mine, a 13 per cent increase in NPV profits came from considering the problem holistically rather than one period after another. Yet despite this possible large reduction in NPV, problems often need to be reduced in size in order to be solved in reasonable time by a commercial package.

Another problem reduction technique is the 1-Dest heuristic, from the paper (Letelier *et al,* 2020). Here the solution of the initial LP is used to fix the destinations each bin can be sent to for all time periods. Proportional amounts sent to each destination, per bin, come from a formulation that uses the LP solution. For this discussion, it is sufficient to know that destinations and proportions sent to each destination is fixed across periods.

Note that the MILP formulation requires a set of variables to say, for each bin and period, part of the bin is sent to a particular destination. Also, because bins are associated with clusters, binary variables are needed to say that a cluster is completely removed by a particular time period, which is used in determining whether below clusters can be accessed, or not. Of course, not all clusters are accessible across all time periods. When considering the maximum size of a formulation, we take this into account and so talk of the maximum number of variables in the formulation instead of the true size which is problem specific.

From the above discussion, the number of variables in a MILP formulation is of the order:

$$|B| \times |D| \times |T| + |C| \times |T|$$

Using the 1-Dest heuristic approach, variables for the destination dimension collapse to a single variable per bin. Further, since each bin of the cluster is being removed at the same rate, these variables collapse to a single variable per cluster. Thus, the number of variables in a 1-Dest MILP formulation is of the order:

$$2 \times |C| \times |T|$$

In recent published literature (Letelier *et al,* 2020; Akhil Ankem *et al,* 2026; Abalos, Brickey and Goycoolea, 2025), |C| is typically less that 500 and |T| is less than 30 meaning the MILP being solved

has at most 30 000 variables of which at most 15 000 variables are binary. This is significantly less than the full formulation where the number of bins is a significant contributing factor. For example, in Abalos, Brickey and Goycoolea (2025) a problem having 273 clusters, 288 073 bins, two destinations and 25 period is cited. After preprocessing, the full formulation has 7 208 650 variables. For the 1-Dest formulation, this drops to at most 13 650 variables of which 6825 are binary.

As can be appreciated from the above examples, the 1-Dest heuristic significantly reduces the number of variables in the MILP when creating a solution. This is necessary in order to make problems small enough to be solved using commercial packages, in a reasonable amount of time.

We note that while the problem formulation is greatly simplified, it has also been greatly restricted. In our testing, we have found cases, in particular problems that have blending requirements, where fixing destinations leads to no solution being found. For example, this can occur when, in the LP solution, a high-grade bin can be sent to the mill over a number of periods. The LP solution can be fractional allowing this to occur and the bin proportion is needed to make grade. In the reduced MILP formulation, this may not occur leading to infeasibility. This issue becomes more prevalent when multiple conflicting blending constraints are present.

To see this, consider the example in Figure 6 and say the five clusters are part of a mining operation. All clusters have positive cost, except cluster D which is slightly negative. The numbers beside each label are the clusters average grade. Say we have a minimum average grade of three per period, and two blocks can be removed per period. Clusters D and E are high-grade and needed to make grade in a particular period. In the LP solution cluster E is accessible, leading to cluster D not being used as it has negative cost. But once destinations based on the LP solution are fixed and integer restrictions are placed on the top three clusters (A, B, C), the minimum grade cannot be met in the first period and no solution is found. A better result comes from using cluster D to make grade in the early periods and using cluster E in later periods. This solution is found if the problem is solved holistically, but not when destinations are pre-fixed.

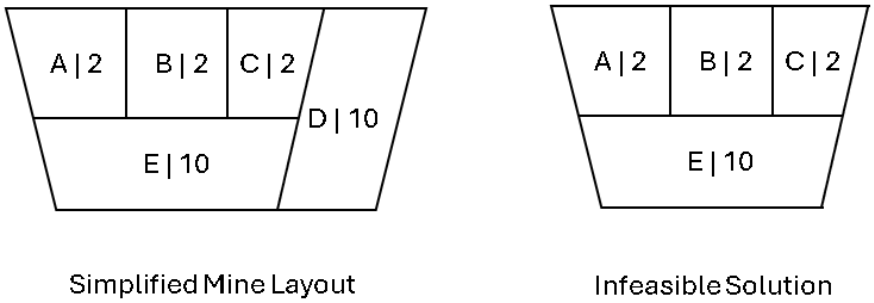


FIG 6 – Example of a simplified mine.

It is important to keep in mind that the method to do this allocation is rules based rather than based on a theoretical result. As such, there is no known way that guarantees the allocation is always correct. Therefore, we believe that for tightly restricted problems, or problems with many different binding constraints, an incorrect allocation may lead to a poor solution or no solution being found.

Considering this issue, we ran a number of experiments. Our goal is to make the MILP formulation less restrictive, to improve on or make feasible the formulation, without significantly increasing the number of variables and so still allowing solution by a commercial package in a reasonable amount of time.

We present three different ways to do this reduction, in the same vein as the 1-Dest heuristic. Firstly, when the bin contains ore, we calculate the 1-Dest bin destination, but also allow the bin to be wholly sent to a stockpile. Our rationale is that we want to use ore at some stage rather than discarding it and so we always want stockpiling to be an option. Similarly, we tried allowing clusters to always be able to be sent to waste, as well as to their 1-Dest calculated destinations.

Lastly, we tried Time Windowing, where we use the LP solution to restrict a cluster so that it cannot be removed before it was started in the linear relaxation, and must be finished by the period it was finished in the linear relaxation. For this method, the problem is restricted in terms of which clusters

and bins are available in any given time period, but bin destinations are not pre-fixed. In our runs, this method gave the best results in terms of discounted cash flow; however, runtimes are much longer as the restricted MILP formulation is much larger when there is a large number of bins. Of course, for a very large number of bins, runtimes will become prohibitive.

In our tests, summarised in Table 5, we compare our implementation of the three different variations. Ten problems were run against each method, results for each problem are given in a row of Table 5. Each case is a bench-based scheduling problem, all of which have different blending and capacity constraints. Problems have a number of phases scheduled by bench. Problem sizes range from 50–300 clusters with up to 30 time periods.

TABLE 5

Results of MILP formulations.

Stockpile-dest	Stockpile-dest runtime (seconds)	Waste-dest	Waste-dest runtime (seconds)	Time window	Time window runtime (seconds)
11.07%	2.7	7.53%	98.6	2.52%	11.44
16.96%	9.1	0.46%	39.4	2.53%	6.8
19.61%	9.1	0.46%	41.7	2.53%	7.04
17.32%	4.9	5.39%	205.8	5.38%	13.17
4.20%	202.3	2.75%	2555.7	2.75%	218.86
5.36%	1296.5	3.79%	1367.2	6.63%	403.7
-	2.6	6.21%	2201.4	6.27%	16.17
0.49%	22.9	3.01%	121.7	3.85%	24.61
5.75%	274.2	3.44%	12374.3	4.15%	786.65
4.42%	38.0	3.36%	46.8	4.51%	35.94

In our tests, the Time Windowing approach always gave a feasible solution, though runtimes were much longer. In the case of the first two variants, the restricted mine planning problem can still be infeasible, and thus unable to provide a solution.

We highlight that for this work, we are not currently adding all known cutting planes. For bench-based scheduling problems, we add Diamond cuts to capture extraction capacity restrictions between clusters and for use in expressing a maximum sinking rate. We also do Probing to restrict the earliest time each cluster can be extracted. Adding other cuts from the literature will undoubtedly improve the LP solution and so the prediction of which destinations each bin is sent to. However, tightening the LP solution is unlikely to guarantee that the destinations chosen exactly match a feasible or optimal solution in all cases.

The main drawback with any of the restricted models, is the lack of flexibility fixing of variables introduces. There is no guarantee that the LP solution will act as a good guide for these fixing decisions. Fixing variables may lead to the problem becoming infeasible or only poor-quality solutions can be found. This limitation is more impactful as problems become more complex. As flexibility is added to try to address this issue, the runtimes increase.

Our investigation leads us to the belief that while restricted formulation methods are useful in many cases, they are heuristic solution methods and so cannot always guarantee that a good solution will be found.

CONCLUSION

This paper presents new modelling for capturing mining direction, along with our analysis and variation of current literature methods for solution generation. We capture mining direction by representing the sequential extraction of bins within a cluster. This addresses a limitation of the 'Same Rate' case which assumes that all bins in a cluster are extracted at the same rate, and does not accurately capture practical mining direction in the schedule. We discussed the implementation of this approach within the BZ algorithm, and how path contraction can be applied to reduce the number of vertices in the flow problem.

We also investigated a range of modifications to the TopoSort heuristic. Our results indicated that these variants can improve the solution quality, achieving better Net Present Value with only marginal increase to runtime. Given their low computational cost and ease of implementation, these variants represent a practical enhancement for generating quality feasible solutions for capacity constrained strategic mine planning problems.

In addition, we examined the 1-Dest heuristic, exploring the impact of pre-fixing decisions of the schedule. We show the results of experimentation that reduces the number of variables in other ways, introducing flexibility at the cost of computational complexity.

Overall, the work shows a methodology for practical mine scheduling. Future work will investigate holistic solution methods as an alternative to heuristic techniques.

REFERENCES

Abalos, P, Brickey, A and Goycoolea, M, 2025. Scheduling by pushbacks: A historical review of optimization approaches for strategic open pit mine planning, *Resources Policy*, 111. http://dx.doi.org/10.2139/ssrn.5138407

Akhil Ankem, V A, Desaulniers, G, Gamache, M and Raymond, V, 2026. New cutting planes for open-pit mine scheduling with multi-period block extraction, *Computers and Operations Research*, 189:107408. https://doi.org/10.1016/j.cor.2026.107408

Bienstock, D and Zuckerberg, M, 2010. Solving **LP** relaxations of large-scale precedence constrained problems, in *Integer Programming and Combinatorial Optimization*, Lecture Notes in Computer Science (Springer: Berlin). https://doi.org/10.1007/978-3-642-13036-6_1

Caccetta, L and Hill, S P, 2003. An application of branch and cut to open pit mine scheduling, *Journal of Global Optimization*, 27(2–3):349–365.

Chicoisne, R, Espinoza, D, Goycoolea, M, Moreno, E and Rubio, E, 2012. A new algorithm for the open-pit mine production scheduling problem, *Operations Research*, 60(3):517–528.

Espinoza, D, Goycoolea, M, Moreno, E and Newman, A M, 2012. Minelib: A library of open pit mining problems, *Annals of Operations Research*, 206(1):93–114.

Hill, S, Bai, V, Rudge, I, Place, D and Poblete, C, 2023. *Strategic Mine Planning, The **BZ** Algorithm and Beyond*. Available from: <https://www.3ds.com/assets/invest/2024-03/geovia-strategic-mine-planning-bz-algorithm-beyond-article-en.pdf>

Hochbaum, D S, 2008. The pseudoflow algorithm: A new algorithm for the maximum-flow problem, *Operations Research*, 56(4):992–1009.

Letelier, O, Espinoza, D, Goycoolea, M, Moreno, E and Muñoz, G, 2020. Production scheduling for strategic open pit mine planning: A mixed-integer programming approach, *Operations Research*, 68(3):769–791.

Moreno, E, Rezakhah, M, Newman, A M and Ferreira, F, 2017. Linear models for stockpiling in open-pit mine production scheduling problems, *European Journal of Operational Research*, 260(1):212–221.

Rezakhah, M, 2017. Open pit mine planning with stockpiling, Operations Research with Engineering, Colorado School of Mines.

Strategic planning in practice – people, projects and decision support

The role and application of techno-economics theory in the development of mining projects

M J Fox-Martin[1] and D J de Villiers[2]

1. Executive – Technical Delivery, Mventech, South Africa.
   Email: michael.fox-martin@mventech.com
2. Executive – Innovation, Mventech, South Africa. Email: jaco.devilliers@mventech.com

ABSTRACT

Mining project development is both an art and a science, demanding nuance and insight underpinned by rigorous technical work. A typical mining project requires a multitude of interdependent decisions, taken across disciplines under imperfect information and in anticipation of an expected economic return. These decisions are frequently made in an expedited, pressurised environment, producing sub-optimal outcomes that undermine the investment and strategic value of the asset. Current approaches remain fragmented across disciplines and project phases, generating an opportunity cost that is rarely understood or quantified.

This paper presents a framework for applying techno-economic theory to mining project development, addressing the content and processes to achieve alignment of the key project decisions with the desired operational outcomes. It argues that success depends on getting both the strategy embedded in a design and its execution right at the same time, and that economic criteria should direct the technical workstreams within each study phase rather than merely evaluate them at the gate.

Given the demands for mining companies to deliver returns in increasingly difficult orebodies, the industry needs a rational, unified framework to guide teams toward the best outcomes for their stakeholders. The framework offered here provides a coherent, economically grounded logic for integrating the fragmented technical and economic threads of project development into a single, repeatable decision-making approach.

INTRODUCTION

Major capital projects frequently do not deliver on their business cases. Research conducted by Flyvbjerg and Gardner (2023) on 16 000 projects across 136 countries revealed a dismal track record of project success, with only 47.9 per cent of projects remaining within budget, 8.5 per cent within both time and budget, and just 0.5 per cent within time and budget while also delivering their intended benefits. Mining projects are even worse; a McKinsey assessment of 35 large mining projects completed between 2002 and 2015 showed that 30 of them went over budget (86 per cent) and 27 were significantly delayed (77 per cent) (Dussud *et al,* 2018).

In mining, value leakage is not confined to schedule and budget, sub-optimal technical decisions can destroy significant value (see for example Lane (1964), Bascetin and Nieto (2007), Poniewierski and Hall (2016), and Asad, Qureshi and Jang (2016)). The consequence of incorrect design criteria may be significant and, in the absence of deliberate measurement, invisible.

While much scholarly work has been done on the technical aspects of mining and mining project development, and to a lesser extent on the integrated technical and economic aspects, widespread adoption of this thinking has not found its way into useable models (Minnitt, 2007).

While project management is generally considered complex, project management in a mining context is even more so:

- Every orebody is different.
- Mining projects involve a multitude of stakeholders from various technical disciplines.
- Decisions are made with imperfect information.
- The economic assumptions vary significantly within the lifespan of a project/operation.

- The strategic intent of the owner has a profound impact on the requirements and approach for the project development.

Against this background, ‘what good looks like’ is subjective and generally poorly understood from a global optimum perspective. Many project teams adopt a generic, forward-pass, technically based project development process without adequate upfront consideration of the technical and economic considerations and framing necessary for operational and investment success.

Project failures in time, budget, or benefits can be attributed to failures in the analysis and design, and/or execution, or any combination thereof. These possible outcomes are synthesised in Table 1. In the context of this paper, project development refers to the life cycle commencing with discovery through the project phases (namely Opportunity, Concept, Pre-Feasibility (PFS), Feasibility (FS), Front-End Engineering Design (FEED), Execution, and Ramp-Up).

TABLE 1

Project development process and content quality.

Project development process quality	**TYPE I ERROR (FALSE POSITIVE)** **Poor strategy, good execution** Executing a flawed strategy with high efficiency. Mine extraction strategy is poorly grounded in technical and economic fundamentals. The project runs smoothly on a shaky hypothesis. Stakeholders believe the plan is robust because execution looks pristine.	**SUCCESS** **Good strategy, good execution** Strategy and delivery are both sound. The optimal mine plan is delivered on-time, on-budget, at the right quality, and delivers sustainable returns over the long run.
	FAILURE **Poor strategy, poor execution** Both strategy and delivery are defective/ suboptimal; any apparent gains are likely noise or misattributed. Apparent early gains disappear under rigorous validation. Difficulties in explaining divergence between plan and reality. Total project collapse.	**TYPE II ERROR (FALSE NEGATIVE)** **Good strategy, poor execution** Sound strategy undermined by poor delivery. The content is strong, but rollout/ execution/resulting operational phase fails to fully realise the potential economic benefits.
	Project development content quality	

The key implications for mining project development are as follows:

- Type I Error is the most dangerous quadrant as this will encourage capital investment in a technically sound-looking, but fundamentally flawed mining plan. Typical examples may include:
  - Over-reliance on inferred resources without adequate infill drilling.
  - Processing plant designed for incorrect metallurgical assumptions.
  - Executing an expansion at peak commodity cycle based on faulty economics.
- Type II Error is the hidden cost quadrant in that a viable orebody or project underperforms, or is ultimately abandoned, because of poor execution. The economic cost is not visible.

The economic value leakage (and remaining economic value) of the above scenarios is depicted schematically in Figure 1. The basic premise is that good execution can never mitigate the value leakage arising from a poorly defined project development strategy.

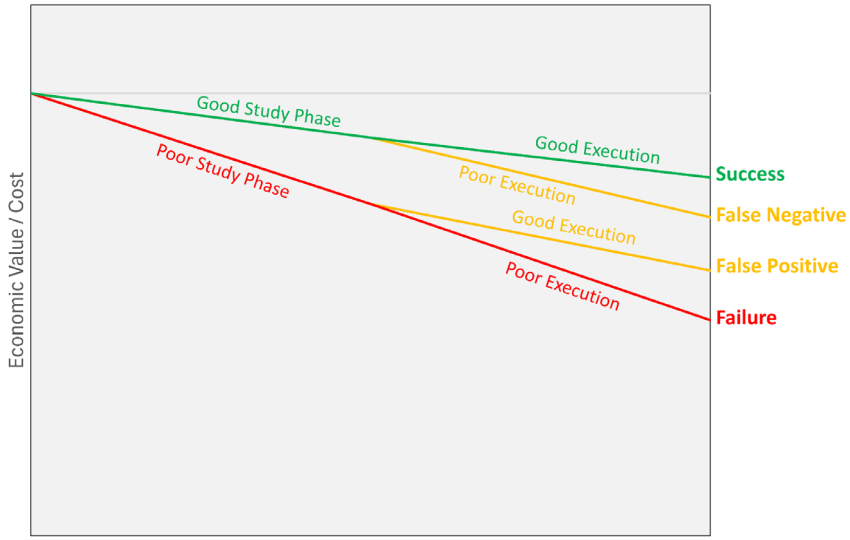


FIG 1 – Synthesised outcomes of project development process and content quality.

While the above concept separates the process and content of project development, these two dimensions have influence on each other, a theme explored further in this paper.

ECONOMIC FRAMEWORK

Traditional economics assumes human beings are perfectly rational and make optimal choices on the assumption that they possess perfect information, unlimited processing power, and enough time. This approach is often adopted in mining projects because it is easy to execute and explain. The problem is that this is not how the world really works, and so we lean on behavioural economics which uses bounded rationality as its foundation. Bounded rationality is the idea that when people make choices, their cognitive ability, time and available information are limited. Because of these limits, humans do not look for the single best choice. Instead, they pick an option that is good enough to solve the problem.

The behavioural economics body of knowledge provides a complementary view dealing with the governance (ie process) and information (ie content) which is useful in understanding the dynamics within mining project development and their impact on the economics thereof.

Table 2 summarises the key differences between traditional and behavioural economics and highlights key considerations as to why decisions in mining projects must necessarily adopt behavioural economic assumptions.

TABLE 2

Traditional economics versus behavioural economics.

	Traditional economics	**Behavioural economics**
Human nature	Perfectly rational and logical	Boundedly rational and emotional
Goal	Optimising (seeking the absolute best outcome)	Satisficing (seeking a good-enough outcome)
Information	Complete, perfect, and costless	Limited, incomplete, and confusing
Errors	Random and quickly corrected by the market	Systematic, predictable, and driven by biases
Computational capacity	Infinite capacity	Limited cognitive bandwidth

Bounded rationality

Bounded rationality is a concept introduced by Simon (1957) as a corrective to the idealised 'rational actor' model in economics.

The core idea is that human decision-making is rational, but only within limits. Classical rationality assumes decision-makers have complete information, unlimited cognitive capacity to process it, and the ability to identify the optimal choice among all alternatives. Bounded rationality rejects this as unrealistic. In practice, people face three fundamental constraints: information is incomplete or costly to obtain, cognitive and computational capacity is limited, and time to decide is finite. Rationality is therefore 'bounded' by these conditions, a situation not uncommon in mining.

A central behavioural consequence identified by Simon (1957) is satisficing, a term he coined by blending 'satisfy' and 'suffice'. Rather than searching exhaustively for the best possible option (optimising), decision-makers set an aspiration level and accept an available alternative that meets it. This is not irrationality; it is a sensible response to the cost of deeper investigation. The distinction that matters here is between bounded rationality and irrationality. Bounded rationality preserves the assumption that people are trying to act sensibly toward their goals; they are simply doing so under real constraints. Satisficing is relevant to the debate regarding net present value (NPV) or all-in sustaining cost (AISC) as defined by World Gold Council (2018), as the primary optimisation criteria discussed later.

Because our minds have limits, we rely on mental shortcuts or rules of thumb that lead to systematic biases:

- Heuristics (mental shortcuts): Rules-of-thumb used to make fast decisions under time pressure. The economic impact is that instead of calculating exact probabilities of risk, people use the availability heuristic, judging risk based only on recent outcomes or personal experiences.
- *Status quo* bias: The tendency to do nothing or stick with a default option because analysing alternatives requires too much mental effort. The economic impact is the potential to default to a sub-economic solution.
- Choice overload: A breakdown in decision-making when a person is presented with too many options. The economic impact is that too many alternatives lead to poor decisions in a potentially rushed approach.
- Framing effects: Decisions are made based on how information is presented, rather than the facts themselves. The economic impact is the potential to promote a sub-economic solution.

While the use of heuristics may result in sub-optimal choices being made, appropriately and responsibly applied they enable effective decision-making.

Bounded rationality is discussed later in terms of the content of an operations strategy (ie the technical aspect of mining project development).

Opportunism

Williamson (1981) built on the work of Simon (1957) by pairing bounded rationality with opportunism, wherein he postulated why contracts are necessarily incomplete: parties simply cannot foresee and specify every future contingency, because their capacity to anticipate and process complexity is limited. Opportunism is defined as self-interested behaviour with guile; meaning parties may pursue their own interests by strategic, dishonest, or deceptive means, especially when contracting and enforcement are imperfect. Opportunism emphasises tactics like lying, cheating, or taking advantage of informational/contractual weaknesses, particularly under uncertainty and hard-to-contract conditions. This incompleteness is what creates the need for governance structures; in other words, mechanisms to handle the gaps and adaptations that arise as circumstances unfold. Williamson proposed a theory of the firm as a governance structure which seeks to economise on bounded rationality while simultaneously safeguarding against opportunism, both of which impact transaction costs. Transaction is defined as the transfer of goods and/or services (an exchange) that takes place across a technologically separable interface between parties (Williamson, 1981). Transactions may

be internal or external to the firm and may not always be in financial terms (e.g. inter-department engagement). In the context of this paper, and mining project development in general, the project management function provides the governance structure, transaction costs include both financial and opportunity costs, and bounded rationality refers to limitations in cognitive ability, time and available information.

Information asymmetry is the enabling condition for opportunism whereby the other party is unable to cheaply verify the correctness or accuracy of information. If information were symmetric and costless to check, guile would be detected and priced out. Williamson's framing separates this along the transaction timeline. Ex ante, asymmetry allows adverse selection in that the informed party misrepresents attributes before the deal is concluded (eg a contractor's true cost base, a vendor's real performance envelope, a developer's genuine confidence in a resource etc). Ex post, it produces moral hazard through hidden action after commitment; where effort, quality, or cost discipline can't be fully monitored. Both are only problems because bounded rationality prevents writing the complete contingent contract that would otherwise close the gap.

Typical examples of information asymmetry in mining project development include:

- Study teams hold superior information, relative to the investment committee, on how much of a Front-End Loading (FEL) deliverable rests on firm data rather than judgement.
- A geological data gap is an asymmetry problem in substance even if it appears purely technical, in that sparse data widens the interpretive latitude a party can exploit, knowingly or not.
- An engineering, procurement and construction management (EPCM) contractor typically holds superior information on engineering quantities, contingency loadings, and schedule realism relative to the owner (ie classic asymmetry).
- Vendors hold superior information on equipment derating and life cycle cost.
- Operational teams hold superior information on operating costs, and may sandbag the estimates informing a project valuation in order to preserve float later.

If we delve into the foundational reasons for cost overruns, delays, and benefit shortfalls, Kahneman (2011) as a psychologist explains these outcomes through cognitive bias (especially optimism bias and the planning fallacy), while Flyvbjerg, Skamris Holm and Buhl (2002), from an economic geography perspective, explain these outcomes in political-economic terms, citing strategic misrepresentation. When viewed in the context of Williamson, Kahneman describes the outcomes in terms of bounded rationality while Flyvbjerg describes the outcomes in terms of opportunism.

Flyvbjerg, Skamris Holm and Buhl (2002) argue that, for the largest overruns, systematic underestimation in large public works is better explained by deception than by delusion. This finding emphasises the importance of process and governance in megaproject development.

The management of opportunism is discussed later in terms of the process of an operations strategy.

OPERATIONAL CAPABILITIES FRAMEWORK

Several authors have expressed concern that while project teams may have adhered to a rigorous stage gate process and prepared an FS that satisfies the FEL criteria, the FS outcomes may deliver a project that is sub-optimal or mis-matched to the desired business outcomes (Hutton, 2022).

The manufacturing industry has undertaken a significant volume of research into the internal capabilities required to achieve competitive advantage. One useful model, the 'Sand Cone' model of Ferdows and de Meyer (1990) presented in Figure 2, has been tested across several industries and contexts and shows promise for application within the mining industry. The central thesis of their Sand Cone model is that sustainable performance is cumulative and sequential, with quality performance forming the foundation, building through dependability, flexibility, throughput, and finally cost efficiency. Their research indicated that these capabilities are reinforcing rather than trade-offs, the implication being that sustainable competitive advantage requires a long-term view and disciplined acquisition of the foundational capabilities. The model implies that, for example, should

the cost capability be pursued without first expanding the preceding foundational capabilities, the Sand Cone will collapse.

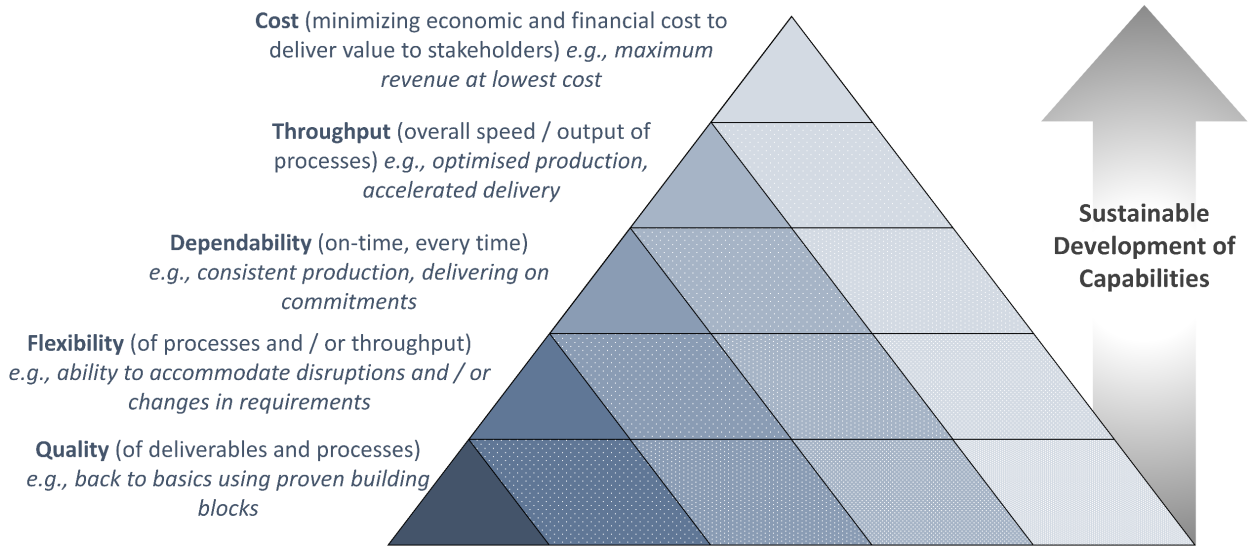


FIG 2 – Cumulative capability sand cone model, adapted for mining (after Ferdows and de Meyer, 1990).

We propose a broader interpretation of 'Cost', encompassing both financial and economic costs, in order to capture the essential requirement to deliver optimal value by also enhancing the revenue value drivers in the development of mining projects. As such, the cost capability incorporates the revenue imperative as an opportunity cost where revenue generation is sub-optimal.

The research has indicated that the order of capabilities does not always follow the sequence as proposed by Ferdows and de Meyer (1990). Schroeder, Shah and Peng (2011), testing the sequence across 189 manufacturing plants, found no universal support for the original ordering: some plants follow it and others do not, with the sequence contingent on strategic choices and plant context. This finding is intuitive as the competitive capabilities must necessarily reference the key uncertainties, economic assumptions, and technical considerations. A concentrator, for example, may place greater emphasis on dependability rather than flexibility, while in a mining operation flexibility can ensure dependability. The project team, supported by subject matter experts, should determine their respective capability criteria to inform their workstreams. The project or study manager is responsible for integrating these criteria, noting the dependencies in the technical design process described in the Project Development Framework section.

We propose that mining is one such context, and we therefore place flexibility beneath dependability. The reason is the locus of variance. In manufacturing the dominant source of variability is the process itself, so conformance to design yields dependability directly. In mining the dominant source is the orebody, which lies upstream of every process the operation controls; an operation may conform fully to its design and still fail to deliver tonnes at grade. Where variance cannot be eliminated at source it must be absorbed, and absorption capacity, in the form of stockpiles, blending, multiple working faces, alternative processing destinations and cut-off latitude, is flexibility. Slack and Lewis (2015) state the same relationship in general terms, observing that flexibility internally supports dependability.

The Sand Cone model has been successfully used as a framework for trade-off and operability assessments.

It is further proposed that the stage-gate criteria developed by the mining houses and their project development consultants and contractors be categorised in accordance with Sand Cone thinking.

OPERATIONS STRATEGY FRAMEWORK

A key theme of this paper encompasses the reconciliation of competitive requirements with the operational resources or capabilities. The development of operations capabilities is strategic in the

sense that firms that become proficient in developing the right capabilities secure competitive advantage. Slack and Lewis (2015) provide a useful definition as follows:

> *Operations Strategy is the total pattern of decisions that shape the long-term capabilities of any type of operation and their contribution to overall strategy, through the reconciliation of market requirements with operational resources.*

The key endeavour, therefore, is to ensure the correct project decisions are taken to provide the optimal capabilities in the future operations for sustainable competitive advantage.

Slack and Lewis (2015) propose an operations strategy framework, presented in Figure 3, which integrates the process and content components of an operations strategy.

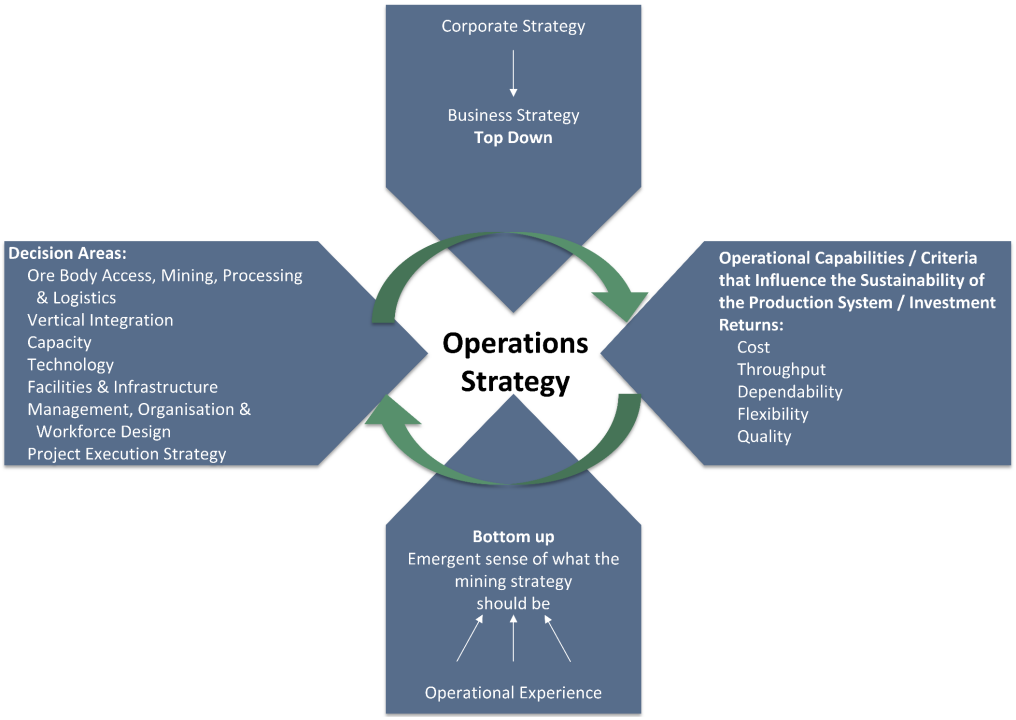


FIG 3 – Operations strategy process and content (adapted from Slack and Lewis, 2015).

Operations strategy content

On the horizontal axis, project teams make numerous technical decisions which have an influence on the competitive capabilities (eg capital expenditure (CAPEX), operating expenditure (OPEX) etc). In the traditional approach the design team produce outcomes that are generally very narrow in their focus, usually discipline specific, and often tenuously linked to the business requirement. In our proposed approach we start with the end in mind, in other words by being clear on the techno-economic criteria and making decisions in line with the requirements for the optimisation of the project. The approach is supported by Hutton (2022) who advocates a techno-economic process to inform the strategy content.

Linking the key decisions as part of the operations strategy (project development in our case) is a complex endeavour; the Sand Cone model is used to simplify this process and enable discrete, fit-for-purpose criteria to be evaluated by the multitude of stakeholders responsible for the many workstreams within the project. In this way the project team can explore and optimise the desired/preferred configuration before committing significant resources and time into the subsequent phases.

The content of the operations strategy, shaped by the multitude of decisions taken therein, is where bounded rationality effects are felt. The use of artificial intelligence for the processing and interpretation of data can significantly mitigate the effects of bounded rationality.

Operations strategy process

On the vertical axis, there is a tension between the top-down goal and the bottom-up sense of what a project can deliver. Through the iterative process, the project team is able to test scenarios, criteria, and assumptions and examine their impact on the project outcomes. Equally important, the process builds alignment between the various stakeholders.

This iterative process should be conducted for all key decisions in each project phase, and indeed a similar approach should be adopted in the operations phase.

The process adopted in the development of the operations strategy is where opportunism takes hold, requiring an effective governance framework to identify poor content (as a consequence of bounded rationality) whilst safeguarding against opportunism.

Artificial intelligence, through reduction of data asymmetry, will play an important role in reducing opportunism and thereby enhancing the governance aspects. Decisions, however, must necessarily be taken by humans, informed by appropriate data.

DECISION FRAMEWORK

Decision quality

Spetzler, Winter and Meyer (2016) set out a Decision Quality (DQ) framework defining the elements necessary for a good decision as presented in Figure 4. It is an extension of decision science and in particular Decision Analysis; a set of concepts and tools that produce clarity about the best choice in an uncertain and dynamic environment. DQ not only uses Decision Analysis to get to the 'right' answer but also engages the relevant parties in the decision process to achieve alignment and commitment to action.

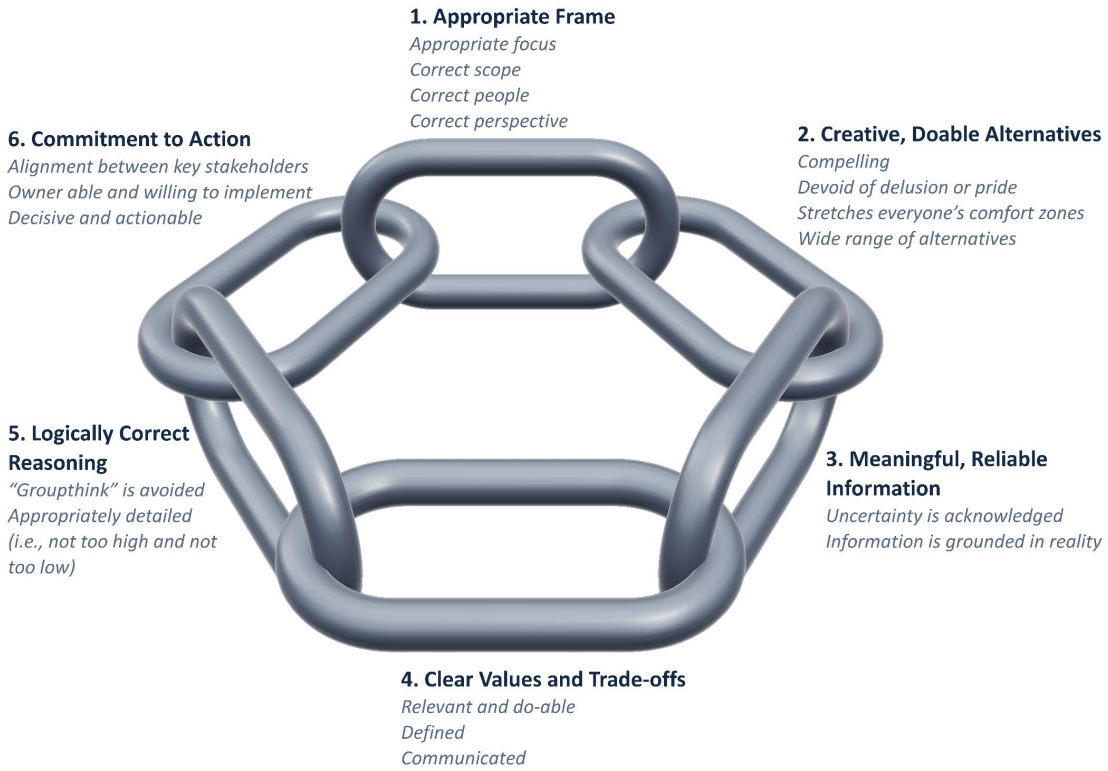


FIG 4 – Decision quality framework (adapted from Spetzler, Winter and Meyer, 2016).

Every decision is dissected into six distinct elements, each of which must be addressed appropriately for a good quality decision.

1. Appropriate frame: A decision frame has three components: Purpose (what you hope to accomplish), Scope (what to include and exclude), and Perspective (your point of view, and

consideration of how others might approach it). Quality in framing is achieved when decision-makers have alignment on purpose, perspective, and scope.

2. Creative alternatives: Alternatives are creative, implementable, have sufficient breadth and detail.
3. Relevant, reliable information: Information may come from data and judgment. The issue with decision-making is that it is forward-looking, and all we have is the data at hand. The data must be reliable and should describe the underlying uncertainties and biases.
4. Clear values and trade-offs: This element focuses on preferences, objectives, and risk tolerance.
5. Sound reasoning: Alternatives, information, and values form the basis of the decision (ie what you can do, what you know, and what you want). Sound reasoning integrates these to help select the alternative with the best prospects.
6. Commitment to action: Without effective action, the time and effort that went into the decision is wasted. In addition to the above five elements, for a quality decision you also need to gain the motivation and commitment to action from the mandated individuals.

Like the links in a chain, the quality of the decision is dependent on the strength of each element, and the decision is only as good as the weakest link.

A central insight of the framework is separating decision quality from outcome quality. Decision quality is a reliable, repeatable way to achieve excellence in decision-making; however, a good decision can still have a bad outcome due to uncertainty.

The thinking in the Operations Strategy Framework, together with the Operational Capabilities Framework, aligns with the DQ framework.

Multi-criteria decision analysis (MCDA)

Belton and Stewart (2002) devised MCDA as a formal approach for the rational and structured analysis of multiple potentially conflicting criteria and provides a structured model that can serve as a focus for discussion and leads to rational, justifiable, and explainable decisions.

The MCDA model allows for mixed sets of data, quantitative and qualitative, including expert opinion, which allows for the participation of all key stakeholders and the outcomes of which generally receive acceptance from all the parties concerned.

We advocate the ‘value measurement model’ proposed by Belton and Stewart (2002), where numerical scores are constructed to represent the degree to which one decision option may be preferred to another. Where actual numerical data is available, it is used directly in the model. To ensure that the approach is understood and receives acceptance, the Sand Cone criteria are reformulated, in Table 3, into terms and language understood by project, technical and operational teams. The Sand Cone logic is, however, retained.

TABLE 3

Multi-criteria decision criteria.

Ferdows and de Meyer (1990) capability	Proposed capability	Criteria	Description
Quality	Safety and Legal	Safety	Will not endanger the health and safety of stakeholders.
		Legal/Statutory Conformance	Conforms to current/anticipated legal obligations.
	Design and Technical	Primary Purpose	Adherence to the primary operating characteristics with reference to the design criteria (ie actual versus design).
		Secondary Purpose	The secondary value adding characteristics (eg water and energy intensity).
		Quality Conformance	Conformance to specification eg mining best cut, material size distribution etc.
		Technology Maturity	The maturity of the technology and near-term pathway for the development of the technology.
		Labour Skills and Availability	The labour requirement for the designed process.
Flexibility	Flexibility	Operational Flexibility	Ability to accommodate uncertainty/unforeseen issues (eg geological features, delay in mining cycle etc).
		Modification Flexibility	Ability to accommodate changes in design, layout, equipment, systems and processes. Links to mechanisation pathway.
		Scalability/ Constraints	Ability to scale production, and constraints on production throughput.
Dependability	Dependability	Variability	Likelihood of delivering the required work, at the right time, at the right quality.
		Operability	Ease of use/operation.
		Maintainability	Ease of servicing/maintenance.
		Constructability	Ease of construction and commissioning.
		Operational Readiness	Extent of operational readiness intervention required.
Throughput	Throughput (Revenue)	Production Throughput Confidence	Confidence in the ability, supported by the technical decisions and noting all constraints, to sustainably deliver the scheduled production throughput.
Cost	Cost	Production Cost Confidence	Confidence in the ability, supported by the technical decisions and noting all constraints, to sustainably deliver scheduled production throughput at the budgeted cost.

The MCDA process is as follows:

- Key stakeholders/representatives are identified ensuring that all relevant disciplines are represented.
- The alternatives, their assumptions and outputs are articulated and explained to the team.
- Agree the weighting of each technical capability (eg Safety and Legal, Design and Technical, Operability, Flexibility and Dependability, Throughput (Revenue), and Cost).
- Agree the identification and weighting of the key criteria within each technical capability.
- For each criterion, identify the measure to be applied, the scale for that measure, and whether the scale is inverse.

- In an integrated session, for each technical capability, rate each alternative relative to a baseline, recording the key findings/reasoning to ensure a common understanding and convergence on the relevance, validity and reliability of the criteria and their associated data.
- The MCDA model normalises the data for each criterion and assigns 0 to the lowest scoring alternative, and the maximum weighted score to the highest scoring alternative with the remaining alternatives receiving a pro-rata score according to their placement between the lowest and highest scoring alternatives.
- The scores are totalled to determine the degree to which each alternative satisfies the selected criteria and allowing for analysis by technical capability.
- Where an alternative scores significantly lower in the foundation layers (ie Safety and Legal, and Design and Technical), this should potentially be flagged as fatally flawed.
- The results are presented to the team and any clarifications or queries addressed as necessary.

The outcome of the above process is a numerical index of the relative preference of each alternative against the agreed preferences and their weighting. This serves as the technical index for the techno-economic assessment.

Economic assessment

The economic assessment is derived from the techno-economic model (if appropriate to the decision), or alternatively from estimates compiled by the project team. Several economic criteria may be used in the economic assessment, depending on the nature of the decision to be made (eg NPV, AISC, margin curve position, life cycle cost).

The selection of criteria is not a modelling convention. It is a framing decision in the sense of the first Decision Quality element, and it determines the direction in which the technical workstreams will be steered.

In practice this decision is seldom made explicitly; it generally follows the corporate protocol which may undermine its usefulness in adjudicating the project for its intended role in the portfolio. Hall (2003) traces the industry's persistent underperformance in part to this disconnect between the measures that attract external attention and those that drive intrinsic value, noting the perverse result that value-adding strategies are often perceived as likely to depress the share price. Where the criterion is never stated, competing criteria are optimised in parallel by different disciplines, and the resulting conflict surfaces at the gate as an unexplained divergence between the technical recommendation and the investment case. This is bounded rationality operating on the content of the operations strategy: the team is acting sensibly within its own frame, but no participant holds the whole frame. We therefore propose that the economic criteria, their relative priority, and the tolerance for trading one against another are agreed and documented at phase initiation, before the technical workstreams commence.

The divergence of NPV and AISC

The most consequential instance of this problem is the relationship between NPV and AISC, because the two metrics are routinely reported side by side and are widely assumed to move together. They do not.

NPV measures value creation and is sensitive to the timing of cash flows. AISC measures unit cost efficiency and is insensitive to timing altogether. Lane (1964) established the consequence formally in the context of cut-off grade: an NPV-maximising cut-off must carry an opportunity cost for the future cash flows foregone by the decision taken now, and is therefore systematically higher than the break-even cut-off, which carries no such term. The penalty is symmetric. Bascetin and Nieto (2007) observe that a cut-off set too low extends the life of the operation without economic justification, while a cut-off set too high consigns economic material to waste; Rendu (2014) develops the relationship between net present value, capacity constraints and opportunity cost in detail. The cut-off grade that minimises unit cost is likewise displaced from break-even, but by a different amount, so the two optima do not in general coincide. Their relative order depends on the operating

margin and on the magnitude of the opportunity cost, and may reverse between deposits and between price environments. The divergence is therefore structural rather than a product of poor estimation, and is shown schematically in Figure 5. An extensive literature extends the result to dynamic price and cost environments, multiple elements and destinations, and conditions of grade uncertainty (Asad, Qureshi and Jang, 2016; Dagdelen and Kawahata, 2008; Githiria and Musingwini, 2019), but the underlying conflict is unchanged.

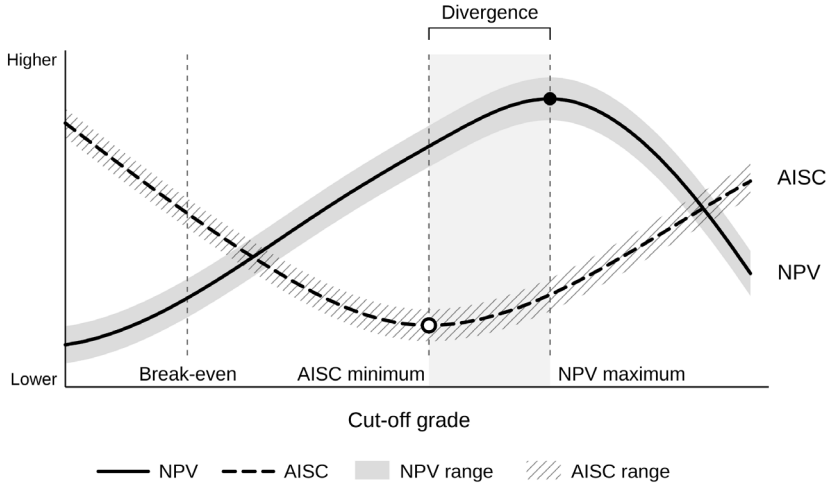


FIG 5 – Schematic divergence of NPV and AISC as a function of cut-off grade, with bands showing the effect of other technical decisions.

The envelopes in Figure 5 represent the range of outcomes arising from the other decisions listed in Table 4. Once that range is admitted, neither optimum holds a fixed position on the cut-off grade axis and the width of the divergence becomes project-specific, which is why the divergence cannot be located by sweeping cut-off grade alone. Whether the two optima are separable at all for a given deposit is itself a question that only the integrated analysis can answer.

TABLE 4

Technical decisions and their divergent economic effects.

Technical decision	Economic lever	Effect on NPV	Effect on AISC
Process and infrastructure design	Installed capability and design margin	Leaner design raises NPV; under-specification erodes it through lost production	Design margin raises the capital charge but stabilises unit cost
Cut-off grade policy	Grade and tonnage delivered per period	Higher early cut-off increases NPV	Reduces early AISC; increases late-life AISC
Throughput selection	Mine life against capital intensity	Higher rate compresses life and increases NPV	Scale economies reduce AISC; feed grade dilution increases it
Stripping schedule	Timing of waste movement	Accelerated stripping increases NPV	Increases near-term AISC
Mining method and access	Development capital and selectivity	Bulk methods favour NPV at scale	Selective methods protect grade and reduce AISC
Sustaining capital timing	Free cash flow profile	Deferral increases NPV	Deferral produces later step-changes in AISC

The significance for project development is that the divergence is created by technical decisions taken early in the design sequence, well before any economic evaluation is performed. Table 4 sets out the principal levers and the direction of their effect.

Each row represents a decision that is presently made on technical grounds within a single discipline, and each carries an economic consequence in two directions at once. A team instructed to reduce operating cost and a team instructed to maximise value will, acting in good faith, design different mines from the same orebody.

Techno-economic trade-off approach for macro decisions

It follows that the techno-economic assessment cannot consist of optimising a single metric. We propose instead that candidate configurations be generated applying the range of the levers in Table 4, and that their outcomes be mapped using NPV, AISC and Technical (ie MCDA) criteria. Figure 6 presents typical outcomes as a bubble graph. The risk-averse case (Alt 0) is technically preferable but delivers the worst NPV. Alt 1 and Alt 4 may be eliminated due to poor technical confidence. Alt 2 and Alt 3 are the preferred candidates. In this way the strategic decision belongs to the owner rather than a technical one belonging to the study team. This follows the scenario-based method advocated by Poniewierski and Hall (2016), and extends the Hill of Value approach of Hall (2014), which maps net present value as a surface over cut-off grade and production rate, by adding a second economic axis.

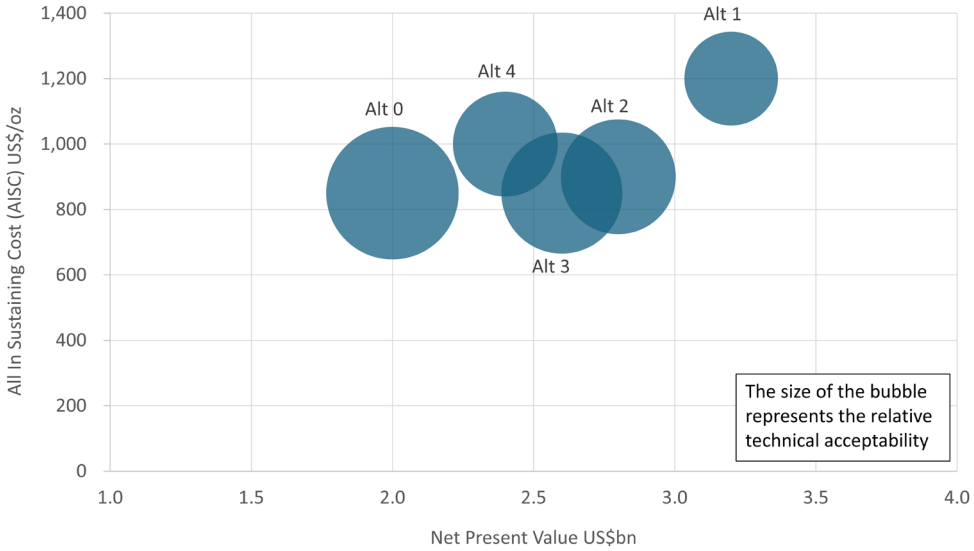


FIG 6 – Techno-economic trade-off of macro decisions.

A project financed substantially by debt requires resilience in the lower half of the cost curve and will accept a lower NPV to obtain it. An owner with a strong balance sheet and a constructive price view will accept a higher cost position to bring value forward. In a weakening price environment the optimal position moves, which is why the criterion must be revisited at each gate rather than fixed at the outset. The phrase ‘within owner risk tolerance’, used in the phase requirements, is given operational meaning here: it is the choice of position within this trade-off space.

This approach is consistent with the multi-objective optimisation methods now appearing in the mine planning literature, in which Pareto-optimal sets are generated for competing objectives rather than a single optimum being sought. We do not claim this as a validated technique. We propose it as an analytical device that renders an existing and largely unmanaged conflict visible at the point where it can still be influenced.

Two findings from the literature bear directly on the argument of this paper. The first is that the direction of the common error is not the one practitioners assume. Poniewierski and Hall (2016) report that across numerous strategy optimisation studies they have not encountered an operation mining above the cash-maximising grade, but have frequently encountered operations mining well below it, subsidising loss-making tonnes with higher-grade ore. The concern that a value-directed

cut-off policy amounts to high-grading is therefore largely misplaced. The second is a matter of sequence. They conclude that cut-off grade should be an outcome of a study conducted to meet the company's value goals, and not a predefined input to that study. That conclusion, reached for a single parameter, is the conclusion this paper reaches for the technical workstreams as a whole.

Techno-economic trade-off for design decisions

Although the technical and economic criteria are linked, it is useful to separate these two dimensions when presenting the outcomes of the MCDA process. Table 5 presents typical technical evaluation outcomes per Technical Area, and Figure 7 presents the integrated technical and economic outcome. In this case the preferred alternative demonstrated superior value but was somewhat less technically acceptable than the reference case. Alternative 3, although fatally flawed, is retained in the table to demonstrate why it is later rejected.

TABLE 5

Technical evaluation of design decisions.

	Alt 0	**Alt 1 Hole into the shaft orepass at 14.5 Level**	**Alt 2 Conveyor belt in additional shaft cross-cut on 16 Level**	**Alt 3 Tramming in additional shaft cross-cut on 16 Level**
	Reference case			
Safety and legal	0.100	0.100	0.100	0.000
Design and technical	0.139	0.139	0.095	0.050
Operability	0.215	0.157	0.276	0.092
Flexibility and dependability	0.182	0.066	0.141	0.075
Throughput	0.061	0.031	0.031	0.000
Cost	0.061	0.031	0.031	0.000
Overall score	**0.759**	**0.523**	**0.673**	**0.217**
Overall ranking	1	3	2	4

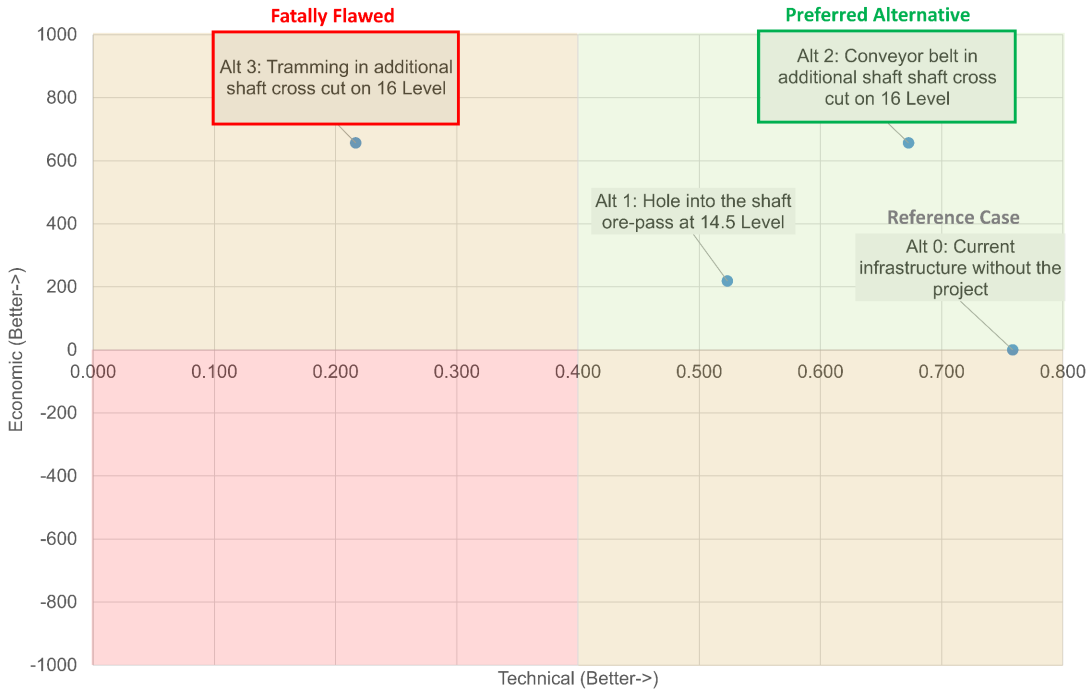


FIG 7 – Techno-economic trade-off of design decisions.

PROJECT DEVELOPMENT FRAMEWORK

Technical development process

Figure 8 presents a generalised schematic process flow of the technical design process within a typical mining study phase. Each element has the potential to impact the business case of the project, either as a direct contributor to revenue, as a constraint to production, a key cost driver, or an accelerator or impediment to timely implementation.

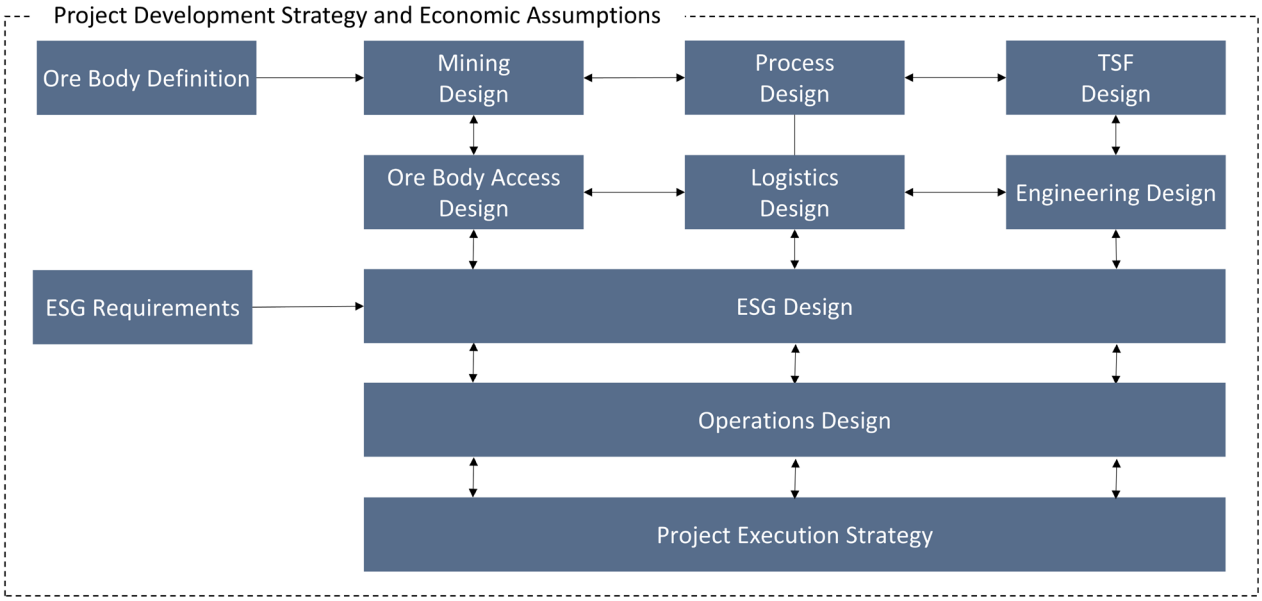


FIG 8 – Generalised schematic process flow of the technical design process.

Within this process there is an implicit hierarchy of decisions, and therefore the maturity of each element necessarily develops sequentially as depicted in Figure 9.

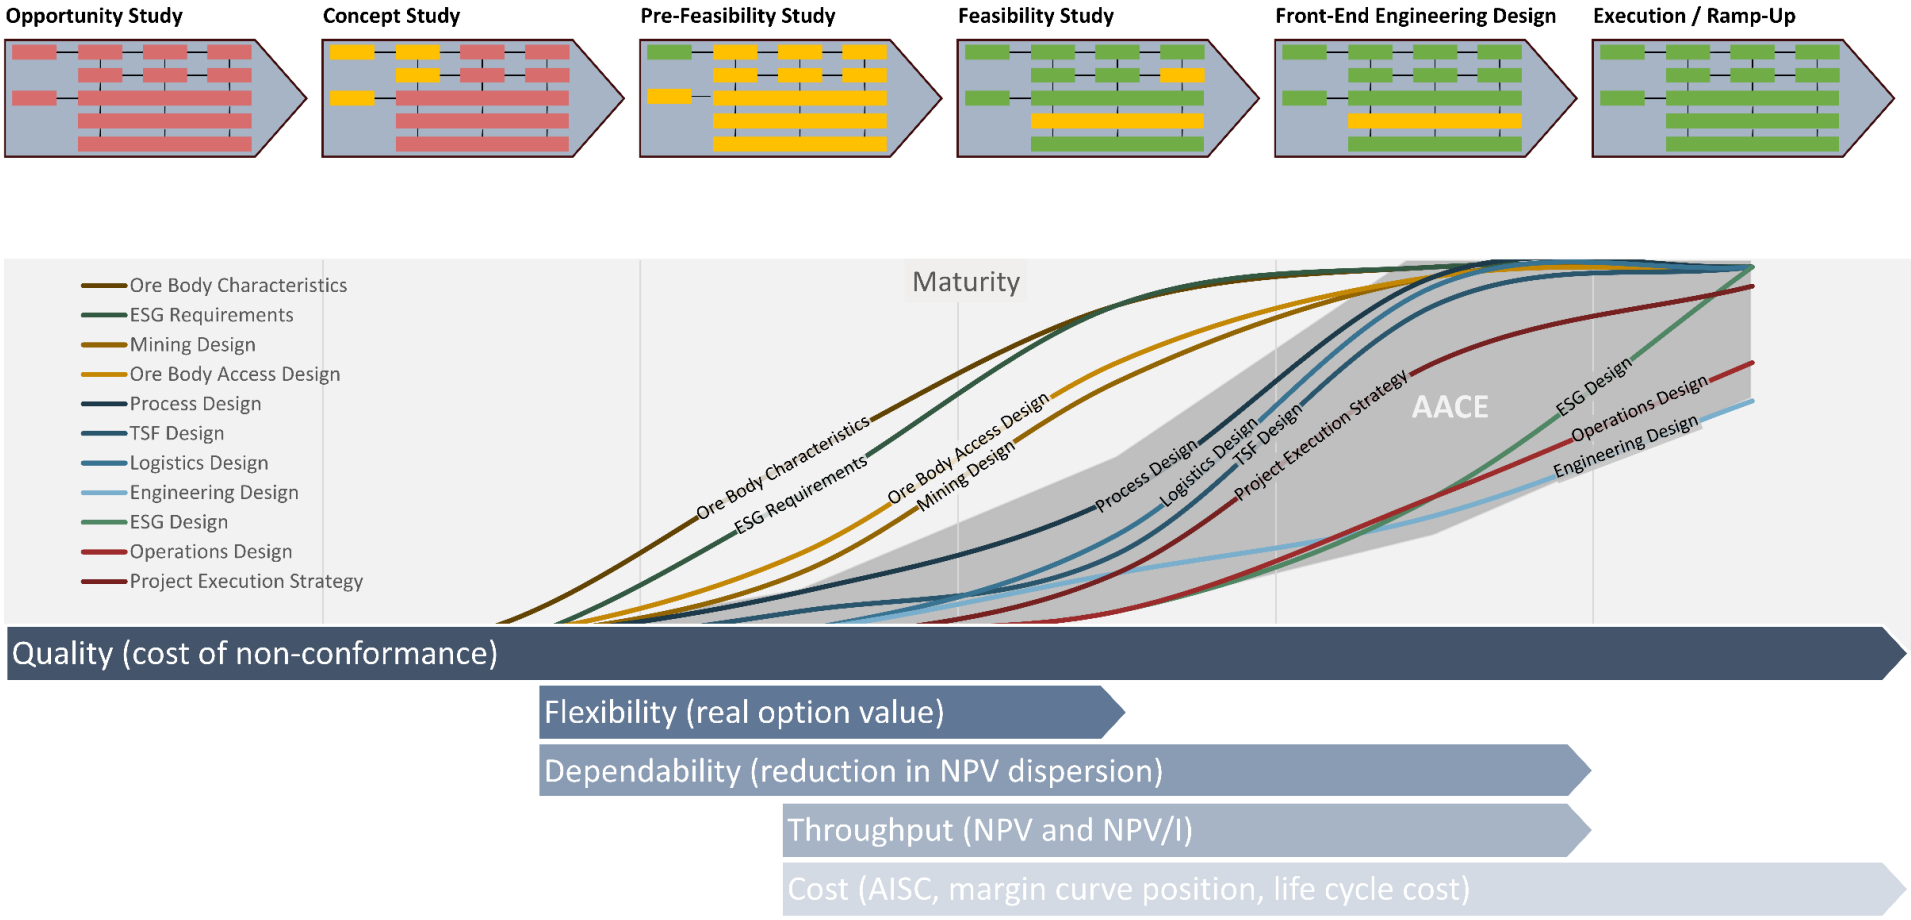


FIG 9 – Synthesised maturity per key technical element and competitive capability.

The implication, in a mining context, is that the early workstreams and disciplines must necessarily be advanced ahead of the AACE International cost estimate class requirements (grey shaded envelope), to ensure sufficient reliability in the data used by the dependent disciplines.

The economic definition maturity is, at best, directly dependent on the maturity of the technical elements within each phase.

This hierarchy accords well with the Sand Cone model:

- Quality is the first priority commencing with the orebody and environmental, social and governance (ESG) requirements and then flowing into the dependent elements. Quality is enhanced through the entire project development process.
- Flexibility is developed in the PFS and FS through appropriate evaluation of alternatives.
- Dependability, which is reliant on flexibility, is developed in PFS through to FEED, culminating in designs that are fit-for-purpose.
- Throughput is finalised with PFS and then designed-in through to FEED.
- Cost, which is directly influenced by all the preceding capabilities, is structurally embedded in PFS and refined through FS, FEED, and Ramp-Up.

Stage-gate process

Stage gating is a project-development approach that divides a project into discrete phases with formal review ‘gates’ between them. At each gate, stakeholders confirm that predefined objectives and deliverables for the preceding phase are complete and decide whether to go ahead, pause, re-work or abandon the project (ie the approving body makes an investment decision based on the strength of the techno-economic outcomes presented to it). This ensures that each stage (with its planned scope, schedule and cost) meets quality and viability criteria before further investment is made.

Stage-gating in mining projects is particularly important because of the high uncertainty and large capital at risk. Initial exploration and scoping involve relatively low cost but high uncertainty, whereas later stages such as engineering and construction require significant investment. By gating the process, companies gather needed data (resource models, market studies, permits etc) at each phase and only commit further funding when justified.

The criteria embodied within the stage-gate framework may be seen as heuristics or ‘rules of thumb’ commonly adopted for project success. To be useful the heuristics should be relevant and applicable to the context and scope of the project in question. Typical stage gate objectives, requirements and outcomes are summarised in Table 6. The stage-gate architecture itself is partly a device for progressively collapsing asymmetry; each phase buys information that constrains the room for opportunistic representation before capital is irreversibly committed. The gating process forces alignment between stakeholders at the defined gates, thereby reducing information asymmetry and the resulting opportunism that may otherwise arise.

Generally, the stage-gate process provides formal, incremental oversight of the project development. Each gate enforces that risks are understood and objectives met before larger commitments are made. By identifying technical and financial issues early, the process helps prevent downstream failures (cost overruns, schedule delays or unviable projects). In short, stage-gating is a high-level governance tool to manage the investment decision under uncertainty.

In practice, the key economic criteria are poorly or seldom used within the various study phases to direct and inform the technical requirements. The economic and financial outcomes are generally determined at the end of each phase as an outcome of the technical workstreams; and where the necessary hurdle rates are not achieved, project teams hastily develop schemes to improve the early net cash flow to satisfy the investment criteria rather than deliver what the project really requires. This invariably leads to project failure or Type I Error.

TABLE 6

Requirements per project development phase.

	Opportunity	Concept	Pre-feasibility (PFS)	Feasibility (FS)	FEED (Front-end-engineering-design)
Accuracy	±50%+	±35–50%	±20–25%	±10–15%	±5–10%
Purpose	Identify and prioritise opportunities: Opportunity assessment. Screen concepts. Strategic prioritisation. Identify data gaps. Initial framing of key value drivers.	Early-stage study to define, profile and screen alternatives: Test basic technical and economic viability. Define value optimisation trade-offs. Strategic and techno-economic profiling for optimum value within owner risk tolerance. Go/no-go on PFS.	Study and techno-economic adjudication of alternatives and selection of preferred option: Conduct appropriate trade-offs. Strategic and techno-economic profiling for optimum value within owner risk tolerance. Confirm preferred option. Preliminary reserve declaration. Go/no-go on FS.	Detailed study of the preferred option to provide adequate definition, confirm viability and support execution approval: Compile designs and scopes for key infrastructure and installations. Full reserve declaration. Develop operational readiness strategy. Confirm economics for financing and construction decision. Go/no-go on FEED/execution.	Definition and pricing of the execution scope to enable contracting: Produce construction-ready engineering packages. Finalise procurement strategy. Finalise EPC/ EPCM contract.
Key techno-economics requirements	Understand the key value drivers and viable extraction pathways of the orebody and determine range of economic potential. Identify candidate mining and access methods enabled by those drivers. Assess pathways based on the economic outlook (eg price cycle, basket price, demand trajectory etc). Define the upside and downside under geological and economic uncertainty.	Determine potential value and range of the opportunity. Determine the relative strategic importance of the opportunity (timing and rank). Determine key drivers of value, their variability and impact to the opportunity. Define level of confidence required in key drivers to make pre-feasibility decision. Define work required to achieve the required pre-feasibility confidence.	Determine single business case to take forward. Provide evidence of trade-off approach and assumptions made. Determine potential value and range of opportunity. Confirm alignment with business strategy. Provide benchmarks of all key drivers for comparison. Determine work that can be done to reduce residual uncertainty in business case.	Optimise proposed solution. Determine best estimate of all related costs and returns. Provide benchmarks of all key drivers for validation. Document assumptions for validation. Finalise scope, schedule and cost of implementation. Compile detailed risk assessment of implementation phase. Prepare detailed project execution plan. Prepare investment case and sensitivity analysis to drivers. Prepare value tracking baseline for the execution and operations phases.	Monitor against techno-economic baseline. Continually review value drivers and market conditions against project baseline.
Resource/ Reserve	Exploration targets Early Inferred	Inferred Some Indicated	Indicated Some Measured	Measured Proved/Probable Reserves	Proved/Probable Reserves No change expected
Engineering depth	Nil; desktop only	~1–3% of total project engineering	~10–15% of total project engineering	~25–40% of total project engineering	~60–100% of total project engineering. All major discipline packages complete.
Typical cost (USD)	Very low Internal effort only	Low. $0.1M – $1M	Moderate. $1M – $10M+	High. $10M – $100M+	Very High. 1–5% of total project CAPEX (eg $20M–$150M+ on a major project).
Duration	Weeks	1–4 months	6–18 months	12–36 months	12–30 months (runs partly concurrent with permitting and early procurement).

Further, project value tracking (of the business case and its drivers) is not routinely undertaken to track value degradation through the various phases. To mitigate these shortcomings, a proposed list of key techno-economic requirements for each of the study phases is presented in Table 6, to ensure sufficient rigour in the economic profiling of the project.

The project development process through the various phases is depicted schematically in Figure 10. The shaded envelopes represent the volume of work undertaken in each phase, the arrows represent the alternatives that may be contemplated, the line type indicates the degree to which each alternative adequately addresses the requirements of each phase, and the length of each arrow represents the depth of work undertaken in each alternative. Stage gates are represented by the rectangular blocks separating each phase.

The traditional approach depicted in diagram (A) employs a funnel in which alternatives are eliminated through the early study phases to arrive at the optimal solution which is elaborated at the FS phase and progresses into execution and operations. Due to the number of alternatives evaluated under this approach, each alternative relies on high level assumptions, and the depth of each alternative is necessarily shallower. The primary challenge with this approach is the time taken to undertake the volume of work required, the arduous assurance process at each stage gate, and the lack of flexibility to accommodate changes in assumptions (which is partially an outcome of the time taken for development), whilst demonstrating diligence in seeking the optimal solution. This approach is generally driven by a forward technical pass with the techno-economics determined at the end of each phase. A key defect with this approach is the absence or weak techno-economic input into the study direction. Timelines lengthen not because the engineering is slow, but because techno-economic requirements are regenerated at each phase rather than carried forward, and because the criteria governing progression are rediscovered within each phase rather than treated as known.

How this works out in practice is demonstrated in diagrams (B) and (C).

The approach represented in diagram (B) is typically experienced in the major mining companies where all the study phases are undertaken. This outcome is a consequence of the challenges implicit in the traditional approach process which results in alternatives and work being carried over to subsequent phases, together with unavoidable re-work. This increases the workload in the subsequent phases which undermines those deliverables and the intent on the particular study phase objectives, but also places pressure on the project development timelines. It is not uncommon for such issues to have to be addressed in the execution and operations phases. Considering that most value is generated in the study phases (by delivering an optimal techno-economic solution), any compromise in the study phases results in an opportunity (ie economic) cost that is not visible.

The approach represented in diagram (C) is typically experienced in the junior mining companies where selected study phases (eg PFS) are eliminated. This invariably occurs due to top-down pressure and on the assumption that the project is not complex and sufficient information is available to progress directly to FS. This results in key trade-off and value engineering rolling over to the FS phase; back-end loading effort, impacting quality due to resource overload and lack of focus. As in (B), technical and project issues are carried over into the execution and operations phases. The economic impact is even more pronounced than in (B).

The proposed approach represented in diagram (D) seeks to address the shortcomings in (A), (B), and (C). In the Concept Phase, deeper techno-economic modelling is undertaken to profile the project to understand the key risks, value drivers, dependencies, and information shortfalls. This allows for a better project development strategy and its associated planning.

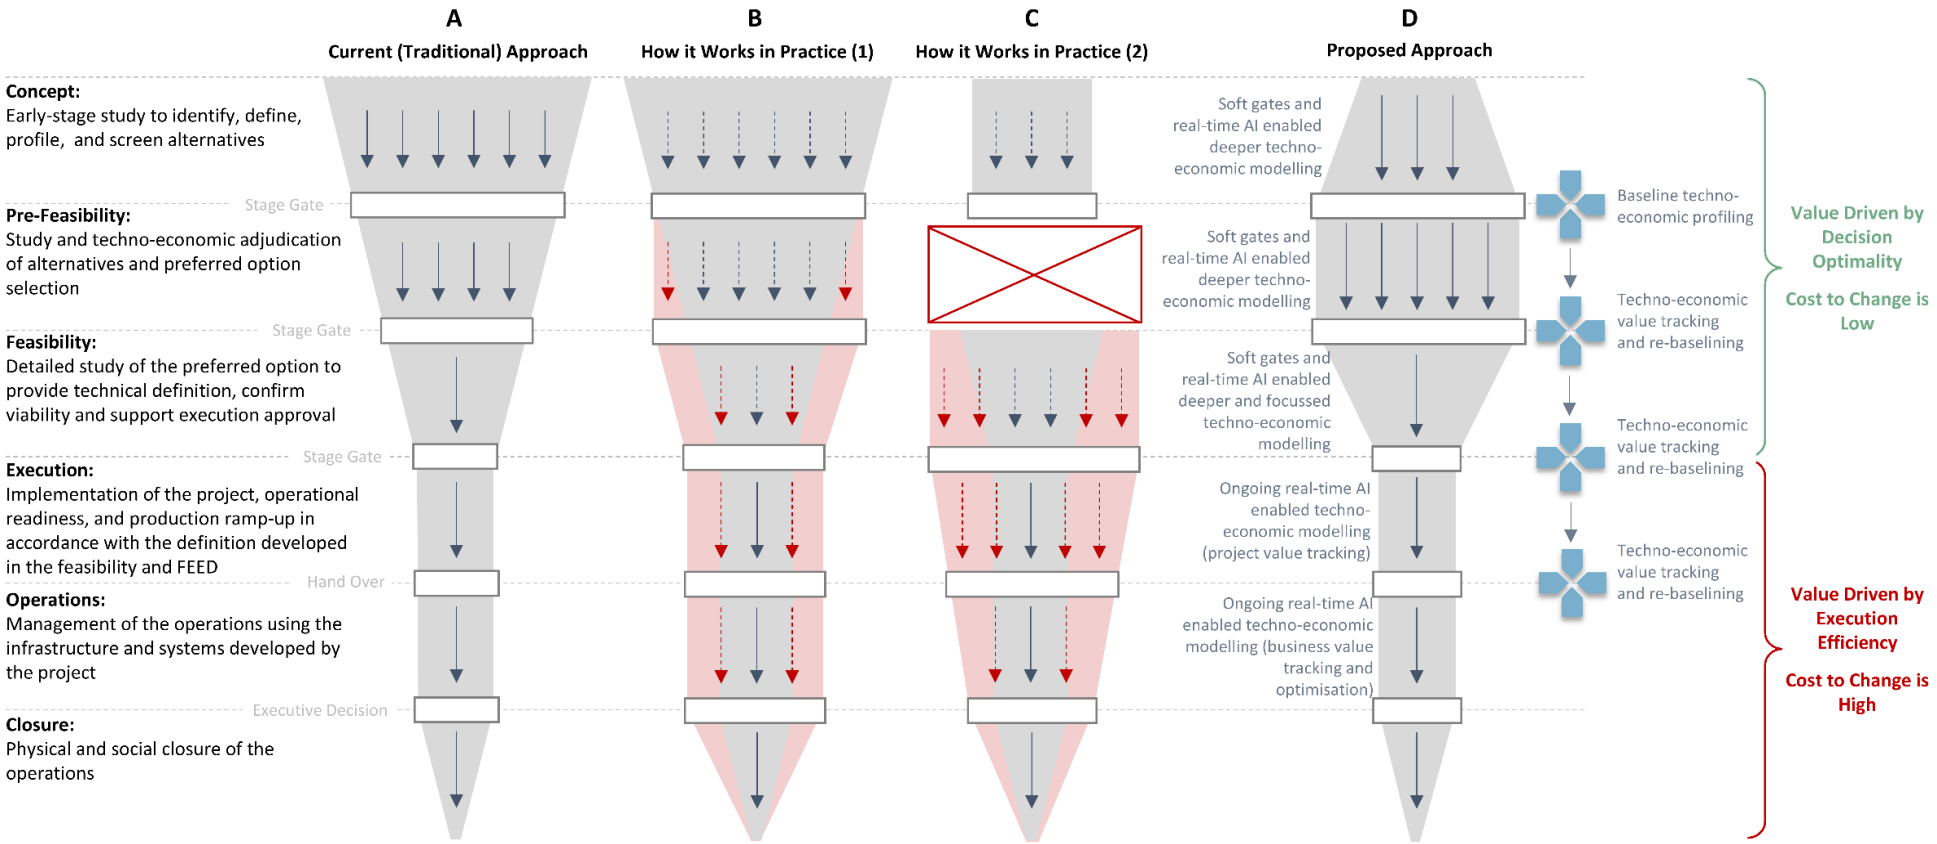


FIG 10 – Mining project development approach in theory, practice, and proposed.

A baseline techno-economic profile is generated which facilitates longitudinal visibility of key metrics and especially value leakage as the project is developed through the various phases. Applying reference-class thinking, as per Flyvbjerg (2021), to mining projects allows for project teams to define a narrower set of proven orebody access and mining methods that adequately define the technical and economic requirements of the project. The criteria distinguishing a mature concept from an immature one, and the work needed to close residual uncertainty at each gate, can be derived from historical studies, engineering work, trade-offs, and existing operations strategies on comparable orebodies. Deliberately classifying these criteria means each project begins from accumulated knowledge rather than a blank sheet, with the value baseline persisting as the study develops so that work informs the next phase rather than being repeated within it. Therefore, when geological and resource certainty is confirmed, the project development need not wait on reconstructing foundations already laid. This frees the study team to concentrate on the few drivers that determine value, rather than re-deriving context. This approach allows for fewer alternatives to be evaluated, but with a greater depth of techno-economic analysis to profile the project and devise a strategy for its development. Rigour is not traded for speed; it is redirected toward the decisions that warrant it.

Using the techno-economic profile generated in the Concept Phase, the PFS can test various technologies, configurations, and strategies in a focused approach. The techno-economic model can, where appropriate, be used to adjudicate decisions based on their potential impact to the outcomes of the project.

Within a particular phase, to avoid technical conflicts and resulting re-work, soft gates are proposed to obtain alignment and agreement at defined points, for example:

- completion of mining design
- completion of process design
- completion of engineering design
- completion of techno-economic modelling.

At culmination of the PFS, a single preferred alternative would have been identified, with the key techno-economic drivers refined and used to inform the FS and subsequent execution strategy. At this point the intent of the project development strategy remains aligned to the traditional model. The FS can therefore focus on ensuring a practical and realistic project can be delivered.

Project value tracking, as an update on the latest techno-economic baseline, provides insights to the project team on value leakage and the reasons for it; indeed, it is proposed to continue value tracking throughout the project and operational life cycle.

This approach, enabled by the advent of artificial intelligence, can accelerate the timeline for some workstreams (eg mining, processing, logistics, engineering, costing etc), however others may still require traditional timelines for the requisite work to be done (eg geological exploration, environmental assessments, regulatory approvals etc).

The time and effort in defining the optimal techno-economic solution will generally result in a more certain and enhanced investment case.

CONCLUSIONS

Mining projects continue to disappoint against their business cases, and the evidence suggests the causes are structural rather than incidental. This paper has argued that the failure is not primarily one of engineering competence, but of the relationship between economic reasoning and technical work. Where techno-economic analysis is treated as an output generated at the end of each study phase, it can only confirm or reject what the technical workstreams have already decided. Treated instead as an input that frames and directs those workstreams, it becomes the mechanism through which a project is steered toward value.

Three conclusions follow. First, strategy and execution are not substitutes. The quadrant framework presented here shows that disciplined delivery of a poorly framed extraction strategy produces the most dangerous outcome of all, because the resulting project appears sound while destroying value. Good execution cannot recover value leaked at the point of definition.

Second, the capabilities that determine operational performance are cumulative. The Sand Cone logic, extended here to treat cost as encompassing both financial and economic cost, offers project teams a means of ordering their trade-off criteria so that cost and throughput objectives are not pursued ahead of the quality, flexibility and dependability foundations on which they depend. Coupled with Slack and Lewis' reconciliation of competitive requirements against operational resources, it provides a vocabulary in which multi-disciplinary teams can express technical decisions in terms that connect to the investment case.

Third, the decision architecture matters as much as the analysis. Bounded rationality and opportunism operate on different dimensions of project development: the first degrades content, the second degrades process. Multi-criteria decision analysis, disciplined by the Decision Quality elements and separated explicitly from the economic assessment, addresses both. It exposes the reasoning behind a preference rather than only its result, which is what allows information asymmetry to be narrowed at each gate rather than carried forward.

The project development approach proposed here follows from these conclusions. Deeper techno-economic modelling in the Concept Phase, a persistent value baseline carried across gates, reference-class thinking to narrow the field of credible alternatives, and soft gates within phases together shift effort from re-deriving context to resolving the few drivers that determine value. Fewer alternatives are examined, but each is examined more deeply. Rigour is not traded for speed; it is redirected toward the decisions that warrant it.

Two areas warrant further work. The first is the formal treatment of trade-offs between competing economic criteria, particularly the tension between net present value and all-in sustaining cost, which arises from cut-off grade policy, process and infrastructure design philosophies, throughput selection and sustaining capital intensity, and which the industry currently resolves by convention rather than by analysis. The second is empirical validation: the framework presented here is grounded in established theory and in the authors' practice, but its claims about value preservation across study phases can only be substantiated through longitudinal measurement on projects that adopt it. We would encourage operators to instrument their studies accordingly, since the opportunity cost this paper describes remains invisible precisely because it is so seldom measured.

REFERENCES

Asad, M W A, Qureshi, M A and Jang, H, 2016. A review of cut-off grade policy models for open pit mining operations, *Resources Policy*, 49:142–152.

Bascetin, A and Nieto, A, 2007. Determination of optimal cut-off grade policy to optimize NPV using a new approach with optimization factor, *Journal of the Southern African Institute of Mining and Metallurgy*, 107(2):87–94.

Belton, V and Stewart, T J, 2002. *Multiple Criteria Decision Analysis*: *An Integrated Approach*, 372 p (Kluwer Academic Publishers: Boston).

Dagdelen, K and Kawahata, K, 2008. Value creation through strategic mine planning and cut-off grade optimization, *Mining Engineering*, 60(1):39–45.

Dussud, M, Kuvshinikov, M, Pikul, P, Price, R and Samek, R, 2018. Avoiding mistakes of the past: A CEO's checklist in a commodity upswing [online]. Available from: <https://www.mckinsey.com/industries/metals-and-mining/our-insights/avoiding-mistakes-of-the-past-a-ceos-checklist-in-a-commodity-upswing> [Accessed: 04 Sept 2026].

Ferdows, K and de Meyer, A, 1990. Lasting improvements in manufacturing performance: in search of a new theory, *Journal of Operations Management*, 9(2):168–184.

Flyvbjerg, B and Gardner, D, 2023. *How Big Things Get Done*: *The Surprising Factors Behind Every Successful Project, from Home Renovations to Space Exploration*, p 8 (Random House: New York).

Flyvbjerg, B, 2021. Top ten behavioral biases in project management: an overview, *Project Management Journal*, 52(6):531–546.

Flyvbjerg, B, Skamris Holm, M K and Buhl, S L, 2002. Underestimating costs in public works projects: error or lie?, *Journal of the American Planning Association*, 68(3):279–295.

Githiria, J and Musingwini, C, 2019. A stochastic cut-off grade optimization model to incorporate uncertainty for improved project value, *Journal of the Southern African Institute of Mining and Metallurgy*, 119(3):217–228.

Hall, B E, 2003. How mining companies improve share price by destroying shareholder value, in Proceedings CIM Mining Conference and Exhibition (Canadian Institute of Mining, Metallurgy and Petroleum: Montreal).

Hall, B E, 2014. *Cut-off Grades and Optimising the Strategic Mine Plan*, Spectrum Series 20, 312 p (The Australasian Institute of Mining and Metallurgy: Melbourne).

Hutton, C, 2022. *Mining Economics Explained*: *A Guide for Boards, Executives, Managers and Investors*, (Reach Publishers: Wandsbeck).

Kahneman, D, 2011. *Thinking, Fast and Slow* (Farrar, Straus and Giroux: New York).

Lane, K F, 1964. Choosing the optimum cut-off grade, *Colorado School of Mines Quarterly*, 59(4):811–829.

Minnitt, R C A, 2007. Frontiers of usefulness: the economics of exhaustible resources, *Journal of the Southern African Institute of Mining and Metallurgy*, 107(8):539–555.

Poniewierski, J and Hall, B, 2016. Break-even is broken, *AusIMM Bulletin*, June (The Australasian Institute of Mining and Metallurgy: Melbourne).

Rendu, J-M, 2014. *An Introduction to Cut-off Grade Estimation*, 2nd edition (Society for Mining, Metallurgy and Exploration: Englewood).

Schroeder, R G, Shah, R and Peng, D X, 2011. The cumulative capability 'sand cone' model revisited: a new perspective for manufacturing strategy, *International Journal of Production Research*, 49(16):4879–4901.

Simon, H A, 1957. *Models of Man: Social and Rational* (John Wiley and Sons: New York).

Slack, N and Lewis, M, 2015. *Operations Strategy*, 4th edition, p 11 (Pearson Education Limited: Harlow).

Spetzler, C, Winter, H and Meyer, J, 2016. *Decision Quality*: *Value Creation from Better Business Decisions* (John Wiley and Sons: Hoboken).

Williamson, O E, 1981. The economics of organization: the transaction cost approach, *American Journal of Sociology*, 87(3):548–577.

World Gold Council, 2018. *Guidance Note on Non-GAAP Metrics: All-In Sustaining Costs and All-In Costs* (World Gold Council: London).

Best practices for uncertainty management in mineral projects

M Silva[1], B Gadelha[2], L Montiel[3] and T Câmara[4]

1. Mining Director, Vale Base Metals, Toronto ON, Canada. Email: mario.silva@vale.com
2. RDP Manager, Vale Base Metals, Sudbury ON, Canada. Email: bruno.gadelha@vale.com
3. Principal Engineer, Vale Base Metals, Toronto ON, Canada. Email: luis.petro@vale.com
4. Specialist Engineer, Vale Base Metals, Belo Horizonte MG, Brazil. Email: tais.camara@vale.com

ABSTRACT

Uncertainty management focuses on understanding the variability within the solution space to strike the right balance between value creation and risk exposure. Across Vale Base Metals operations, several uncertainty-driven methodologies have been deployed to capture upside potential by explicitly modelling and incorporating uncertainty.

A stochastic mine-planning optimisation framework applied to an iron oxide copper–gold (IOCG) deposit in Carajás demonstrates how geostatistical simulations better represent the deposit's intrinsic variability than traditional estimated models. Validation against blasthole production data shows that simulations more accurately reproduce local grade–tonnage relationships. By integrating these simulations, the stochastic optimiser produces schedules that more realistically reflect geological risk.

Geostatistical simulations have also been used to optimise infill drilling accounting for quarterly and annual production plans in a nickel laterite deposit. Likewise, a stochastic blending-risk approach has been implemented in deep underground operations, where seismic and heat-related interruptions were incorporated into risk-based production plans. This approach prioritised slowing the sink rate and capturing near-term opportunities rather than rapidly mining deeper to maximise grade, as suggested by deterministic models.

Despite the value stochastic optimisation brings to operations and projects, several barriers limit broader adoption. These include limited understanding of uncertainty and its impact on project value, inadequate geological domaining that focuses solely on grade uncertainty, a shortage of mine-planning professionals trained in these methods, overly complex simulation frameworks, and organisational silos that hinder cross-department collaboration.

INTRODUCTION

In mining, as in any other field, projects and operations are inherently exposed to uncertainty. When uncertainty is present, different decision-making approaches can be adopted. One option is to ignore uncertainty and assume that the input parameters accurately represent reality, leading to decisions based on a presumed accurate prediction of future outcomes. However, it is well documented that deterministic-based decisions do not deliver expected results when uncertainty exists (Dimitrakopoulos, 2011; Montiel, Dimitrakopoulos and Kawahata, 2016; Câmara, Scheffer Leal and de Lemos Peroni, 2020).

Mining processes and related decisions are influenced by multiple sources of uncertainty, including commodity prices, deposit volumes and grades, operating costs, metallurgical recoveries, seismicity, and more. An alternative to the deterministic approach is to explicitly incorporate uncertainty into the decision-making process. Rather than attempting to predict the most probable or average scenario, the focus shifts to modelling uncertainty, integrating it into optimization, and managing its impact on outcomes. Due to the non-linear relationships between inputs and outputs in mining systems, an average-type scenario does not yield a solution that performs well on average when uncertainty is present.

A key question therefore arises: which sources of uncertainty are driving poor performance in the process or plan? Addressing this requires a thorough analysis of the decision context and a clear understanding of how uncertainty affects expected outcomes. Effective uncertainty management involves understanding the uncertainty within the solution space to achieve an optimal balance

between value and risk. Instead of identifying a single optimal solution for a fixed input scenario—one that may perform poorly when conditions change—a stochastic approach identifies solutions that remain resilient across a range of possible outcomes, as illustrated in Figure 1.

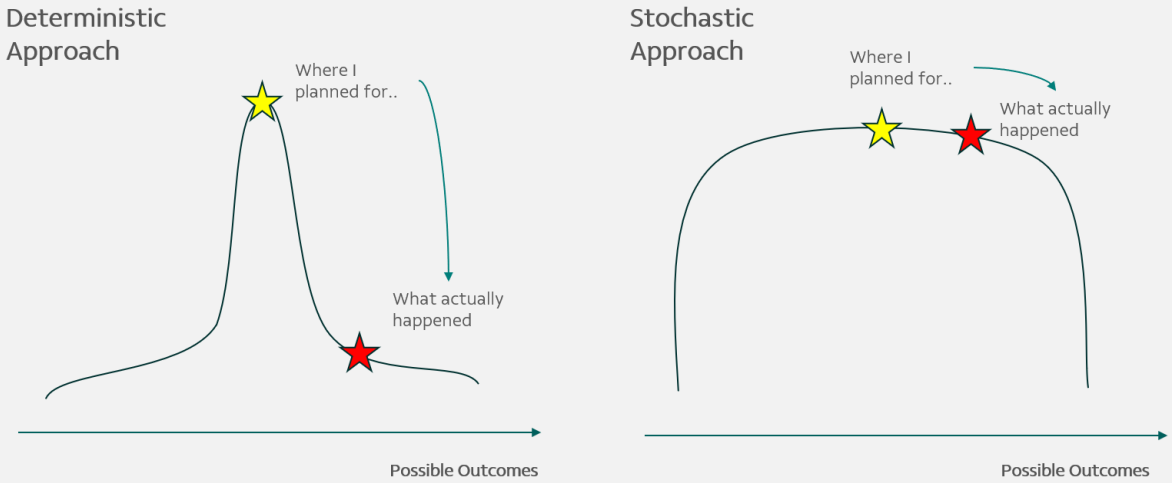


FIG 1 – Deterministic versus stochastic approach.

This paper presents three applications of uncertainty management at Vale Base Metals. The first demonstrates the use of a stochastic mine planning optimiser for an iron oxide copper–gold (IOCG) deposit in Carajás, Brazil. The second highlights the application of geostatistical simulations to optimise infill drilling in a nickel laterite deposit. The third illustrates how incorporating seismicity uncertainty leads to more reliable underground mine plans. Conclusions and final remarks are provided in the closing section.

RECONCILING POWER OF STOCHASTIC PLANNING IN AN OPERATING MINE

A stochastic mine planning optimisation framework was applied and reconciled at a multi-pit copper operation within an IOCG deposit in Carajás, Brazil. A set of categorical and grade simulations was generated using the same 2001 drill hole database that supported the kriged resource model used for the feasibility-level mine plan.

General overview

The mine has been in operation for more than a decade, providing extensive grade control data that enables reconciliation between the estimated block model and the geostatistical simulations. Grade control is based on blast-hole drilling on 16 m benches with 2 m of sub-drilling. Drilling follows a triangular pattern with a typical spacing of 5–8 m (Figures 2 and 3).

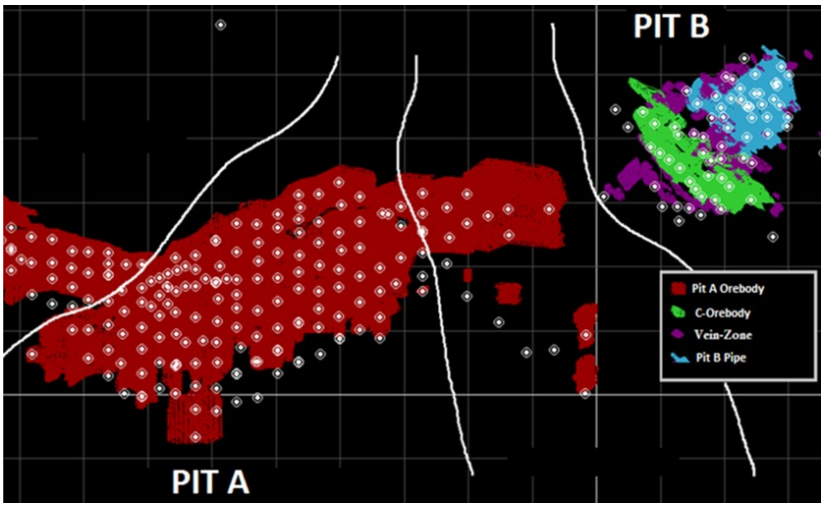


FIG 2 – General overview and drill hole locations.

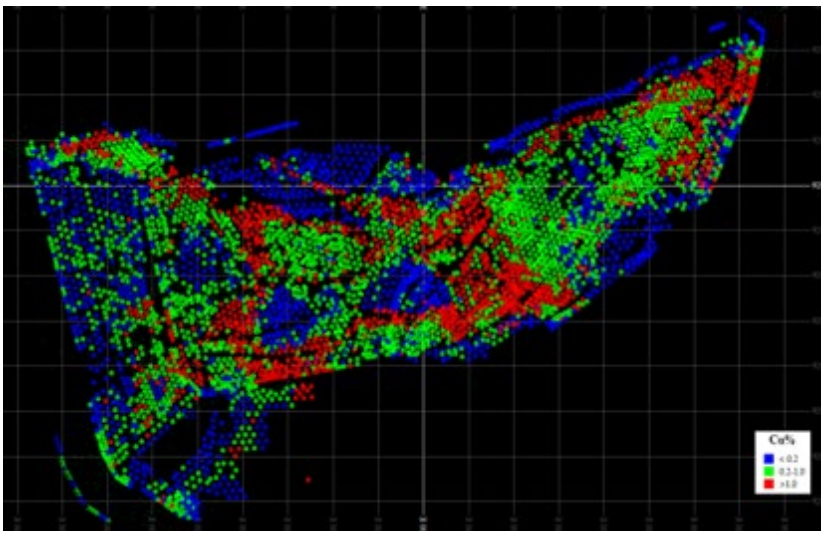


FIG 3 – Grade control data in Pit A.

Blocks in each mined bench were estimated using the nearest-neighbour (NN) method with 16 m composites and a 20 m search radius. This blasthole model—derived from dense grade control data—was compared with both the kriged model and the geostatistical simulations (Figure 4). The simulations show significantly closer agreement with the blasthole model than the kriged estimate, demonstrating their ability to reproduce the true intrinsic variability of the deposit, whereas the kriged model exhibits the expected smoothing effect.

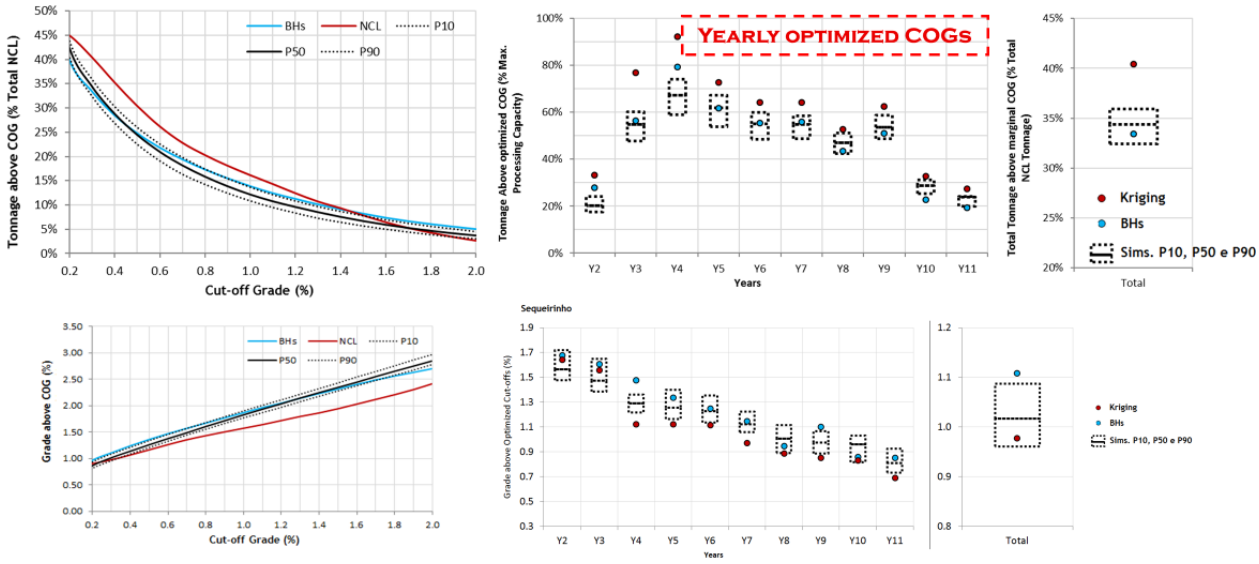


FIG 4 – Blasthole model versus kriging versus simulations – Pit A.

Risk analysis

After validating the simulations against the blasthole model, the next step was to quantify the geological risk associated with the feasibility-level mine plan. A risk analysis was performed by implementing the Whittle-derived schedule—based on the kriged model—across all geostatistical realisations (Figure 5). Risk was expressed using P_{10} and P_{90} percentiles: P_{10} represents the level below which 10 per cent of outcomes fall, while P_{90} represents the level below which 90 per cent of outcomes fall. The P_{10}–P_{90} range captures the most probable 80 per cent of outcomes.

The analysis shows a consistent production shortfall in Pit A across all years, indicating that actual mill feed is expected to be lower than predicted by the kriged model. This shortfall propagates to the total mill feed, with direct implications for the operation's forecasted value.

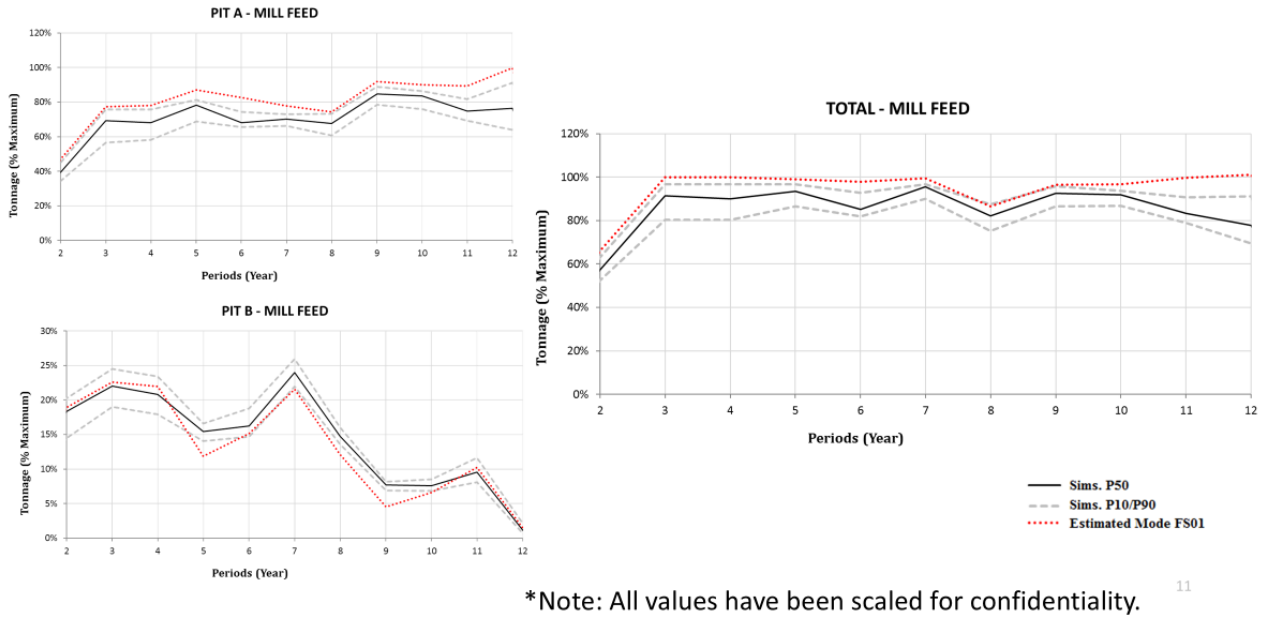


FIG 5 – Risk analysis deterministic schedule.

Stochastic optimisation

The stochastic optimisation framework uses geostatistical simulations as direct inputs. Its objective is not only to maximise net present value but also to explicitly manage geological risk. By deferring high-risk areas—particularly during the payback period—the approach enhances the robustness of early-life production, when reliable performance is most critical and additional information will be collected as mining progresses.

The method does not require predefined ultimate pit limits, nested shells, or pushbacks. Instead, it schedules individual blocks directly, simultaneously optimising extraction timing, destination policies (eg plant versus stockpile), and processing decisions (eg stockpile reclaim). Cut-off grade policies are therefore optimised concurrently with scheduling, rather than in separate steps.

In its current form, the optimiser does not generate fully smooth operational schedules, and maximum sink rates are only locally controlled. A post-processing step is required to convert the optimised output into operationally feasible geometries.

Figure 6 compares the geometries of the Whittle schedule, the stochastic schedule, and the actual mined geometry in Pit A. The stochastic schedule shows notable differences from the deterministic one, prioritising areas that the deterministic plan defers.

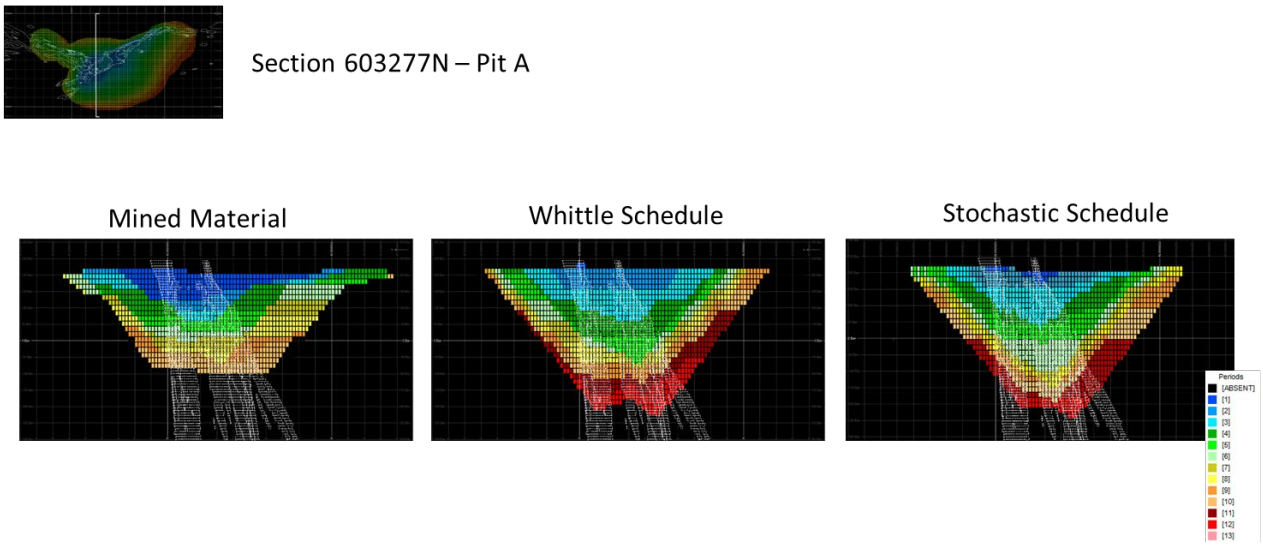


FIG 6 – Comparison of schedules geometry.

Simulations indicate that the deterministic schedule will not supply sufficient mill feed in several years (Figure 7). In contrast, the stochastic schedule explicitly manages the risk of under-feeding the mill, improving reliability of plant utilisation.

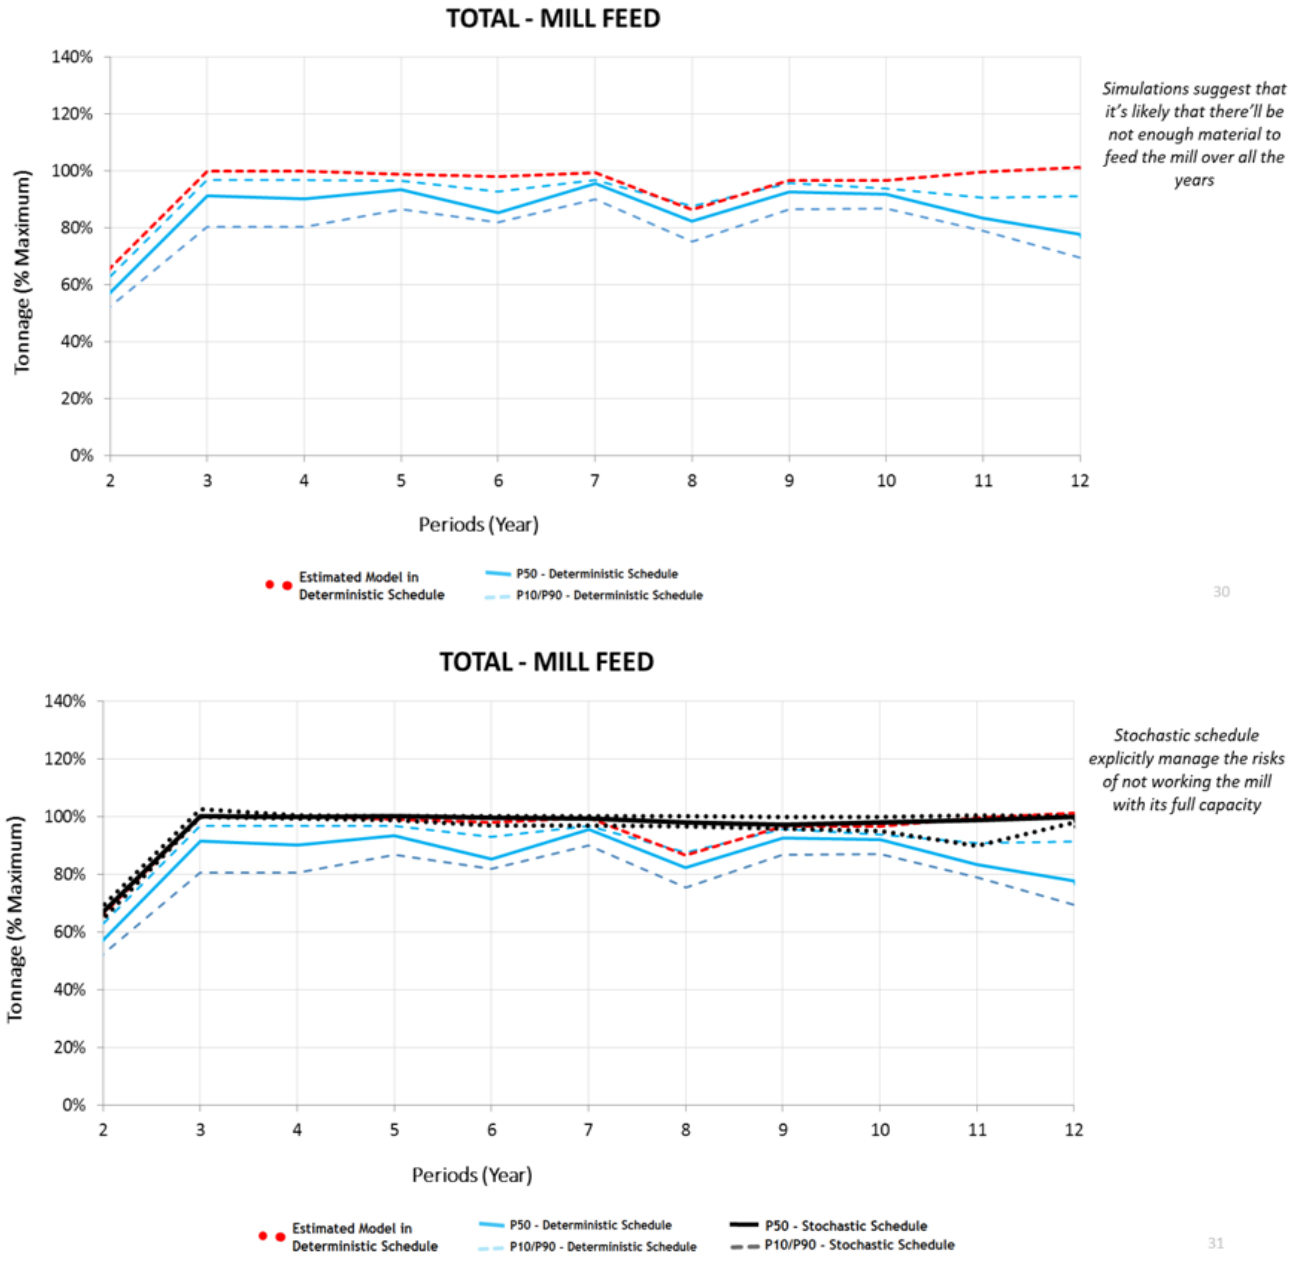


FIG 7 – Back Analysis of stochastic versus deterministic.

Takeaways

This case study demonstrates that average-type models, such as kriged estimates, produce smoothed representations of the deposit that lead to unrealistic mine schedules and unmet production expectations, as shown in the risk analysis. Comparisons with the blasthole model confirm that geostatistical simulations better reproduce the intrinsic variability of the deposit and provide more accurate local grade-tonnage relationships.

The stochastic optimiser generates more realistic and resilient schedules by explicitly incorporating geological uncertainty. Unlike traditional approaches, it does not require predefined ultimate pit limits or pushbacks, and it optimises destination policies simultaneously with extraction decisions.

OPTIMISING INFILL DRILLING

Optimising drilling pattern

Uncertainty management at Vale Base Metals also includes the use of geostatistical simulation to optimise infill exploration drilling. The core principle is that mineral resource classification is fundamentally tied to the level of confidence in the economic potential of the deposit. Consequently, classification criteria must reflect confidence in key production parameters—such as production rate, head grade, and life-of-mine projections.

As resource classification progresses from inferred to measured categories, the required confidence threshold increases, which intuitively corresponds to smaller evaluation volumes or shorter evaluation periods. Benchmark criteria are therefore established based on confidence in both the overall project life (global tonnage and grade) and the time-dependent production metrics used in discounted cash-flow analysis.

These benchmarks follow AMEC's recommended guidelines, which are considered reasonable in most situations:

- Inferred Resource: ±15 per cent on global tonnage and global grade.
- Indicated Resource: ±15 per cent on production rate and head grade for a volume equivalent to approximately one year of mining.
- Measured Resource: ±15 per cent on production rate and head grade for a volume equivalent to approximately one quarter of mining.

To achieve these confidence levels, a 'Virtual Reality' (VR) Back Analysis framework is applied to determine the optimal infill drilling spacing. In this approach, nickel laterite layers are modelled using PluriGaussian simulation, while nickel grades are simulated using the Turning Bands method. One simulation realisation is selected as the VR scenario. Multiple drilling pattern configurations are then tested on this VR scenario, and the resulting synthetic drill data are used to generate updated estimated models and new sets of simulated realisations (Figure 8).

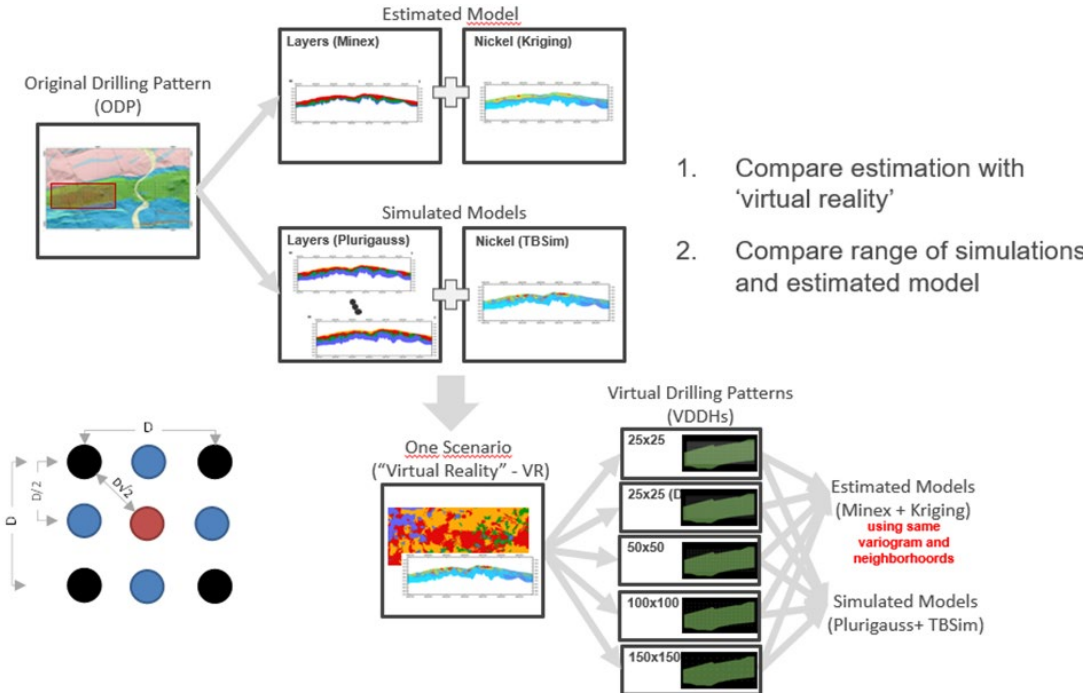


FIG 8 – VR Back Analysis methodology for infill drilling.

The results of this VR Back Analysis workflow are summarised in Figure 9. The study concluded that a 50 m × 50 m (D) and 50 m × 50 m were sufficient to achieve the required confidence thresholds for measured and indicated resource categories respectively.

This methodology effectively captures the spatial uncertainty inherent in the deposit. In some areas, the drilling grid was tightened to better define spatial continuity and improve model reliability.

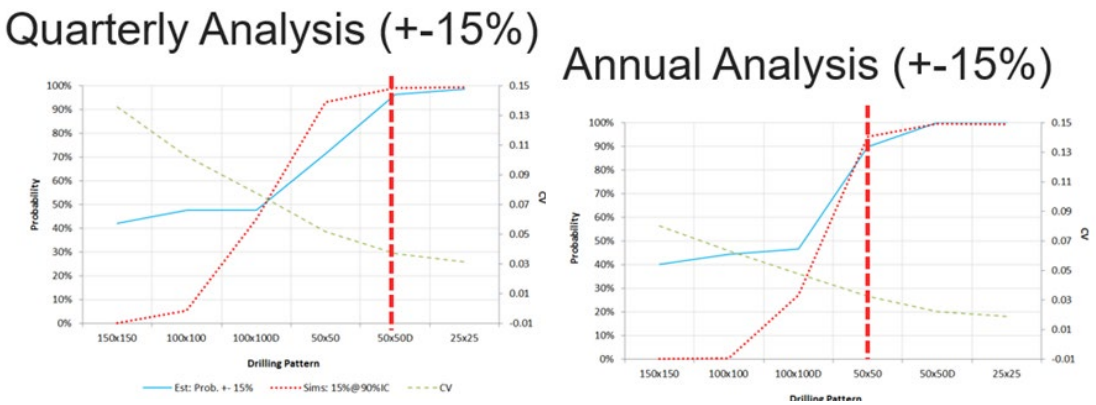


FIG 9 – Results of VR Back Analysis methodology.

DEEP MINING

The shift moves from a deterministic strategy—driving rapid deep mining to maximise grade but suffering major production interruptions from extended seismic no-entry windows and heightened heat-rest requirements, ultimately resulting in unreliable output—to a stochastic strategy that balances risk and grade by moderating sink rates and capitalising on accessible opportunities (Figure 10). By simulating production schedules that incorporate seismic and heat-related disruptions, the operation achieves more reliable performance and a higher expected value.

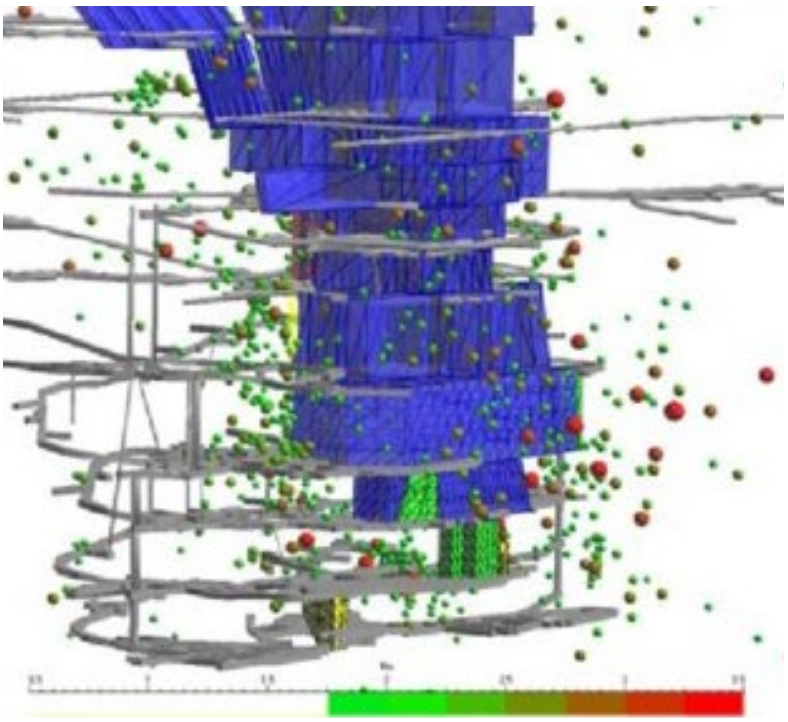


FIG 10 – Seismic and heat interruption events in deep underground mine.

CONCLUSIONS AND CLOSING REMARKS

The examples of uncertainty-based methodologies applied at Vale Base Metals demonstrate clear advantages over traditional optimisation frameworks that overlook uncertainty; however, several challenges continue to limit broader adoption. These include a limited understanding of uncertainty and how it affects project outcomes, particularly when determining which variables are truly influential. Geological domaining also remains a weakness, as efforts often concentrate on grade uncertainty while relying on poorly defined geological boundaries. Compounding this is a shortage

of qualified mine-planning professionals, which makes it difficult to consistently deliver even foundational work and leaves little capacity for more advanced stochastic analyses. Many simulation and assessment models are overly complex and difficult to maintain, and organisational misalignment between valuation, technical, and study teams further hampers effective integration. Additionally, poor practices and misconceptions surrounding the use of stochastic models—such as the overreliance on simple Monte Carlo simulations that can distort the communication of uncertainty—continue to undermine the effectiveness of uncertainty-driven decision-making.

REFERENCES

Câmara, T R, Scheffer Leal, R and de Lemos Peroni, R, 2020. Accounting for operational dilution by incorporating geological uncertainties in short-term mine planning, *Dyna*, 87(213):178–183.

Dimitrakopoulos, R, 2011. Stochastic optimization for strategic mine planning: A decade of developments, *J Min Sci,* 47:138–150.

Montiel, L, Dimitrakopoulos, R and Kawahata, K, 2016. Globally optimising open-pit and underground mining operations under geological uncertainty, *Mining Technology*, 125(1):2–14.

From file-based workflows to decision support systems – a framework for scalable mine planning in the cloud era

L Victor[1] and J Mackenzie[2]

1. Senior Technical Lead – Platforms, Maptek Pty Ltd, Adelaide SA 5065.
   Email: luke.victor@maptek.com.au
2. Global Strategy Manager, Maptek Pty Ltd, Brisbane Qld 4000.
   Email: james.mackenzie@maptek.com.au

ABSTRACT

Over the last decade, there has been a consistent and accelerating increase in the quantities of mining data stored and the rate of mining data generation. Yet many mine planning workflows have remained unchanged, constrained by fragmented file systems, reliant upon individuals for rigorous procedure-following, limited by manual data reconciliation and desktop-bound processing. The challenge facing the industry is shifting from data collection and accessibility towards trusted 'single sources of truth' while largely ignoring the need for trusted, repeatable, transparent and auditable processes to connect data acquisition to decision-making.

This paper proposes a practical framework to transition from manual desktop-centric workflows to decision support systems. At the core of this framework is a structured interrogation of every planning process through five foundational questions: Why is a given decision pathway followed? Under what conditions should that decision change or be reviewed? Where does this data originate? How is it used and transformed? Where does this information go next?

By using these questions, planning activities can be split into three coordinated layers:

1. Revisioned data foundations – Data sources stored with full and recoverable revision history for robust data recovery, edit and lineage tracking for auditability and confidence in data, and access controls for security and data safety.
2. Scalable compute pipelines – Repeatable processing and data transformations executed within a controlled environment to remove variability. Automated computation and validation to reduce latency between data recording and human analysis.
3. Orchestrated decision logic – Structured workflows to encode data validation, repeatable processes and human-in-the-loop review points, converting standard operating procedures and tacit experience into transparent institutional knowledge.

This paper draws upon structured industry engagements and early-stage implementation programs with large-scale mining operations to explore these methods of shaping existing operational processes into decision support systems. Rather than replacing mining professionals, decision support systems require a shift by both organisations and individuals toward analysis and systemic improvement. Questions of data ownership, correct methodology, exception handling, approvals and authority are central to this transformation. This paper examines common adoption barriers, including resistance to losing file-based autonomy, and outlines strategies for progressive implementation which preserve professional judgement and human-in-the-loop decision-making.

As mines pursue increasingly integrated and automated planning systems, the ability to encode data lineage and computation together with transparent decision governance will become critical to faster, more repeatable decision-making. The framework presented offers a pathway for organisations seeking to move beyond individuals manually following standard operating procedures toward scalable, resilient decision architectures to support governance of planning processes in uncertain operating environments.

INTRODUCTION

Significant investment over recent years within the mining industry has been dedicated to capturing and generating increasing quantities of data. It is observed that tier 1 mining companies have access

to petabytes of existing data (CCIWA, 2021) and produce new data at rates of multiple terabytes per minute (McLean, 2018).

Raw data is produced from a wide range of measurement sources such as geological drilling records, measure while drilling (MWD) records, chemical assays, geospatial survey measurements, in field recording systems, equipment dispatch systems, and smart equipment logging.

Data is also produced by a range of post capture analyses such as geological modelling, mine design, scheduling, reconciliation of designs against as-mined data, preventative equipment maintenance systems and optimisation systems.

Research shows that while this data has allowed for significant improvements to analysis and optimisation, decisions in the mining industry continue to rely upon fragmented workflows and manual interpretation of multiple information streams (Kruger, Sparrow and Defty, 2023).

At the same time, the pipeline of new mining engineers remains constrained and uneven. Australian graduate numbers (Australian Network of Mining Engineering Educators (ANMEE), 2025) remain below their mid-2010s peak despite a recent recovery as seen in Figure 1, while US Mining and Mineral Engineering completions (Datausa.io, 2022) continue to fluctuate.

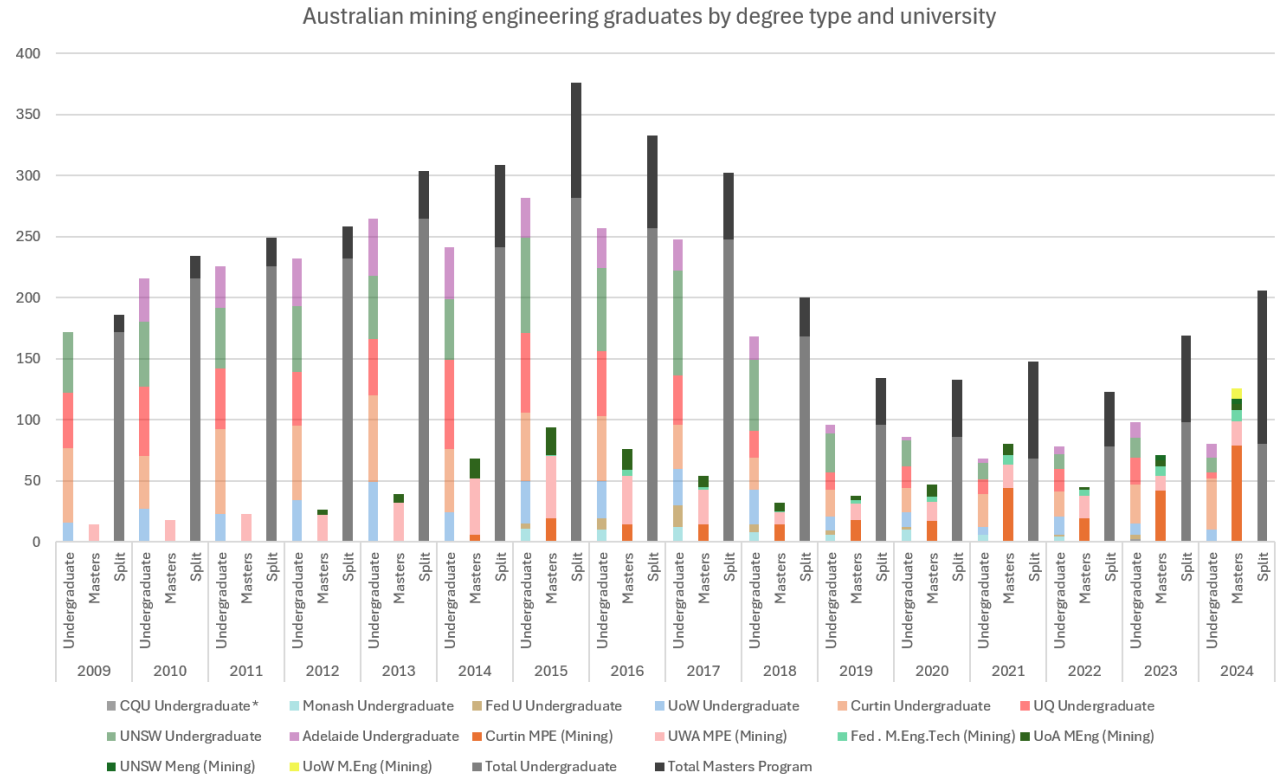


FIG 1 – Australian mining engineering graduates by degree type and university. Adapted from ANMEE (2025).

These trends suggest that the combination of increasing quantities of data produced, existing manual interpretation and handling of data based on written procedures, and a decreasing pool of professional graduates will lead to reduced operational performance as mining complexity and demand continue to increase.

To shorten planning cycles, trusted, auditable and repeatable decision-making is required.

Examples include deciding whether to revise a strategic plan following a material price or reserve change, regenerate a schedule after an approved model or design revision, update an orebody model after a drilling campaign or initiate corrective action when reconciliation identifies a material variance.

LIMITATIONS OF PROCESS CENTRIC PLANNING

A push to shorten planning cycles by the industry leads to a requirement for reducing the time between taking a measurement and making a decision based upon that measurement. This paper refers to the delay between measurement and decision as ‘decision latency’.

Mining data growth

As mining companies increased the quantity of systems and processes generating data, data storage has become a focus of mining technology and software providers with terms such as ‘single source of truth’ becoming commonplace.

Many measurement systems provide automatic measurements and store data to connected databases. This decreases the time between measurement and data being accessible but also increases the frequency with which measurements are made and data becomes stale.

The introduction of new measurement systems also leads to an increase in the number of disparate data streams which mining professionals must interact with and consider.

Better decisions require better data management

The result of having access to many different data sources, and rapidly changing or real-time data from those data sources is decision latency. This can occur because of many reasons:

- Data is stored in multiple disparate systems, each of which requires a person to manually access data or make copies locally. This can lead to multiple file versions existing for a single source of data, of which some may be out of date.
- As new sources of data are brought online and procedures are updated on where to access data, often the individual decision-maker is not accessing the direct source of data but rather a copy made by others. This can lead to gaps existing between the copy and the latest most up to date information.

In the author’s professional experience this gap can extend to months, and sometimes different individuals performing the same task refer to different copied and edited versions of the same data source.

Consider for example a quarterly strategic schedule evaluation which begins with the scheduler taking a copy of the approved geological model. Recently, one month of work has gone into updating the model, but the geology department does not finalise the update until shortly after the scheduler has made their copy. The schedule is developed and is technically correct against its input data. The decision latency between the last geological update and the schedule review in which it is used now spans a month of geological modelling, plus three months of waiting for the next schedule evaluation plus the length of time spent developing the next evaluation.

Following versus improving processes

To reduce the impact of individuals utilising different data sources and methods, traditional planning workflows often follow an explicitly written standard operating procedure. Written procedures are seen as a tool to collectively capture, endorse and improve upon best practices. Usually, procedures are written by experienced mining professionals and officially approved or endorsed for use company wide by management. These procedures typically capture much of the experience of contributing professionals, but rigid procedures can prevent the use of expert judgement and creative solutions in exceptional cases.

Steps included often range from which shared network drives and folders to access data from, what steps to follow and tools to use, and a series of provided assumptions and parameters to be used as a reference during a task. What is omitted from written procedures is expert knowledge of exceptional circumstances which should trigger deviations from the standard input parameters. This expert knowledge is the most valuable addition that experienced professionals provide.

INTERROGATING A DECISION

Orchestration of decision logic requires significant understanding of the underlying data, transformations and logic which are often provided by a combination of standard operating procedure and expert technical advice.

The following five questions may be applied to uncover the data, transformation and human analysis needs for a decision to be repeatably, reliably made. Often, including in the example explored later in this whitepaper, these questions open up information about further data and transformations which must also be interrogated.

As an example, we can explore each question against the following decision: Does the current strategic schedule need to be revised when new inputs are generated?

Why is a given decision pathway followed?

This question exposes the underlying logic which supports each decision pathway. Investigating and documenting the logic underpinning a given decision provides an opportunity to make decisions more transparent for auditing purposes. Transparent decision logic also provides onboarding opportunities for new members of a decision-making group as well as exposing any underlying assumptions for future continuous improvement efforts. In some simple scenarios, analysis of decision logic will allow for some decisions to be automated.

For our example decision, the existing schedule is used downstream for planning and execution. If the assumptions and inputs do materially affect the net present value of the project, a review will be necessary to ensure resources are spent efficiently.

Under what conditions should that decision change or be reviewed?

Many decisions have more than two possible outcomes. While a standard operating procedure may only outline one or a small number of most likely outcomes, they may also fail to describe all of the potential anomalies where an expert's input may be required to achieve a good outcome, or define what differentiates good and bad outcomes. Understanding when each possible outcome should and should not be selected can help to identify areas of ambiguity where expert analysis is required, establishing what circumstances indicate that a decision should be escalated for more careful consideration. Exploring this question may also expose areas where continuous improvement efforts may provide some benefit.

Returning to the example decision, the goal of this question is to determine triggers for when a change to our inputs will be significant enough to warrant reviewing the schedule. Triggers for this case might include price variations greater than 5 per cent, any change to pit design or ore modelling exceeding 100 kt difference in ore reserves, or a variance in production rate exceeding 3 per cent.

Where does this data originate?

When looking at a decision, it is important to identify where each piece of input information is sourced. This step may be as simple as identifying the system which is recording the data, or as complex as discovering a trail of files copied, imported and exported through multiple software packages and teams. In the latter case, opportunities to reduce the handling and re-handling of data can improve decision latency.

In the case of our example decision, data may originate from many sources. Price estimates or sale contracts, financial reporting, geological models, mine designs, mineral reserving and equipment productivity data from site. In each case the purpose of this question is to identify where the correct source for these inputs is and how to effectively access them.

How is it used and transformed?

After tracing the source of any given data, it is also important to understand what transformations the data undergoes before reaching the decision-maker. Mining data which are stored in disparate systems may undergo multiple export and import steps, data validation and calculations before being used for a given decision. The impact of those transformations should be documented and

understood from an auditing and technical de-risking perspective, and opportunities to reduce unnecessary data handling may be uncovered while investigating this question. This question also exposes opportunities to automate data transformations and further reduce decision latency by performing transformations where the data is stored instead of transferring inputs and outputs.

This question might identify for our example decision how to take the selling price from the most recent sale contract, what export and import processes to run to bring pit designs and reserves into the scheduling package, and how to average the latest equipment productivity data for input to the schedule. This is also where details of the procedure for generating a schedule can be captured and codified.

Where does this information go next?

While it is important to ensure that decisions are made correctly with all the latest available information, the decision itself is of no use unless it is communicated and actioned appropriately. This question allows the decision and any processed outputs to be passed on to the next steps and downstream teams and processes. As mining companies typically rely upon a large suite of disparate systems and software tools, this step may expose further validation and transformation steps for automation as well as identifying data sharing and access requirements for other teams.

Using our example decision, identifying the budgeting, infrastructure planning, operational planning and management teams which rely upon the schedule allows for these stakeholders to be updated when the decision must be changed.

REFRAMING PROCESSES AS DECISION SYSTEMS

This paper proposes that several areas of improvement are available to reduce decision latency compared to traditional written procedure methods for mining decisions. Decisions in the mining environment can be conceptualised as having three main components: revisioned data foundations, scalable compute pipelines and orchestrated decision logic.

Revisioned data foundations

While many mining companies have implemented revisioned document controls and approvals processes, it is common for mining data to be controlled by manual copying between shared folders and rigorous naming conventions. To reduce the risk of accidental deletion or overwrite of critical data, write permissions are often tightly controlled by IT teams who are separated from the technical staff requiring access. This can lead to delays in accessing data and duplication and copying of data between multiple shared locations, as well as errors in selecting the most up to date source of data.

One method for reducing reliance upon manual data transfer is use of a data management system which natively supports revisions on mining data and records details of each revision created. Such systems allow mining data to be directly referenced, without relying upon naming conventions and date checking as data can be updated without changing file names. By managing updates to data as revisions on a single object, ambiguity in which source of data is most up to date is reduced. By ensuring that data can be reverted to earlier revisions, the risk of data deletion or inadvertent edits is also reduced. By recording details of edits made, the management and auditing of changes to the data becomes easier and more transparent.

An example of revisioned data might be a single file where the latest approved schedule is stored. Each time a new approved schedule is created, it is created as a new revision on the existing file and details of the geological model, pit design, cost assumptions and other inputs used are recorded with that revision. Older revisions are stored in the same data source, and their results, assumptions and inputs used are still available for reference if needed.

Scalable compute pipelines

The existing practice of using written procedures directly corresponds to the need for repeatable, reliable outputs from a given set of input data. By transitioning from manually followed steps in a document to a script or automation performing data transformations, decision latency can be reduced by reducing reliance upon mouse and keyboard inputs.

Building data transformations in a controlled, virtualised environment, removes several risks such as variable hardware, software versions, networks and operating systems. Ensuring that compute can scale up and down with demand can also reduce the time taken for analyses to run (Mackenzie and Moncrieff, nd), further reducing decision latency.

Some examples of repeatable compute tasks might include validation checks on new data coming in from drilling campaigns, running a pit optimisation or schedule with a fixed set of input data sources or comparing monthly actuals against plan.

Orchestrated decision logic

It may be helpful to imagine the decision process and the flow of data as a liquid moving through pipes. For this heading, consider newly recorded data as cold water entering a reservoir. Picture a decision that must be made as a glass being filled at a tap at the end of the pipes. It is not necessary for the cold water to travel instantaneously through every pipe each time a drop of water is collected, however when a glass must be filled at the tap, the amount of time taken to travel from the reservoir to the glass determines how warm the water has become.

The same is true for decisions. Decisions do not always need to be revisited each time a new piece of data is recorded. For example, it is not practical to re-estimate the grades of a deposit each time a drill hole is logged. However, when a decision does need to be made, every delay between the recording of a piece of data and analysing the outcomes adds to the decision latency.

Within many standard operating procedures, there are steps which require thought and analysis from a mining professional and steps which are simply repeating pre-defined actions in a sequence or copying data and parameters. By linking the source of the data directly to where the data is stored, automating the repeatable validation and transformation steps, and notifying a human or group of humans when analysis is required, it is possible to reduce decision latency. A system which watches for when a decision is required and automatically runs the repeatable data transformations at that time can reduce repetitive manual effort and ensures that the latest available data is used at the time a decision must be made. An illustrative comparison diagram is shown in Figure 2.

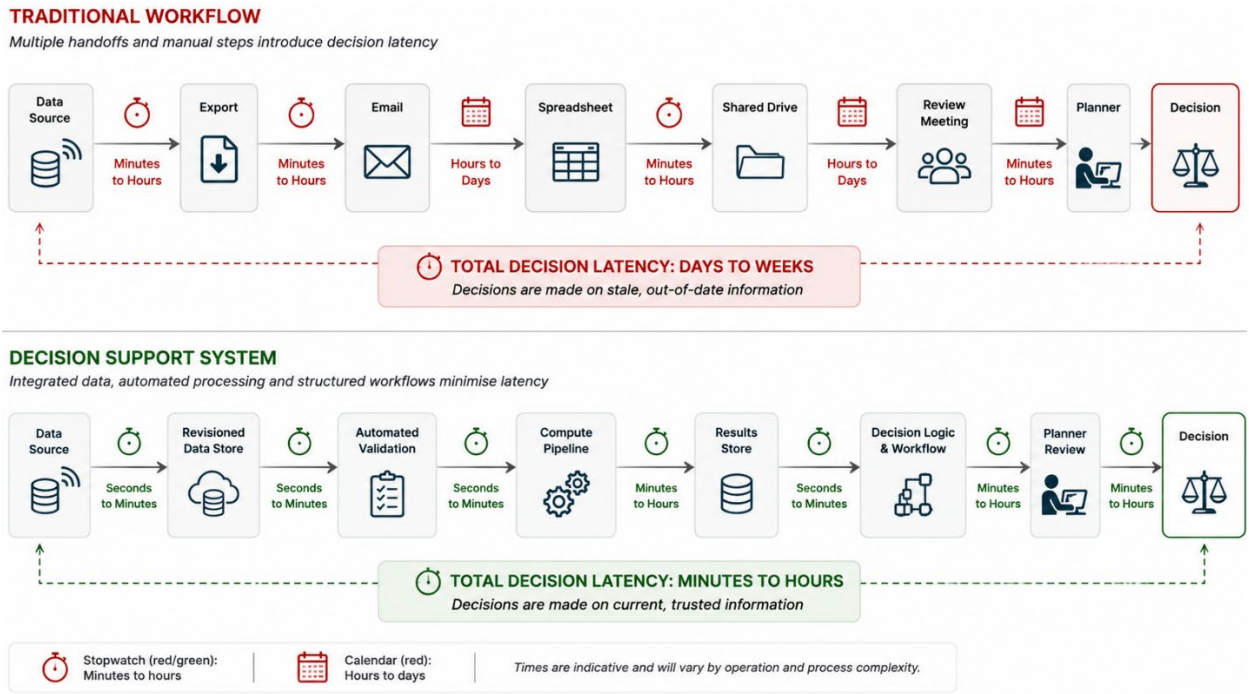


FIG 2 – Illustrative comparison of traditional workflow to decision support system workflow.

Using the earlier example of re-estimating grades based on new drilling, by automating the drilling import and validation and watching for completion of a planned set of infill or extension drill holes, an orchestrated system could present the geologist with the validated drilling, on the day that the drilling campaign reports its last hole, ready for immediate estimation.

DECISION MODELLING EXAMPLE – VOID MODELLING FOR TIER 1 MINER

The following example explores a body of work which the author's employer has undertaken for a tier 1 mining company. It is not the purpose of this paper to explore the risks of voids in open pit mining operations nor offer specific advice relevant to managing those risks, however for the purposes of establishing the project the following general statements will be offered:

- Pre-existing voids in an open pit mining operation present a real and significant danger of injury to operators and damage to mining equipment.
- Handling of incidents related to voids encountered during mining may cause lengthy delays and additional costs.
- Preventative measures typically rely upon targeted drilling for identification, and controls in identified risk zones typically involve reduced productivity which would be costly to universally apply across an entire mine.

Full details of the client's systems and tools have been deliberately left vague to avoid inadvertently identifying the client.

Existing problem

The original scope of work identified the need for a systematic approach to predicting subsurface voids and producing outputs in a format readable by existing client systems. The scope of work called out the need to incorporate measure while drilling (MWD) and reverse circulation grade control (RCGC) data which was already being recorded and generate trigger action response plans in elevated risk areas, and recommended standardising data flows for business improvement reasons.

Interrogating the required decision

For the purposes of this example, the decision required for the project can be stated as 'which areas require special treatment for subsurface voids?'. The following headings explore the problem through the lens of the five foundational questions.

Why is a given decision pathway followed?

In the case of this project, the existing reasoning was already clearly defined for two available pathways.

1. Additional controls are implemented for mining in identified elevated risk areas to prevent injuries, damage to equipment and downtime on mining benches.
2. Additional controls are not implemented for mining in other areas to improve mining rate and reduce costs.

Under what conditions should that decision change or be reviewed?

This question is arguably the most important for the void modelling project. In this case the purpose of the project was to identify in an automatic fashion the triggers for changing which decision pathway is followed. This objective was achieved using a geological domaining machine learning algorithm to model high and low risk areas in three dimensions.

Where does this data originate?

In this case, exploring the source of the data led to identifying cloud-hosted storage systems for MWD and RCGC data, as well as existing procedures for handling elevated risk areas. In order to minimise decision latency, specific views to the existing storage systems were sought for this project such that data can be queried directly using application programming interface (API) calls without relying upon manual export and copy operations.

How is it used and transformed?

The answer to this question depends in this case upon the proposed solution for making the decision. The solution implemented utilised a machine learning algorithm working on cloud native architecture

to perform the modelling of void risk areas. The specific requirements of the algorithm dictated the transformation requirements to get data into the correct input format and transformations were automated to ensure manual data handling is not required.

Where does this information go next?

Investigating downstream uses of the decision informed the need for breaking elevated risk areas down from three dimensional models to two dimensional polygons. Export file formats and downstream software systems were identified and automated data transformations implemented to ensure that outputs are made available where they are needed including geographic information system databases and desktop visualisation software.

By exploring all five of the foundational questions, a solution was developed, represented in Figure 3, which successfully imports data automatically from third party databases, runs validation and transformation steps without human input, produces a 3D model of high and low risk zones, produces 2D polygon outputs for visualisation and input to downstream systems and ensures that outputs are accessible by the relevant downstream parties.

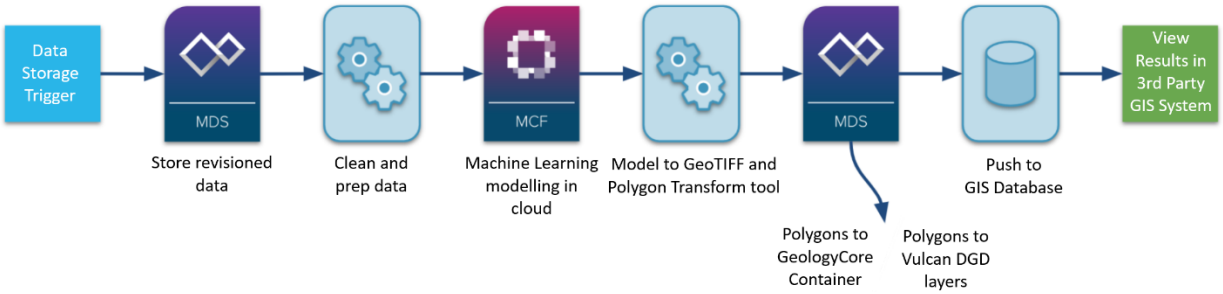


FIG 3 – Conceptual diagram of void modelling project solution.

ADOPTION CHALLENGES

Introducing new methods of handling data and making decisions can introduce significant changes to business processes and individual workers' responsibilities. Resistance to such changes can come from multiple stakeholders and can manifest in different ways. This section explores some of the potential causes and challenges associated with rethinking how mining decisions get made.

Cultural

While standard operating procedures already exist to bring about repeatable processes, in general terms many aspects of a mining professional's day are not explicitly covered within them. This leaves significant room for teams and individuals to develop their own ways of handling particular tasks and problems, which can be difficult to relinquish control over.

One point of resistance which can be experienced while implementing change is the loss of individual autonomy, especially around file ownership. When moving from a file based system where each group or department maintains their own copies of data and models to one in which revisioned data is shared in a structured and controlled manner, customisations and transformations which are performed upon that data could be threatened. This is one issue to be aware of when interrogating both how data is used and transformed as well as where it goes next after a decision is made. By understanding and integrating these customisations, the shift towards a decision support system can become a benefit to teams and individuals who have historically made manual transformations and copies of data. It is also important to recognise that manually created independent copies of data introduce decision latency as they rely upon an individual or group to update data when new inputs become available.

In some cases, individuals or teams may perceive the introduction of transparent decision logic as an attempt to expose them to external scrutiny. In these cases it can be helpful to frame the decision support systems as generating supporting evidence that processes are being followed correctly and identifying any issues or errors stemming from external sources such as incorrect data, or

fundamental flaws in the standard operating procedures. This transition can also be seen as an opportunity to improve upon existing systems and procedures.

One further challenge which can be encountered is a perception that new automations will result in job losses. While it is true that automation is used to reduce the required work hours for a given task, industry observations indicate that the use case for decision support measures is to improve outcomes rather than reducing workforces. Often, the greatest value which experienced professionals provide to their employers lies in applying their knowledge and expertise to analysing and solving problems and not in executing routine procedural tasks.

Organisational

At an organisational level, ownership, approval for use and responsibility for execution of individual processes and standard operating procedures are usually well defined. Ownership of specific data and its use for those processes is frequently less clearly defined. To ensure that decision-makers are supported with the most up to date, correct source of data for any given decision, clear boundaries of data ownership are required. These boundaries refer to both the responsibility for creating and maintaining data sources and the responsibility for managing any changes to the processes which modify, validate and build upon them.

Mining organisations traditionally ensure that operational execution is gated behind signature based approvals processes to prevent unsafe or inefficient use of resources. This provides auditability of when decisions were made or changed as well as individual accountability. Moving away from signature based approvals processes can present a challenge from a change management perspective. Two main paths exist for managing approvals with a decision support system. Either a business may expend significant effort into robust decision support systems, such that the system itself can be signed off as approved for use, or include explicit human in the loop approvals systems which require an authenticated, approved individual to provide approval for a decision to be used in the field each time.

Professional opportunity

Beginning to implement decision support systems in a business provides several opportunities for experienced professionals to add value to their teams. Reducing time expended on data handover and repetitive execution unlocks time for analysis and process improvement efforts. By closely investigating data flows and transformation steps, multiple improvements can be made to further reduce double-handling of data.

Making business processes and decision logic transparent means that context can be gleaned by individual professionals about how their work depends upon inputs from others and where their work will impact on downstream processes. Decision transparency also offers opportunities for less experienced team members to better understand their role.

THE DATA FIRST MINE

With consistently growing quantities and sources of data over time, it is natural to focus on storage, security and accessibility of data. This is especially true as issues of capacity, speed and price tend to be directly linked to quantity of data. By putting too much focus on the storage and transfer of data, it is possible to miss the value of the data itself which lies in enabling analysis and process improvements. By focusing on the decisions and analysis and working backwards to determine the sources of data and transformations required to best support the decision-maker, improving valuable decisions becomes the primary objective.

A mental image that can assist with visualising decision support systems is treating the system as infrastructure. To produce good decisions, robust infrastructure must deliver the correct data in the correct formats to inform the decision-maker. Reducing decision latency and unnecessary processing and duplication of data are a matter of simplifying and optimising the infrastructure serving each decision. This optimisation can take several forms, including reusing the same data for multiple decisions, automating transformations and compute steps, preventing data transformations from executing when decisions are not required, setting automatic triggers for starting or revisiting

decisions, and ensuring that humans are notified when anomalies are encountered which require expert assessment.

By directly creating human in the loop decision processes, mining businesses can ensure that professional expertise is applied to the most valuable problems available at any given time while automation frees up additional time for valuable analysis and process improvement efforts. Making decision architecture transparent ensures that tacit knowledge is captured and can be readily accessed for onboarding of new team members.

CONCLUSIONS

Increasing mineral demand and complexity of deposits require more efficient methods of making decisions to support a constrained pool of mining professionals and graduates. While ongoing growth in data creation and storage volumes continues to be an important factor for mining companies, the primary challenge and differentiator is increasingly making the best use of the available data and compute resources.

Traditional digital transformation efforts often focus heavily on data generation and storage solutions. This paper proposes the opposite approach: begin with the decision, assess the options and appropriate triggers for each option, then identify the data sources, transformations and governance required to support the decision-maker.

A method of interrogating existing standard operating procedures and decision processes is proposed based upon the five questions:

1. Why is a given pathway followed? List possible pathways and appropriate triggers for each option.
2. Under what conditions should that decision change or be reviewed? Identify areas where escalation to human experts is required or continuous improvement efforts may be useful.
3. Where does this data originate? Eliminate unnecessary duplication and manual handling of data.
4. How is it used and transformed? Automate repeatable transformation and compute steps, reducing reliance on procedure-following.
5. Where does this information go next? Further expand automation and integration efforts to ensure that downstream decision-makers and operations teams receive appropriate outputs.

This paper proposes that a good decision support system requires three key components:

1. Data governance to ensure that data can be relied upon to be recorded consistently, reliably stored in the correct locations and accessible to all relevant downstream processes.
2. Repeatable compute in controlled environments to ensure that captured data is automatically validated, transformed and ready for analysis in the correct formats, and reduce delays cause by waiting on human personnel.
3. Transparent decision pathways to ensure that appropriate data and transformations are used consistently, decisions are readily viewed for auditing of processes and continuous improvement and professional expertise is made visible for onboarding and teaching purposes.

Human expertise is often most effectively applied to complex analytical tasks and comparing possible solutions rather than following repeatable steps in a process or manually transferring data. By reframing the purpose of mining data as supporting decisions and building data transfer and compute infrastructure to ensure that decision-makers receive validated, recent information each time a decision must be made, decision-makers are empowered to make better decisions.

ACKNOWLEDGEMENTS

The authors would like to acknowledge Christie Myburgh for productive and thoughtful discussions regarding decision support systems and for providing recommendations for relevant reading to inform the development of this paper.

Artificial intelligence tools were used during the preparation of this paper. The ChatGPT (GPT-5.5) large language model was used to assist with structural review, language and editorial suggestions and formatting guidance, as well as production of Figure 2. All technical content, interpretations, conclusions and final editorial decisions remain the responsibility of the authors.

REFERENCES

Australian Network of Mining Engineering Educators (ANMEE), 2025. 2024 Australian Graduate Mining Engineering Statistics Report [online]. Available from: <https://www.ausimm.com/globalassets/downloads/181225_appendix-1_ausimm-professional-pathways-whitepaper_final.pdf> [Accessed: 12 Jun 2026].

CCIWA, 2021. BHP's transformation future proofs for change [online]. Available from: <https://cciwa.com/business-toolbox/growth/bhps-transformation-future-proofs-for-change/> [Accessed: 12 Jun 2026].

Datausa.io, 2022. Mining and Mineral Engineering | Data USA [online]. Available from: <https://datausa.io/profile/cip/mining-mineral-engineering> [Accessed: 12 Jun 2026].

Kruger, K, Sparrow, D and Defty, T, 2023. An integration platform for smart, integrated mining, in Proceedings of the Southern African Institute of Mining and Metallurgy Digitalization in Mining Conference 2023.

Mackenzie, J and Moncrieff, J, nd. The mining data ecosystem: revolutionising modern mining, unpublished.

McLean, A, 2018. Rio Tinto preparing for the Mine of the Future with automation | ZDNet [online]. Available from: <https://www.zdnet.com/article/rio-tinto-preparing-for-the-mine-of-the-future-with-automation/> [Accessed: 12 Jun 2026].

AI-enabled strategic planning and why people are the mining industry's real asset

B Yaqub[1]

1. MAusIMM, Principal Engineer, Strategic Production Planning, Rio Tinto Copper, Perth WA 6000. Email: belinda.yaqub@riotinto.com

ABSTRACT

Artificial Intelligence (AI) is increasingly discussed in strategic mine planning for its ability to improve speed, automation and analytical reach. This focus is useful, but incomplete. Orebody value is not an outcome of computation alone. It depends on people linking fragmented technical, operational and financial evidence into a coherent decision pathway. As mining organisations invest in AI capability, they must also preserve the planning expertise needed to frame questions, interpret outputs and govern the decisions those outputs inform.

This paper presents a human-led view of AI in strategic mine planning through two recent case studies. The first shows how contextual scoping helps teams verbalise strategic context and define the planning problem before analysis begins. The second examines a development-stage resource project that considers potential value and future growth beyond an initial development. Together, the cases show that AI adds most value after people have made the problem visible, linked disciplines and applied judgement to define what the analysis should solve.

The paper proposes a practical framework for human-led AI in strategic planning. AI can accelerate analysis, surface relevant evidence, map linkages and improve comparison. It does not replace the human capability to frame the problem, own assumptions, challenge weak logic and judge what is operationally, socially and strategically acceptable. The paper concludes that AI is most valuable when embedded in expert-led systems that preserve judgement, cross-disciplinary challenge and accountability for the final recommendation.

INTRODUCTION

Artificial Intelligence (AI) is increasingly discussed in mining as a source of computational speed, automation and predictive capability (Corrigan and Ikonnikova, 2024; Codoceo-Contreras, Rybak and Hassall, 2024). Across the mining value chain, AI is being applied to data processing, pattern recognition, forecasting, optimisation and reporting. Much of this discussion is useful, but it tends to emphasise what AI can calculate rather than how planning decisions are made. In strategic mine planning, value does not emerge from a single model or isolated algorithm. It emerges when technical, operational, infrastructure, schedule and financial understanding are integrated into a defensible view of where value lies. AI can support this analysis, but mine planning strategy remains a human-led endeavour.

For this reason, the central proposition of this paper is that strategic planning first depends on people defining the right question. They must determine what problem is being solved, what matters and which trade-offs require judgement. AI can then strengthen the analysis, but only after people have framed the problem. Human expertise connects fragmented analyses into a coherent decision pathway. It also determines what is operationally, strategically and socially acceptable. This becomes especially evident in live technical dialogue. Experts speak across disciplines, recognise linkages, challenge assumptions and expose tensions in real time. That exchange remains fundamental to planning and is not easily reduced to a computational task (Kumar *et al,* 2024; Gomez *et al,* 2025).

Throughout this paper, judgement refers to the human capability to define what matters, frame the question being solved and determine which trade-offs are acceptable in context. It is not a substitute for evidence, but the discipline that gives evidence relevance.

This paper explores that proposition through two recent strategic planning case studies. The first shows how a contextual scoping workshop creates the human interaction needed to verbalise strategic context, align perspectives and define the planning problem before analysis begins. The

second examines a development-stage resource project that considers potential value and future growth beyond an initial development. Together, the cases show how AI can improve analytical reach while subject matter experts remain responsible for judgement, strategic coherence and decision credibility.

STRATEGIC PLANNING CASE STUDIES

The following case studies illustrate how strategic planning in mining depends on first defining what matters and what question is being solved, then integrating technical, operational, social and financial considerations to address it. Together, they show how people frame the problem, test assumptions and interpret trade-offs across complex planning environments, while AI strengthens the process without displacing that human role, as shown in Table 1. These examples then provide the basis for the human-led AI framework proposed in the next section.

TABLE 1

Comparison of the two strategic planning case studies.

Criterion	Case study 1 – Contextual scoping workshop	Case study 2 – Development-stage resource project
Purpose	Use collaborative scoping to verbalise strategic context and define the planning scope.	Use value framing to define the investment context and test expansion credibility.
Strategic question	What problem should the planning analysis solve before options are tested?	What factors should be considered when assessing future growth beyond an initial development?
Main planning tension	Convert competing perspectives, constraints and objectives into a shared planning scope.	Balance value, acceptability, optionality and execution readiness.
Human contribution	Surface tensions, translate across disciplines and define a credible analysis scope.	Frame growth logic, test constraints and judge whether expansion is credible.
AI contribution	Summarise inputs, identify themes, trace dependencies and structure planning questions.	Organise studies, benchmark options, distinguish readiness from optionality and improve traceability.
What AI cannot replace	Judge which questions, assumptions and disagreements matter before analysis begins.	Judge whether the evidence is strategically credible, not just analytically complete.
Link to framework	Shows why people must frame the problem before AI-supported analysis begins.	Shows why people must judge strategic credibility before AI-supported evidence informs growth.

Case study 1 – Contextual scoping to frame the strategic planning problem

Case study 1 reflects a contextual scoping workshop at a brownfields site where operational realities, objectives and constraints had evolved over time. The workshop brought teams together before detailed analysis began, giving participants an opportunity to surface issues, verbalise the strategic context and define the broader system view of a complex operating asset. Its purpose was not to select a preferred case, but to create alignment and define the problem the planning analysis needed to solve. These workshops matter because strategic planning is not simply an analytical task. Analysis provides evidence, but strategy emerges through structured dialogue between people who see the asset from different perspectives.

In these workshops the tension between perspectives was not abstract; it was present in the room. Technical specialists pushed for options that made sense from a resource, infrastructure and schedule logic. Asset owners and business leaders tested those options against capital discipline,

delivery risk and long-term value. Site leaders challenged whether the pathway was practically viable. Social, political and stakeholder considerations shaped what could credibly progress. No single perspective was wrong; each pulls the strategy in a different direction.

This is where strategic planning became human work. It was not a process of determining where the highest value would lie or accepting the technically strongest option in isolation. With every contribution, trade-offs surfaced and assumptions were challenged. The planning question was reframed into something the organisation could judge. Strategic planners held this process together by translating between disciplines, identifying where positions were genuinely in conflict and helping the group define an analysis scope that was both technically meaningful and organisationally credible.

AI is well positioned to summarise technical inputs, expose inconsistencies and trace dependencies across vast operational bounds, however, the decisive judgement remained human. Planners had to determine which assumptions were credible, which disagreements mattered and which questions the analysis should answer. In this case AI is well positioned to accelerate synthesis by summarising contributions and identifying themes. It could also make linkages visible and translate workshop discussion into structured planning questions. It could not replace people's responsibility to verbalise context, align perspectives and define the problem being solved.

Case study 2 – Navigating resource and project bounds, execution readiness and stakeholder acceptability

The second case concerns a development-stage resource project progressing through internal feasibility while also moving toward the later decision and agreement stages required for execution. Its strategic importance as a growth option in a future-facing commodity is not enough on its own; the project must show that growth can be value accretive, not merely larger in scale. The planning challenge is therefore sharper than in Case study 1: the evidence needed to support an initial development may differ from that needed to assess future growth across value, technical maturity and stakeholder acceptability.

The project sits at the intersection of value, capital discipline, technical maturity and stakeholder acceptability. Business owners must assess value, scale and risk across development options, while the project team must progress an executable investment case and preserve credible growth optionality. At the same time, local rights holders and other external parties are concerned not only with whether development proceeds, but with how it proceeds, including infrastructure location, water sourcing, landscape sensitivity, heritage considerations and future opportunity. When specialists verbalise these tensions and link issues across mining, processing, water, infrastructure, logistics and finance, the problem statement becomes clear, and work can begin on framing the possible answers.

AI is well positioned to support this environment by providing targeted background reading and organising prior study material in a more accessible way. It can help identify current technical and market expectations, benchmark alternative development scales, separate execution readiness from strategic optionality and improve traceability across assumptions. Yet, as in the first case, it cannot substitute for judgement. People remain central because they understand the physical system, approval pathways, cultural constraints, stakeholder relationships and the consequences of selecting an option that is technically possible but strategically weak. AI may help sharpen the case, but people remain the real asset because they understand what must be protected, what can be changed and what future the project is intended to create.

Across both cases, the value of AI sits less in replacing strategic judgement than in improving the speed, traceability and comparability of the information that informs it. In both cases, AI adds value only after people have made the problem visible, linked the relevant disciplines and applied judgement to define what the analysis should solve.

A HUMAN-LED AI FRAMEWORK FOR STRATEGIC PLANNING

These two case studies suggest that AI is most useful when embedded within a deliberately human-led strategic planning process that enables specialists to test assumptions, challenge logic and

integrate perspectives, consistent with broader thinking on human-in-the-loop systems (Kumar *et al,* 2024). The framework clarifies where AI can accelerate analysis, surface relevant evidence, map linkages and improve comparison while preserving human capability to frame the problem, own assumptions, challenge weak logic and judge what is operationally, socially and strategically acceptable (Figure 1).

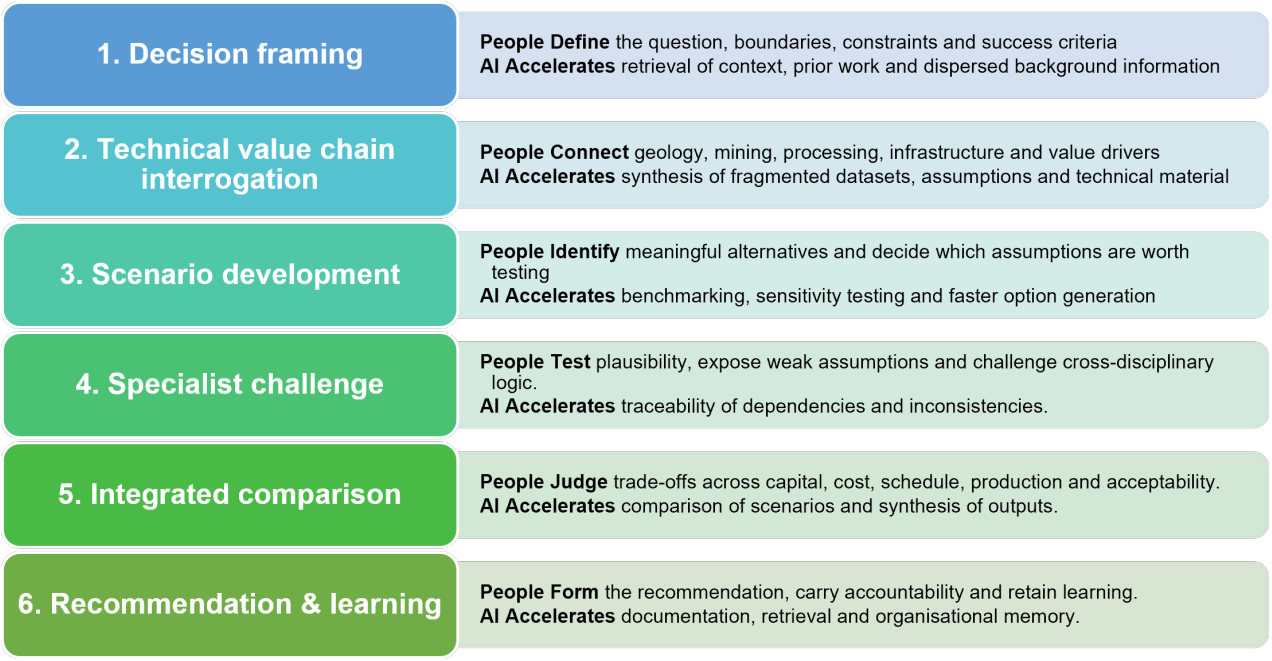


FIG 1 – Human-led AI framework for strategic mine planning.

Taken together, the framework shows that people and AI play complementary roles across the planning chain. People frame the problem, own assumptions, test plausibility, apply judgement and remain accountable for the recommendation. They must also verify AI outputs before use, using expertise to prompt effectively and improve reliability. Viewed as a team member with specific strengths and limitations, AI can help create stronger business processes when its capabilities are understood and governed by people.

The risk is that organisations rely on AI without preserving the capability needed to shape inputs and challenge outputs. A human-led AI framework protects decision quality by keeping faster analysis anchored in expert judgement.

CONCLUSION

These case studies and the proposed framework point to a broader implication for the mining industry. If AI adoption is discussed only in terms of speed, automation or model performance, then the most critical part of strategic planning is overlooked. Strategic decisions depend on integrating technical, economic and operational data that are often partial, uncertain and distributed across teams. The apparent maturity of an AI output does not remove the need for judgement or for understanding its context. In many cases, it may increase the need for informed challenge, because confident outputs can conceal weak assumptions or misplaced causal logic (Gomez *et al,* 2025).

This has practical implications for how mining organisations build capability. If technical knowledge is reduced while AI capability is expanded, organisations may unintentionally weaken the judgement needed to use AI well and sustain competitive advantage. Conversely, when AI is embedded within a human-led planning process, it can increase analytical reach, improve data management and create more room for high-value technical discussion. The challenge is therefore not *if* AI should be used in strategic planning, but *how* it should be utilised and governed so that expertise remains central and decision outcomes are improved.

AI has clear value in strategic mine planning, particularly when used across the analytical chain from orebody knowledge through to comparison of financial outcomes. However, its strongest contribution

is not the automation of judgement; it is the extension of human analytical reach. The evidence from these case studies suggests that people remain the integrating asset in strategic planning. They hold the knowledge to explain how an operation works, why it works that way and how it could work differently. For mining organisations seeking to adopt AI effectively, the priority should be to design human-led systems in which AI strengthens expertise rather than substituting for it.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the many teams involved in strategic planning across Rio Tinto, whose work and collaboration underpin the ideas reflected in this paper.

We also thank the many collaborators who contributed to the projects referenced here with their valuable technical knowledge and shared insight.

Finally, we acknowledge Rio Tinto for permission to publish this paper.

USE OF ARTIFICIAL INTELLIGENCE DISCLOSURE STATEMENT

In the preparation of this work, the author used Copilot to assist in the drafting and formatting of this paper. The author has reviewed and takes full responsibility for the paper's content.

REFERENCES

Codoceo-Contreras, L, Rybak, N and Hassall, M, 2024. Exploring the impacts of automation in the mining industry: A systematic review using natural language processing, *Mining Technology*, 133(3):191–213.

Corrigan, C C and Ikonnikova, S A, 2024. A review of the use of AI in the mining industry: Insights and ethical considerations for multi-objective optimization, *The Extractive Industries and Society*, 17:101440.

Gomez, C, Cho, S M, Ke, S, Huang, C-M and Unberath, M, 2025. Human-AI collaboration is not very collaborative yet: A taxonomy of interaction patterns in AI-assisted decision making from a systematic review, *Frontiers in Computer Science*, 6:1521066.

Kumar, S, Datta, S, Singh, V, Datta, D, Singh, S K and Sharma, R, 2024. Applications, challenges and future directions of human-in-the-loop learning, *IEEE Access*, 12:75735–75760.

Underground mine design and method selection

Transforming mine planning – auto-design rules for rapid underground layouts in Deswik and Datamine

A Saxena[1] and J R Town[2]

1. Principal Mining Engineer, AMC Consultants Pty Ltd, Melbourne Vic 3000, Australia. Email: asaxena@amcconsultants.com
2. Principal Mining Engineer, AMC Consultants (UK) Limited, Reading, Berkshire RG1 3EU, United Kingdom. Email: jtown@amcconsultants.com

ABSTRACT

Mine design is a fundamental component of mine planning, forming the backbone of every mining project from its inception through implementation and throughout the life of the mine. Accurate and efficient design is essential for evaluating project feasibility and ensuring operational success. Traditionally, underground mine design involves manual creation of strings and shapes within specialised software such as Deswik and Datamine. While these tools are powerful, the process of generating preliminary designs can be time-consuming, often delaying critical stages of mining studies and decision-making, and consuming resources that could be used elsewhere.

With increasing pressure to complete mining studies faster and test more strategic scenarios, there is a growing need for repeatable, transparent and auditable design workflows. Recognising this need, AMC Consultants (AMC) has leveraged the capabilities of Deswik and Datamine to develop automated design rules that streamline the underground mine design process. These rules enable engineers to rapidly generate preliminary layouts and shapes, reducing reliance on manual drafting and minimising repetitive tasks. The process can also be used to produce inputs into AMC's Hill of Value® strategic optimisation framework, reducing the effort involved in producing mine designs for multiple options or cut-offs, whilst still delivering maximum strategic value.

For these automated design rules to be reliable, each rule set must be supported by clearly defined input data requirements, design assumptions, parameter ranges, geometry constraints, output classifications and quality assurance (QA) checks. Required controls include validation of stope access connectivity, minimum gradient and turning-radius compliance, development-category reconciliation, exception reporting, version control of rule sets, documentation of manual edits, and senior engineering review before outputs are used for scheduling, cost estimation or study reporting.

The implementation of these automated rules has resulted in substantial efficiency gains. After initial rule set-up, individual inventory designs can be generated in minutes rather than days, while project-specific rule modification for advanced users typically requires only a few hours. This improvement accelerates project timelines and enhances the ability to conduct multiple design iterations, optimise layouts, and improve overall project evaluation. By integrating automation into established mine design platforms, AMC demonstrates how innovative approaches can transform traditional workflows, delivering faster, more consistent, and cost-effective outcomes for underground mining projects.

This paper explores the methodology behind developing automated design rules in Deswik and Datamine, highlights the practical benefits observed in real-world applications, and discusses the engineering controls required to use rule-based design outputs responsibly. The findings underscore the importance of transparent assumptions, repeatable workflows, validation checks and senior engineering review when applying automated design methods to underground mining studies.

INTRODUCTION

Underground mine design is a highly iterative, multidisciplinary process that underpins every stage of a mining project, from scoping and pre-feasibility studies through to detailed design, execution and life-of-mine planning. Design quality directly influences capital expenditure, operating cost, schedule, dilution, recovery and, ultimately, project value. Despite its strategic importance, preliminary underground design has historically remained a manual, drafting-intensive activity

performed by experienced engineers using specialist software such as Deswik.CAD (by Deswik Mining Consultants (Australia) Pty Ltd) and Datamine Studio UG (by Datamine Australia Pty Ltd).

Manual design workflows present several challenges. They are time-consuming, with a typical greenfield decline and stope layout requiring several hours to days and sometimes weeks of engineering effort. They are also difficult to reproduce, as design outcomes depend on the individual engineer's judgement, leading to variability between designs and between iterations. Furthermore, the manual nature of the process limits the number of design scenarios that can realistically be evaluated, constraining the ability for project teams to test sensitivities to cut-off grade, mining method, geotechnical conditions, and infrastructure placement.

Recent advances in scripting, parametric design and workflow automation have created an opportunity to make preliminary underground design more repeatable, auditable and efficient. Both Deswik and Datamine now provide tools that allow engineers to encode design assumptions, geometric constraints and engineering rules into controlled workflows — Deswik through rules-based design tools, and Datamine Studio UG through macro and process automation capabilities.

This paper presents AMC Consultants' (AMC) framework for automated underground mine design through three software case studies: two layouts generated using parametric rules in Deswik and one layout produced in Datamine Studio UG.

METHODOLOGY

The automated design workflow developed by AMC follows a five-stage process: (i) input data preparation; (ii) parametric rule definition; (iii) automated string and shape generation; (iv) validation and review; and (v) export of designs to scheduling and economic evaluation tools. The methodology is platform-agnostic in principle but has been implemented in two software environments — Deswik.CAD and Datamine Studio UG — to demonstrate transferability.

Reproducibility and workflow control

For the workflow to be reproducible, the project file should retain the input data sets, rule-set version, parameter table, design assumptions, output layers, validation checks and export settings used for each scenario. The recommended process is to prepare inputs, define project parameters, run the automated rules, export development and infrastructure quantities, complete quality assurance (QA) checks, document exceptions, and then transfer validated outputs into scheduling or economic-evaluation tools. Maintaining this record reduces reliance on individual judgement and allows designs to be re-run when cut-off grade, stope optimisation, level spacing or infrastructure assumptions change.

Input data requirements

Automated design rules require a consistent set of geological, geotechnical, and engineering inputs. The minimum data set comprises: a resource block model with grade and density attributes; a stope optimisation output (typically Deswik.SO or Datamine MSO) providing economic stope shapes; geotechnical domains defining maximum stope dimensions and pillar requirements; topographic and existing infrastructure surfaces; solids showing any areas already mined; and a set of design parameters covering decline gradient, drive dimensions, turning radius, level spacing, and access strategy.

Parametric rule definition

The core of the workflow is the parametric rule set. In Deswik, rules are defined using the Auto Designer toolkit, allowing engineers to encode standard mining geometry, for example a 1:7 gradient for horizontal development, a 5.5 m × 5.5 m drive profile and a 25 m turning radius, as reusable templates. In Datamine Studio UG, equivalent functionality is delivered through the Design and Process Automation modules, where rules are expressed as macros that operate on string and wireframe objects. In both cases, rules are version-controlled and parameterised so that a single rule set can be applied to multiple projects with only the input data changing. Similar parametric design approaches have been reported in the literature (Villa and Mikulich, 2024).

Automated string and shape generation

Once inputs and rules are defined, the software generates the design in a deterministic sequence: the mine access location is fixed, a decline or shaft access is established, level accesses are placed at defined intervals, level drives are extended to each stope, and stope shapes are imported from the optimisation output. Ventilation raises, escapeways, and ore and waste passes are inserted according to predefined spacing rules. A typical full-layout workflow for a 1500 m vertical extent layout completes in 2–5 days, compared with 1–2 weeks for an equivalent manual design. Creation of design strings in the Deswik case studies were completed within a few minutes, once the rule set was established.

Validation and review

Assessment criteria

To support consistent assessment of automated underground layouts, each generated design is evaluated against a defined set of engineering and study criteria. These criteria include compliance with project geometry constraints, completeness of access to stopes and infrastructure, practicality of decline and level layouts, suitability for scheduling and economic evaluation, repeatability of the workflow, and the level of manual refinement required before the design can be used in a study deliverable. The criteria provide a transparent basis for deciding whether an automated design is suitable for scoping, pre-feasibility, feasibility, or operational planning use.

Automated designs are not a substitute for engineering judgement. Each generated layout is reviewed by a senior mining engineer to confirm geotechnical compatibility, constructability, and alignment with the project's strategic objectives. The validation step typically identifies refinements that are fed back into the rule set, progressively improving the quality of automated outputs over successive iterations.

The validation process follows a structured quality-assurance framework rather than relying only on visual review. Checks include comparison of generated strings against the input Mineable Shape Optimizer (MSO) or stope optimisation shapes, reconciliation of total development metres by category, confirmation of decline gradient and minimum turning radius, review of level spacing and access-drive connectivity, assessment of stope access practicality, and confirmation that ventilation, escapeway, orepass, waste pass, and service infrastructure assumptions have been applied consistently. Outputs are also compared with conventional manual designs or benchmark layouts where available. Any discrepancy is classified as either an input-data issue, rule-logic issue, or project-specific engineering refinement, and the rule set is updated where appropriate.

Where automated design completeness or development proportions are reported, percentages are calculated on a development-metre basis by comparing the first-pass automated layout with the final engineering-reviewed study layout.

CASE STUDIES – DESWIK AUTO-DESIGN RULES

Two case studies are presented to illustrate the application of Deswik automated design rules to large underground deposits. In each case, the underlying rule set was drawn from AMC's internal Deswik design rule library, supplemented with project-specific parameters provided by the client.

Case Study 1 – Orebody A

Project context

Orebody A is a long orebody amenable to longhole open stoping (Figure 1).

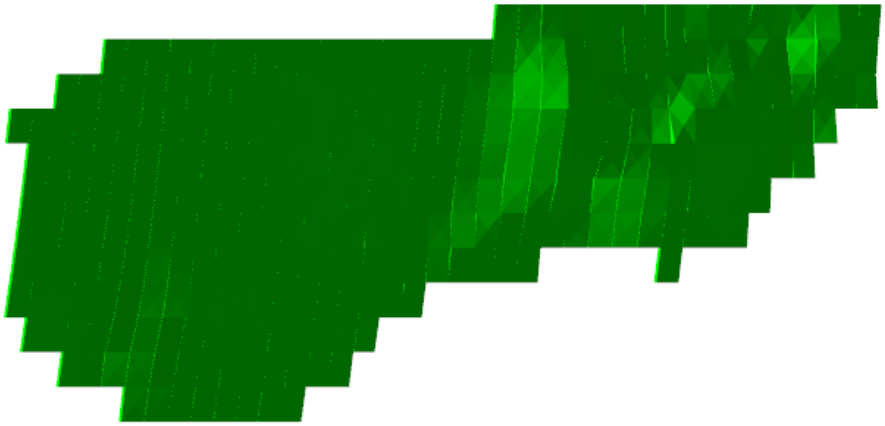


FIG 1 – Case Study 1 – View of orebody.

Auto-design rule set applied

After running MSO on the geological block model to create inventories at a range of parameters, a bespoke set of auto-development rules was developed to create mine designs for each inventory.

The process starts by extracting horizontal slices from the MSO stope shapes and extracting polylines that represent the centreline of each stope. These polylines are then connected via nearest-neighbour searches and extended to create stope access drives. The ends of the stope access drives are offset and the resultant line smoothed using B-spline resampling, and extended to required lengths, to create level access drives. Parallel offsets are created along the level access drives to create cross-cuts between declines and levels and inform ventilation design. Attribute calculations mark polylines for deletion, which are removed in final cleanup steps, while intermediate results are moved to designated output layers.

Set-up time for rules was approximately 12–16 hrs for initial rule set-up, depending on orebody complexity, input-data consistency and mining-method configuration. For advanced Deswik users, subsequent project-specific modification of an existing rule set typically requires approximately 2–4 hrs.

The workflow was implemented as a structured rule set to create the required auto-development layouts. The automated rule set used for Case Study 1 grouped the workflow into broad functional stages, including stope-slice preparation, access-string generation, filtering and cleanup, connectivity checks, smoothing, offset creation, extension of development strings, and final classification of outputs. The rule logic was configured to transform optimisation outputs into repeatable first-pass development layouts while retaining sufficient flexibility for project-specific review and refinement.

Results

The results of applying the auto-development rules and subsequent manual refinement to the different cut-off-based inventories are shown in Table 1 and summarised in Figure 2.

TABLE 1

Case Study 1 – Development kilometres per inventory.

Cut-off grade (%)	0.8	1.0	1.2	1.4	1.6	1.8	2.0	2.2	2.4
Infrastructure (km)	13.4	13.4	13.4	13.4	13.4	13.5	13.5	13.1	13.2
Access (km)	61.4	60.5	58.6	56.9	56.4	55.4	54.5	50.4	48.6
In-stope development (km)	32.0	31.1	29.9	28.9	28.1	27.0	26.2	24.6	22.8
Total development (km)	106.9	105.0	101.9	99.2	97.9	95.8	94.2	88.2	84.6
Ore tonnes per development metre (t/m)	234	211	207	201	188	173	164	163	157

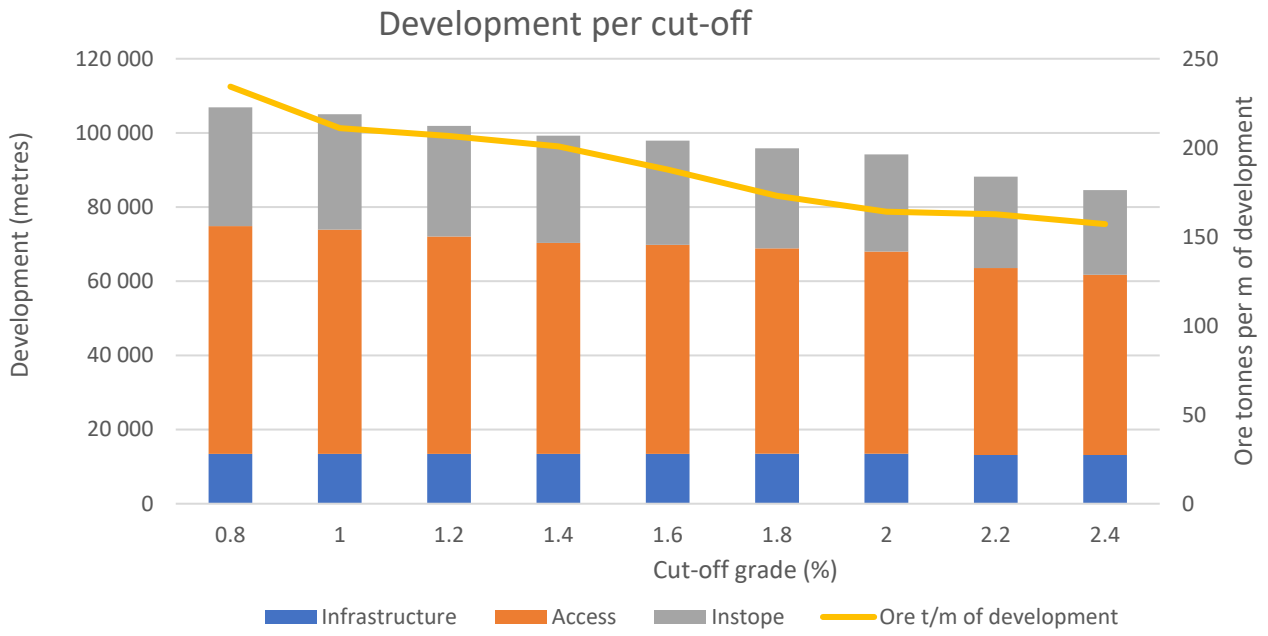


FIG 2 – Case Study 1 – Mining inventory (development and ore tonnes/metre).

Ore tonnes per metre of development vary across cut-off scenarios, indicating that extrapolation from a single inventory would be unreliable. Indicative ore tonnes per development metre range from approximately 234 ore tonnes per development metre (t/m) at the 0.8 per cent cut-off to 157 t/m at the 2.4 per cent cut-off (Table 1).

Figure 3 shows a visualisation of the automatically created development strings.

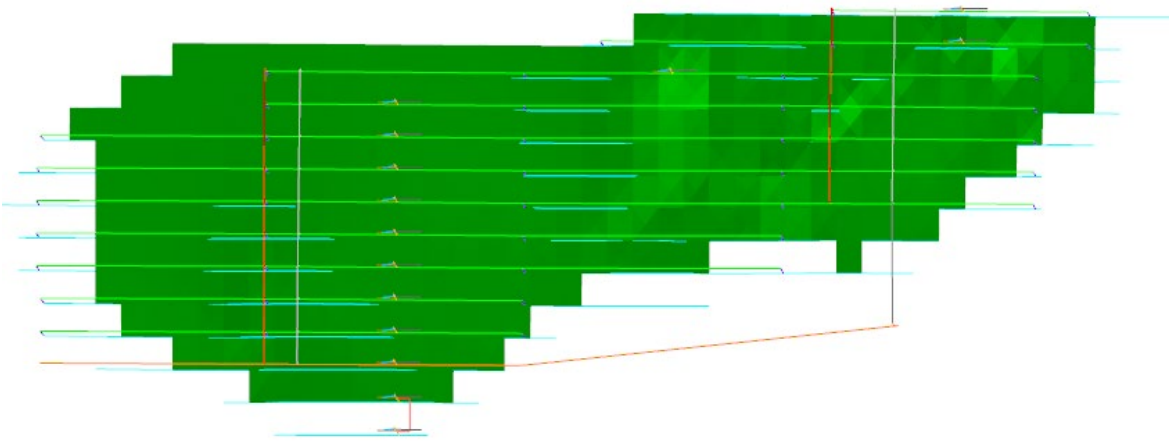


FIG 3 – Case Study 1 – Illustration of auto-development design.

Figure 4 shows a visualisation of the final refined design development strings.

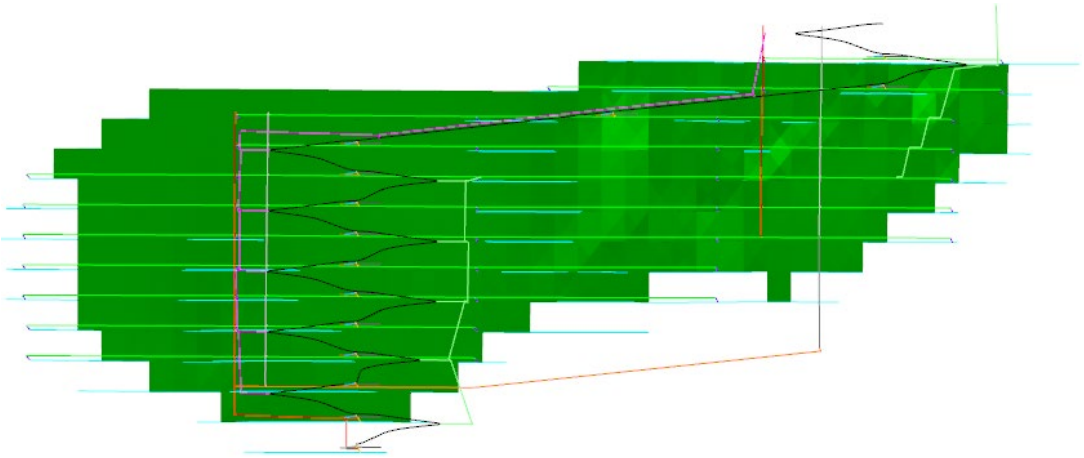


FIG 4 – Case Study 1 – Illustration of development design after manual refinement.

A comparison between the automated 0.8 per cent cut-off layout and the final engineering-reviewed design was completed to quantify design maturity and the extent of manual refinement required. The final design totalled 106.9 km of development, comprising 13.4 km of infrastructure development, 61.4 km of access development and 32.0 km of in-stope development. The automated design produced the majority of this layout, with subsequent manual refinement focused on decline and infrastructure connections, ventilation connections, level-to-level links and local engineering adjustments. Overall, manual additions and edits represented approximately 15–20 per cent of the final study layout.

On this basis, the automated workflow generated approximately 76 per cent of the final study design before manual refinement. When decline development is excluded, recognising that the decline remained largely dependent on manual engineering judgement and interactive design, the automated workflow captured approximately 86 per cent of the non-decline development layout. The validation outcomes before correction and after refinement are summarised in Table 2.

TABLE 2

Case Study 1 – Validation outcomes before correction and after refinement.

Validation measure	Initial automated output before correction	After refinement and engineering review
Stopes successfully accessed	Approximately 95–98% of stopes were accessed directly by generated level or stope-access development. Top-level development required manual design or additional rules after the first-pass automated layout.	Remaining stopes were connected where practical or flagged for redesign. Manual strings were created where required.
Disconnected/ overlapped strings	Approximately 1% of generated strings required connectivity/overlapping review.	Disconnected strings were corrected, removed or classified as non-schedulable exceptions.
Geometry-rule violations	Approximately 5% of strings required review against gradient, turning-radius, stand-off or access rules. The main issues related to ventilation connections and level-to-level connections.	Violations were corrected during engineering review or recorded as rule-set exceptions.
Exceptions requiring manual intervention	Exceptions were mainly associated with decline design, level tie-ins and local access interactions.	Decline and access connections were refined manually where project-specific strategy was not captured by the rule set.

Validation measure	Initial automated output before correction	After refinement and engineering review
Manual edits required	Manual edits represented approximately 15–20% of the final study layout. These edits focused mainly on decline design and connections, ventilation connections, level-to-level links and strategic access decisions not fully captured by the first-pass automated rules.	Edits were documented, included in the final design where required, and identified as candidate improvements for future rule-set updates where the changes were repeatable.
QA and review effort	Structured engineering QA and senior review were required before study use.	Typical QA and review effort was approximately 16 hrs.
Manual or benchmark comparison	An equivalent manual benchmark design would typically require approximately 5–10 engineering days for complex orebody.	For the 0.8 cut-off scenario, the automated output generated approximately 76% of the final 106.9 km study design before manual refinement. The main variation related to decline and infrastructure connections, representing approximately 13.4 km, and ventilation connections and other refinements, representing approximately 12.5 km. Excluding decline development, which remained a predominantly manual design component, the automated rules captured approximately 86% of the remaining final design.

Case Study 2 – Orebody B

Project context

The deposit is a large, base metal orebody amenable to longhole open stoping (Figure 5).

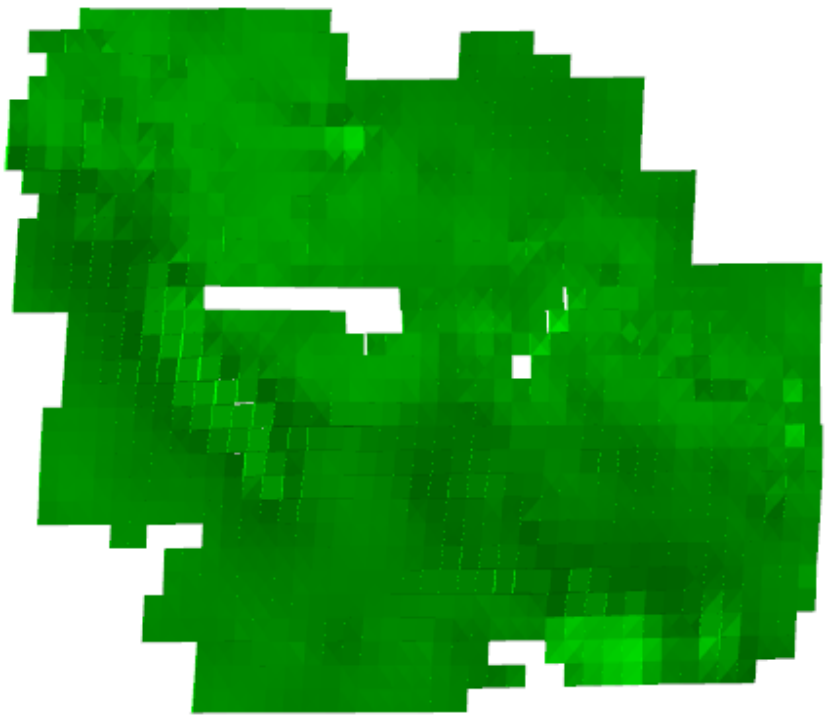


FIG 5 – Case Study 2 – View of orebody.

Auto-design rule set applied

Set-up time for rules was approximately 12–16 hrs for initial rule set-up, depending on orebody complexity, input-data consistency and whether both longitudinal and transverse stope access logic was required. For Deswik advanced users, subsequent project-specific modification of an existing rule set typically requires approximately 2–4 hrs.

The auto-design workflow follows the same broad structure as Case Study 1; however, additional logic was incorporated to distinguish between transverse and longitudinal stope access requirements. This increased the level of configuration required, but the methodology remained consistent: optimisation outputs were converted into repeatable access layouts through a controlled sequence of filtering, connectivity, smoothing, offsetting, extension and classification steps.

Results

The results of applying the auto-development rules and subsequent manual refinement to the different cut-off-based inventories are shown in Table 3 and summarised in Figure 6.

TABLE 3

Case Study 2 – Development kilometres per inventory.

Cut-off grade (%)	0.8	1.0	1.2	1.4	1.6	1.8	2.0	2.2	2.4
Infrastructure (km)	11.4	11.4	11.4	11.4	11.4	11.4	11.4	11.4	11.7
Access (km)	109.8	107.4	104.8	100.3	97.6	96.0	95.1	92.1	88.5
In-stope development (km)	51.0	49.5	47.8	45.7	44.3	42.8	41.4	40.1	38.5
Total development (km)	172.2	168.3	164.1	157.4	153.3	150.1	147.9	143.6	138.7
Ore tonnes per development metre (t/m)	296	291	285	281	271	261	251	246	102

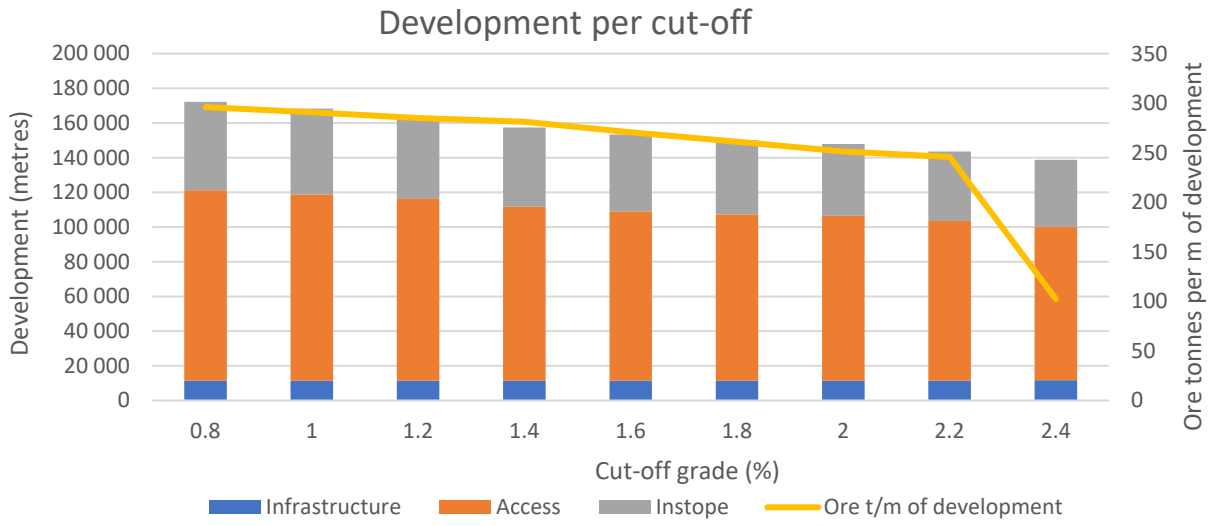


FIG 6 – Case Study 2 – Mining inventory (development and ore tonnes/m).

A comparison between the Case Study 2 automated layout and the final engineering-reviewed design was also completed to quantify the level of refinement required for a more complex orebody with transverse and longitudinal access requirements. The initial automated output totalled approximately 129.3 km of development, compared with approximately 172.2 km in the final refined study design. This indicates that the first-pass automated output represented approximately 75 per cent of the final engineering-reviewed development metres. The variance was mainly associated with infrastructure and decline development, including approximately 11.4 km of decline-related development, together with ventilation connections, local access rationalisation and addition or adjustment of development strings required in the final layout.

The comparison indicates that the automated workflow generated a conservative first-pass layout that intentionally prioritised access coverage and scenario completeness. Engineering review then optimised the layout to a more practical final design by refining decline connections, transverse-access tie-ins, ventilation links and local access development. After allowing for ventilation connections and other refinement items, the automated design represented approximately 80 per cent of the final practical layout. This result is useful because it demonstrates both the value and limitation of rule-based automation: the method can rapidly generate a comprehensive study layout, however senior engineering review remains necessary to complete infrastructure and refine non-preferred development before scheduling and costing.

Figure 7 shows a visualisation of the automatically created development strings.

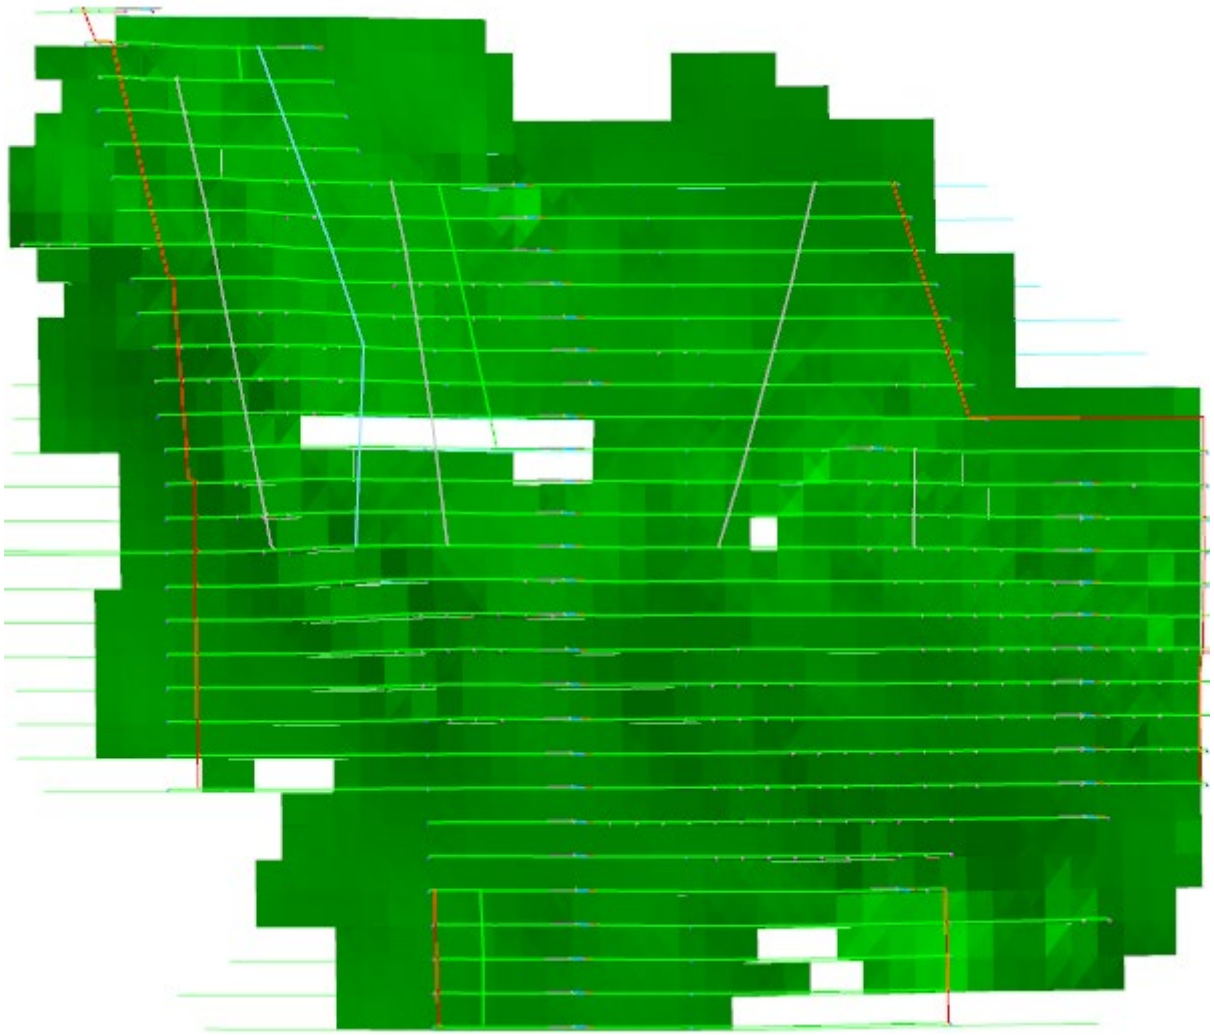


FIG 7 – Case Study 2 – Illustration of auto-development design.

Figure 8 shows a visualisation of the automatically created development strings after manual refinement.

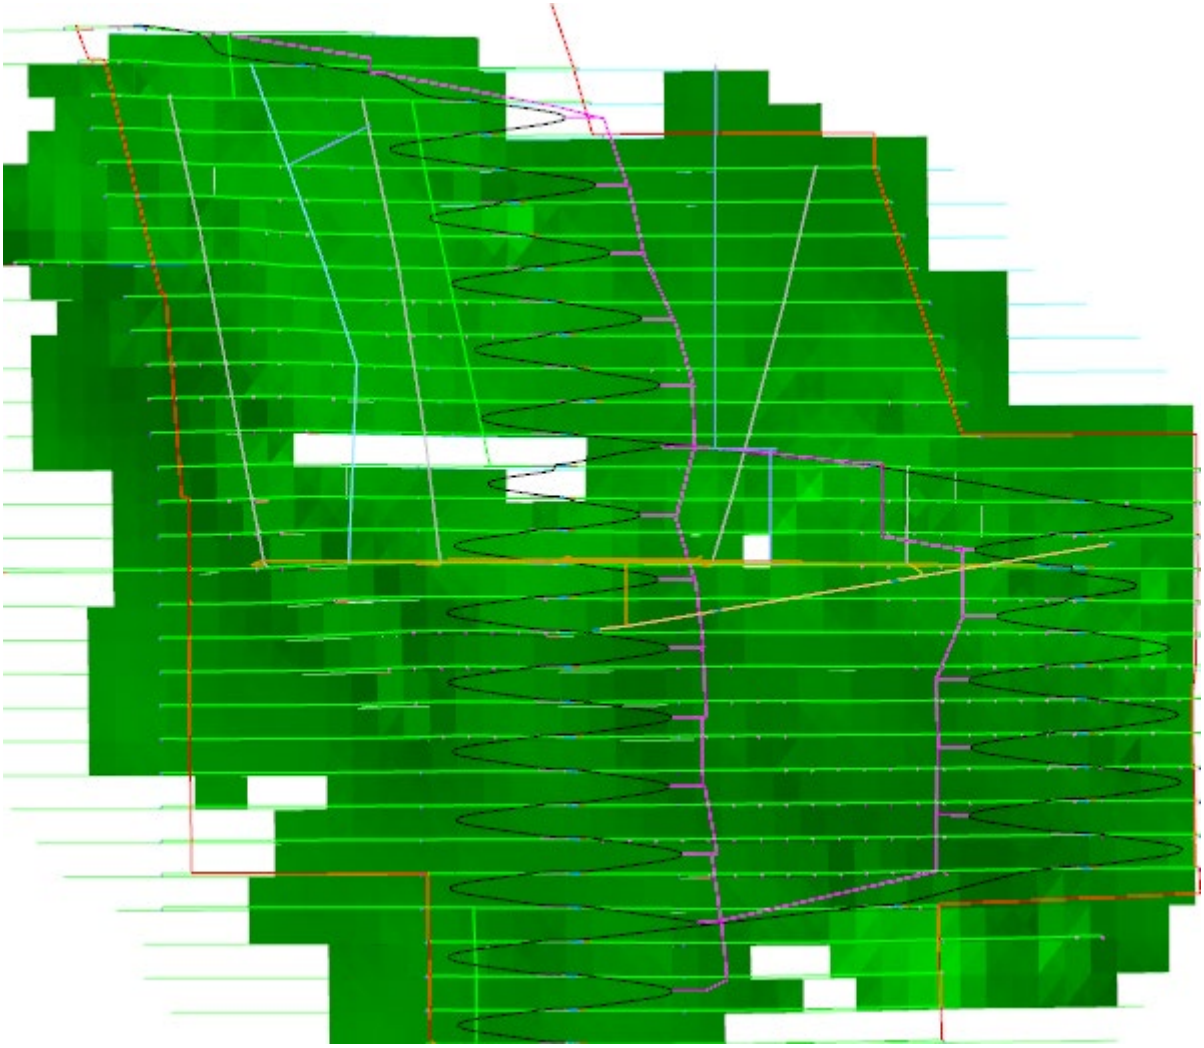


FIG 8 – Case Study 2 – Illustration of development design after manual refinement.

The validation outcomes before correction and after refinement are summarised in Table 4.

TABLE 4

Case Study 2 – Validation outcomes before correction and after refinement.

Validation measure	Initial automated output before correction	After refinement and engineering review
Stopes successfully accessed	Approximately 95–98% of stopes were accessed directly by generated level or stope-access development. Top-level development required manual design or additional rules after the first-pass automated layout.	Remaining access issues were reviewed and resolved where practical.
Disconnected strings	Approximately 2% of strings required connectivity review.	Disconnected strings were corrected, removed or classified as exceptions before export.
Geometry-rule violations	Approximately 6% of strings required checking against geometry constraints.	Violations were corrected, documented or incorporated into rule-set refinements.
Exceptions requiring manual intervention	Exceptions were mainly associated with decline design, transverse-access tie-ins and local access interactions.	Manual intervention refined the decline layout and strategic access connections.

Validation measure	Initial automated output before correction	After refinement and engineering review
Manual edits required	Manual refinement focused on decline connections, transverse-access tie-ins, ventilation links and rationalisation of local development. The initial automated output totalled approximately 129.3 km, compared with approximately 172.2 km in the final engineering-reviewed design, indicating that the first-pass automated output represented approximately 75% of the final development metres.	Repeatable edits were captured as candidate improvements to the rule set. The final design retained the required access and infrastructure logic while adding or adjusting development required for the preferred study layout.
QA and review effort	Structured engineering QA and senior review were required.	Typical QA and review effort was approximately 8–16 hrs.
Manual or benchmark comparison	An equivalent manual benchmark design would typically require approximately 5–15 engineering days.	The automated output generated approximately 129.3 km of development before refinement, compared with approximately 172.2 km in the final study design. This indicates that the first-pass automated output represented approximately 75% of the final engineering-reviewed design. After allowing for ventilation connections and other refinement items, the automated output effectively accounted for approximately 80% of the final design development metres. The remaining difference was primarily associated with decline and infrastructure connections, including approximately 11.4 km of decline-related development, and the addition or adjustment of local access and ventilation development.

CASE STUDIES – DATAMINE STUDIO UG AUTO-DESIGN

The following case study illustrates an equivalent automated design workflow implemented in Datamine Studio UG. The Datamine workflow uses the software's native Design and Process Automation tools, with macros encoding the same geometric and engineering rules used in the Deswik examples.

Case Study 3 – Orebody C

Project context

The deposit is a large, base metal orebody amenable to longhole open stoping, with a vertical extent of approximately 700 m and a strike length of 2400 m (Figure 9).

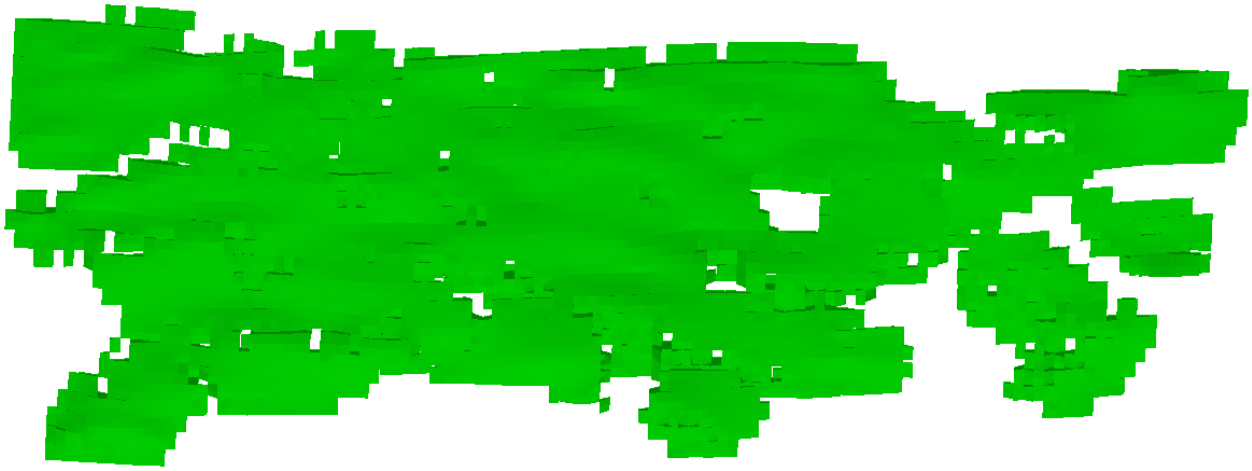


FIG 9 – Case Study 3 – View of orebody.

Auto-design rule set applied

Set-up time for rules was approximately 12–16 hrs for initial configuration, depending on orebody geometry, input-data consistency and mining-method requirements. For advanced Datamine users, subsequent project-specific modification of an existing rule set typically requires approximately 2–4 hrs. Each project still requires some adjustment to reflect orebody dip, strike direction, access strategy and mining method.

In Datamine Studio UG, the auto-design rules sit within the Planning workflow. The input to the rules comprises MSO shapes generated from cut-off grade or net smelter return (NSR) value, together with project-specific design parameters such as level spacing, drive offsets, standoff distances and access strategy. The following design sequence was created to generate the automated development layout.

Design sequence

Generalised Mine Design Sequence for Automated Design Rules:

- Ore Development:
  - Generate ore drives and ore cross-cuts from the MSO shapes.
- Primary Access Development:
  - Create footwall and/or hanging wall drives based on the specified standoff distance from the orebody contact.
- Stope Access Development:
  - Generate waste cross-cuts for transverse stope layouts.
- Ventilation Development:
  - Create ventilation cross-cuts extending from the footwall drive.
- Vertical Connectivity:
  - Establish level-to-level connections using raises located at a predefined distance from the footwall drive.
- Decline Development:
  - Generate the main decline using the Interactive Decline tool.
- Level Access Connections:
  - Create approach cross-cuts connecting the decline to each level’s footwall drive.
- Mine Infrastructure Development:

- Generate supporting infrastructure:
  - cuddies
  - sumps
  - service drive
  - ventilation drives
  - other required mine services and operational excavations.

The Datamine workflow used a configured sequence of automated design actions to generate and classify development strings. Input filters and output attributes were used to control where development was created, how outputs were categorised, and which strings were retained for downstream scheduling and infrastructure estimation.

Example

For example, approach cross-cuts were generated by identifying suitable connection locations between the decline and level access drives, then applying rule-based geometry to create practical level connections. The same workflow principles were applied across ore development, primary access, ventilation connections, vertical infrastructure, decline access and supporting mine infrastructure.

Results

The Datamine Case Study was completed for a single cut-off scenario rather than for a series of cut-off scenarios. Its purpose is therefore to demonstrate transferability of the automated design methodology to Datamine Studio UG, and to compare the first-pass auto-design output with the modified engineering-reviewed design used for sequencing, scheduling and infrastructure estimation.

Figure 10 shows an illustration of the automatically created development strings.

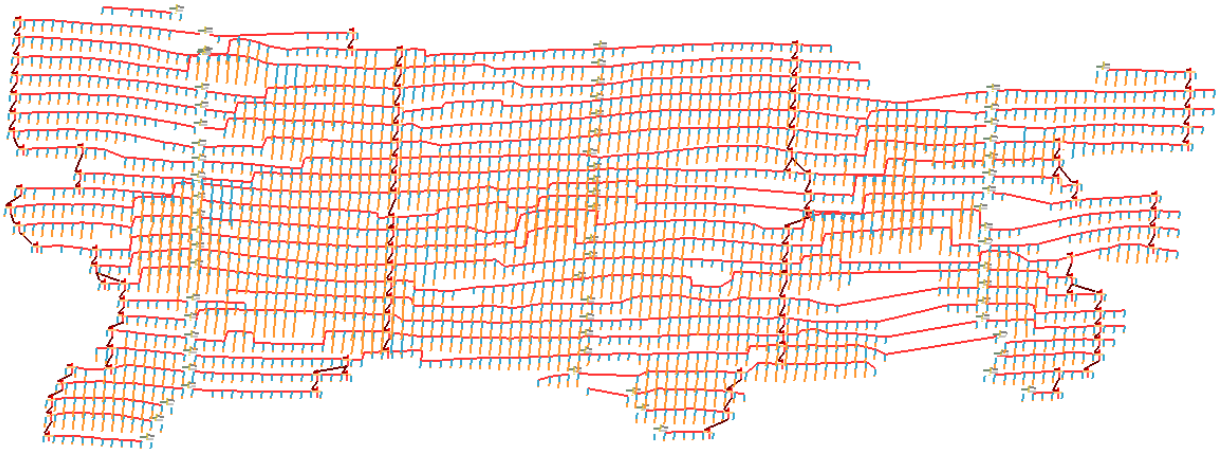


FIG 10 – Case Study 3 – Illustration of auto-development design.

Figure 11 shows an illustration of the development design after manual refinement.

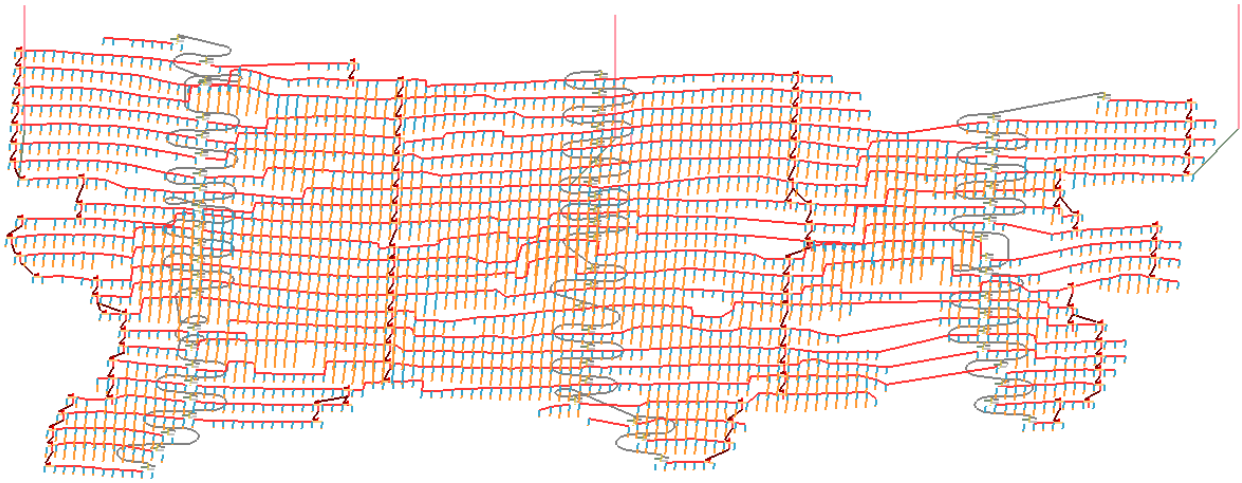


FIG 11 – Case Study 3 – Illustration of development design after manual refinement.

Table 5 compares the auto-design and modified auto-design development quantities for the 0.5 per cent cut-off scenario.

TABLE 5

Case Study 3 – Comparison of auto-design and modified auto-design development quantities for the 0.5% cut-off scenario.

Activity	Auto-design	Modified auto-design
Cut-off (%)	0.5	0.5
Ore development (km)	83.0	83.0
Primary access development (km)	47.3	61.7
Stope access development (km)	63.0	63.0
Ventilation development (km)	9.2	10.1
Vertical development (km)	2.6	2.6
Decline development (km)	0.0	11.4
Infrastructure development (km)	3.0	3.7
Total development (km)	208.1	235.5

The Case Study 3 comparison shows that the Datamine workflow generated a substantial first-pass layout, with subsequent manual refinement adding the decline, adjusting primary access, ventilation and infrastructure connections, and preparing the design for scheduling and infrastructure estimation. This single-scenario comparison supports the paper's transferability argument, while the Deswik examples demonstrate repeated cut-off-based inventory generation. The validation outcomes, based on development-metre and percentage changes from auto-design to modified design, are summarised in Table 6.

TABLE 6

Case Study 3 – Validation outcomes based on development-metre and percentage changes from auto-design to modified design.

Validation measure	Initial automated output before correction	After refinement and engineering review
Stopes successfully accessed	The first-pass auto-design generated 208.1 km of development for the 0.5% cut-off grade scenario.	The modified engineering-reviewed design totalled 235.5 km, an increase of 27.4 km or approximately 13% relative to the first-pass auto-design.
Primary access development	Primary access increased from 47.3 km in the auto-design to 61.7 km in the modified design.	Increase of 14.4 km, equivalent to approximately 30% additional primary access development after engineering review.
Decline development	The auto-design did not include decline development in the measured output.	The modified design included 11.4 km of decline development, representing approximately 42% of the total 27.4 km net increase between the auto-design and modified design.
Ventilation and infrastructure development	Ventilation increased from 9.2 km to 10.1 km, while infrastructure increased from 3.0 km to 3.7 km.	Ventilation increased by 0.9 km, or approximately 10%, and infrastructure increased by 0.7 km, or approximately 23%, reflecting refinement of service and connection layouts.
Unchanged development categories	Ore development, stope access development and vertical development were unchanged between the auto-design and modified design.	These categories retained 100% of the first-pass automated quantities, indicating that refinement was concentrated in decline, primary access, ventilation and infrastructure connections.
QA and review effort	The auto-design retained 208.1 km of development before modification.	Approximately 88% of the modified design was represented in the first-pass auto-design by development metres. The remaining 12% was added or refined through engineering review.
Overall interpretation	The Datamine auto-design output provided a substantial first-pass layout for the single 0.5% cut-off scenario.	The modified design shows that most ore, stope access and vertical development was retained, while engineering review added decline development and refined primary access, ventilation and infrastructure connections before scheduling and estimation.

ADVANTAGES AND LIMITATIONS

Auto-design tools are particularly useful for iterative mine-planning studies where designers need to create multiple scenarios for different cut-off grades, level-spacing assumptions, dilution and recovery factors, or infrastructure configurations.

The main limitations are associated with complex multi-lens geometries, local access interactions and scheduling constraints that require engineering judgement during sequencing and scheduling. In these cases, the automated output provides a strong first-pass layout but still requires manual review and refinement.

RESULTS AND DISCUSSION

Efficiency and comparison with manual design

Table 7 summarises the indicative time comparison between conventional manual design workflows and the automated workflows described in the case studies. The values are intended to illustrate the order-of-magnitude efficiency improvement observed in study work rather than prescribe a universal productivity benchmark.

TABLE 7

Decomposed manual versus automated design workflow time comparison.

Workflow stage	Deswik CS1 (Orebody A)	Deswik CS2 (Orebody B)	Datamine CS3 (Orebody C)
Initial rule set-up	~12–16 hrs	~12–16 hrs	~12–18 hrs
Project-specific modification	~2–4 hrs	~2–4 hrs	~4–8 hrs
Software execution (run time)	<5 min per inventory	<5 min per inventory	Minutes once inputs and rules configured
QA and validation	~16–24 hrs	~16–24 hrs	~16–24 hrs
Manual refinement	~15–20% of final layout (16 hrs)	~20–25% of final layout (18 hrs)	~12% of modified design (12 hrs)
Senior engineering review	8 hrs	8 hrs	8 hrs
Equivalent manual benchmark	>1–2 weeks	>1–2 weeks	>1–2 weeks

LIMITATIONS AND REQUIRED MANUAL REFINEMENT

The principal benefit of automation is not only speed but also repeatability and consistency. Automated outputs are typically suitable as high-quality first-pass study layouts when the input data are clean and the rule set is aligned with the selected mining method. Manual refinement is still required in areas where local engineering judgement is important, including portal or shaft placement, local decline tie-ins, geotechnical stand-off adjustments, ventilation-network practicality, emergency egress, materials-handling interfaces and detailed construction sequencing. The amount of refinement required should be recorded during each project so that the design team can quantify whether the rule set is improving over successive iterations.

Limitations by deposit geometry

Automated design rules are most effective where the orebody geometry and mining method can be represented by repeatable geometric logic. Performance is reduced where deposits contain complex multi-lens arrangements, abrupt dip or strike changes, fault offsets, highly irregular MSO shapes, inconsistent stope continuity, interacting access drives, or local infrastructure constraints that cannot be captured by simple rules. In these cases, the automated layout should be treated as a first-pass design requiring targeted engineering review rather than as a final design product.

Across the Deswik and Datamine case studies, automated design rules delivered consistent reductions in design time relative to conventional manual drafting workflows. The case studies show that repeatable rule-based design can generate preliminary underground layouts in minutes rather than days, while enabling rapid comparison of multiple cut-off or inventory scenarios. Table 8 summarises the automated mine design case study workflows and outcomes across the Deswik and Datamine case studies.

TABLE 8

Summary of automated mine design case study workflows and outcomes.

Workflow	Platform	Application	Rule set / method	Key outcome
Case Study 1	Deswik.CAD	Orebody A	100 automated development rules.	Rapid generation of multiple cut-off based development layouts, with each inventory design generated in under a minute.
Case Study 2	Deswik.CAD	Orebody B	180 rules. Generalised automated development workflow with additional transverse-access logic.	Automated design workflow extended to manage the additional complexity associated with transverse and longitudinal stope access requirements.
Case Study 3	Datamine Studio UG	Orebody C	Planning-based auto-design workflow using configured design-action categories.	Generated development strings suitable for sequencing, scheduling and infrastructure estimation.

Three observations emerge from the case study portfolio. First, time savings from automation increase with the scale and repetition of the design task, and are greatest for large orebodies with relatively low geological complexity, where a consistent rule set can be applied repeatedly across multiple levels, lenses and inventory cases. Conversely, geometrically complex deposits require more targeted manual refinement, which reduces the net time saving. Second, the quality of the automated design depends critically on the quality of the input data, particularly the stope-optimisation output and the consistency of strings, wireframes and design attributes. Third, automated designs are best treated as high-quality first-pass layouts that allow engineers to focus on judgement-intensive refinement, trade-off interpretation and validation rather than routine drafting. The Case Study 1 and Case Study 2 comparisons also show that automated design should be assessed by both run time and the proportion of first-pass development retained or refined in the final engineering-reviewed design.

STRATEGIC IMPLICATIONS

Rapid design generation improves the value of mine planning studies by allowing more scenarios to be tested within the same study period. This is particularly important for cut-off grade analysis, level-spacing selection, mining-method comparison, infrastructure intensity assessment and Hill of Value® evaluations. Rather than extrapolating development quantities from a single layout, the design team can generate scenario-specific development metres, access requirements and scheduling inputs. This improves the reliability of capital and operating cost estimates, supports more robust net present value (NPV) comparison, and helps identify options that would otherwise be excluded because of manual design effort.

CONCLUSIONS

Automated design rules in Deswik and Datamine Studio UG can materially reduce underground mine design time across large underground deposits and multiple mining-design scenarios. Once the initial rule set has been established, designs that would conventionally require several engineering days can be generated in seconds or minutes, enabling more rapid scenario testing and more consistent first-pass layouts. The three software case studies presented — two in Deswik and one in Datamine — demonstrate the transferability of the approach and provide a practical template that practitioners can adapt to their own projects.

Future work will focus on three areas: expansion of the rule library to cover additional mining methods, including block caving and sub-level caving; tighter integration between automated design outputs and the Hill of Value® strategic optimisation framework; and further development of validation controls, exception reporting and design-governance processes.

ACKNOWLEDGEMENTS

The authors acknowledge AMC Consultants for supporting the development of the automated design framework described in this paper and acknowledge the contributions of project teams across multiple offices in refining the rule sets and case study workflows. The authors also acknowledge Deswik and Datamine as the developers of the mine design software platforms (Deswik.CAD and Datamine Studio UG) used in this study.

REFERENCES

Villa, D and Mikulich, D, 2024. Parametric mine design applied to block caving, in *Proceedings MassMin 2024: International Conference and Exhibition on Mass Mining* (eds: D Johansson and H Schunnesson), pp 110–121 (Luleå University of Technology: Luleå).

Underground mining method selection for narrow veins at depth at Cuiabá Mine

R P S Silva[1] and M M Oliveira[2]

1. Long-Term Mine Planning, Technical Services Latin America, AngloGold Ashanti, Nova Lima MG, Brazil. Email: rpsilva@aga.gold
2. Professor, PhD, Federal University of Minas Gerais, Belo Horizonte MG, Brazil. Email: michelmelo@demin.ufmg.br

ABSTRACT

Mining method selection in mature underground operations should be periodically reassessed as geological, geotechnical and operational conditions change with depth. This paper presents a comparative study to support selection of the most appropriate underground mining method for narrow-vein mineralisation at depth at Cuiabá Mine, an underground gold operation in Brazil. The study was motivated by changes observed in deeper narrow-vein areas, including reduced thickness and lower dip, which have increased dilution, reduced recovery and challenged the long-term suitability of the current method.

Four mining scenarios were assessed within a fixed study area comprising ten mining panels: sublevel stoping with two sublevels and no crush pillars; sublevel stoping with three sublevels and no crush pillars; sublevel stoping with three sublevels and crush pillars; and drift and fill. The methodology followed an integrated mine planning workflow, including modifying factors based on historical performance, cut-off grade review, mineable shape optimisation, mine design, preliminary economic assessment, production sequencing and discounted cash flow modelling. The scenarios were compared in terms of run-of-mine tonnes, run-of-mine grade, recovered metal, mine life, revenue, costs, undiscounted profit and net present value.

The results show that the sublevel stoping scenario with three sublevels and crush pillars generated the highest net present value, at US$33.7 million, and is the most attractive option for short- to medium-term value maximisation under the assumptions adopted. In contrast, drift and fill maximised run-of-mine grade, recovered gold and undiscounted profit, reaching 8.09 g/t Au, 169.7 thousand ounces and US$91.6 million respectively, but delivered a lower net present value of US$31.4 million due to its longer mine life and the effect of a 10 per cent discount rate.

The study demonstrates that underground mining method selection for narrow veins at depth depends on the strategic objective of the business.

INTRODUCTION

Mining method selection is one of the most important decisions in underground mine planning. It defines not only the physical layout of the mine, but also the expected levels of dilution, mining recovery, productivity, development requirements, operational risk and economic value. Although the initial selection of a mining method is typically made during project studies, this decision should not be considered permanent. In mature underground operations, geological knowledge, geotechnical conditions, cost structures, commodity prices, operational performance and corporate priorities may change over time. As a result, mining methods should be periodically reassessed to confirm whether the current approach remains the most appropriate for the prevailing conditions.

Underground mining methods are commonly selected based on the geometry of the orebody, geomechanical conditions, grade distribution, access requirements, productivity targets and economic criteria (Hamrin, 1998). From a rock engineering perspective, excavation performance is also influenced by rock mass quality, stress conditions, structural controls and the interaction between the excavation geometry and surrounding ground conditions (Hoek, 2007). These aspects become increasingly important as underground operations advance to greater depths and face more constrained mining conditions.

This reassessment becomes particularly relevant in narrow-vein deposits. In this type of mineralisation, the economic performance of the operation is highly sensitive to excavation accuracy

and dilution control. Small deviations from the planned mining limits may introduce a significant amount of waste into the run-of-mine (ROM) material, reducing the average grade and increasing processing and handling costs. Stewart (2005) highlights that dilution and losses in narrow-vein mines are often controlled by equipment limitations, blasting performance, orebody geometry and hanging wall conditions. At the same time, excessive selectivity may reduce productivity, increase development intensity and extend the production schedule. The selection of a mining method for narrow veins therefore involves a fundamental trade-off between productivity and selectivity, with direct implications for mineral inventory utilisation and financial performance.

This paper presents a case study from an underground gold mine in Brazil, where narrow-vein areas are becoming increasingly challenging with depth. The operation is divided into major ore zones and narrow-vein zones. The major ore zones are characterised by higher grades and greater mineralised thickness, whereas the narrow-vein zones present lower grades and typical mineralised thicknesses of approximately 1.5 m to 3.0 m. The mine currently uses sublevel stoping as its main underground mining method, although cut-and-fill was previously used in the operation. As mining advances to deeper areas, changes in the geometry of the narrow-vein mineralisation have been observed, including a reduction in average thickness and a reduction in dip. These changes have contributed to higher operational dilution and lower mining recovery, raising the question of whether the current mining method remains the most suitable option for the deeper narrow-vein areas.

The geometry of the orebody is a critical factor in this assessment. Sublevel stoping is generally associated with higher productivity and lower unit mining costs, but it requires adequate continuity, sufficient dip and acceptable stope stability to perform effectively. In narrow and less steep orebodies, the method may become more exposed to dilution and recovery losses, particularly when the excavation geometry does not closely follow the mineralised envelope. Drift and fill, on the other hand, is a more selective method and may provide better control of dilution and ore recovery in geometrically constrained areas. However, this selectivity is usually achieved at the cost of lower productivity, higher development and backfill requirements, and longer production duration. The use of backfill is an important component of drift and fill mining, supporting the extraction sequence and contributing to local stability when mining is undertaken in a selective and sequential manner (Grobler *et al,* 2019).

The purpose of this study is to support the selection of the most appropriate underground mining method for narrow-vein mineralisation at depth at Cuiabá Mine. Four mining scenarios were assessed in a fixed study area comprising ten mining panels: sublevel stoping with two sublevels and no crush pillars; sublevel stoping with three sublevels and no crush pillars; sublevel stoping with three sublevels and crush pillars; and drift and fill. The comparison was designed to evaluate the technical and economic implications of each alternative, including dilution, mining recovery, productivity, mine life, recovered metal and net present value.

The evaluation followed an integrated mine planning workflow. Modifying factors were defined based on historical operating performance, including mining recovery, operational dilution and metallurgical recovery. Unit operating costs and economic assumptions were updated, and method-specific cut-off grades were reviewed. Stope optimisation was then carried out, followed by mine design, preliminary economic assessment, definition of activity precedencies, production scheduling at maximum practical capacity and discounted cash flow modelling. The scenarios were compared in terms of ROM tonnes, ROM grade, recovered metal, mine life, revenue, costs, undiscounted profit and net present value.

The main contribution of this paper is to demonstrate that underground mining method selection in mature narrow-vein operations should be supported by periodic quantitative reassessment. A method that maximises net present value over the short to medium term may not be the same method that maximises recovered metal and mineral inventory utilisation over a longer time frame. For this reason, mining method selection should be treated as a recurring strategic planning exercise, rather than as a fixed operational definition.

METHODOLOGY

The methodology was structured to compare underground mining alternatives within a fixed study area, ensuring that differences in results were primarily driven by mining method, design

configuration, modifying factors, production rate and sequencing strategy. The workflow followed a conventional mine planning process, starting with the definition of technical and economic assumptions, followed by mineable shape optimisation, mine design, preliminary economic assessment, production scheduling and discounted cash flow modelling. Similar integrated approaches are commonly used to connect stope boundary definition, mine design and production scheduling in underground mine planning studies (Little, Knights and Topal, 2013; Copland and Nehring, 2016).

Case study and study area

Cuiabá Mine is an underground gold operation located in Sabará, Minas Gerais, Brazil. The mine has been in operation for more than 30 years and is currently mined mainly by sublevel stoping, although cut-and-fill was previously used in the operation. The mine produces approximately 1.4 Mt of run-of-mine material per annum at an average grade of 5.5 g/t Au, corresponding to approximately 250 koz before mineral processing. Ore is transported to surface through a shaft, while waste is transported through both the shaft and the ramp system.

Cuiabá Mine is one of the deepest underground mining operations in Brazil, reaching approximately 1600 m depth. The production areas are divided into two main groups: major ore zones, including Fonte Grande Sul and Serrotinho, and narrow-vein zones, including Balancão, Canta Galo and Galinheiro, which are the focus of this study.

The location of Cuiabá Mine and the narrow-vein zones considered in this study are shown in Figure 1.

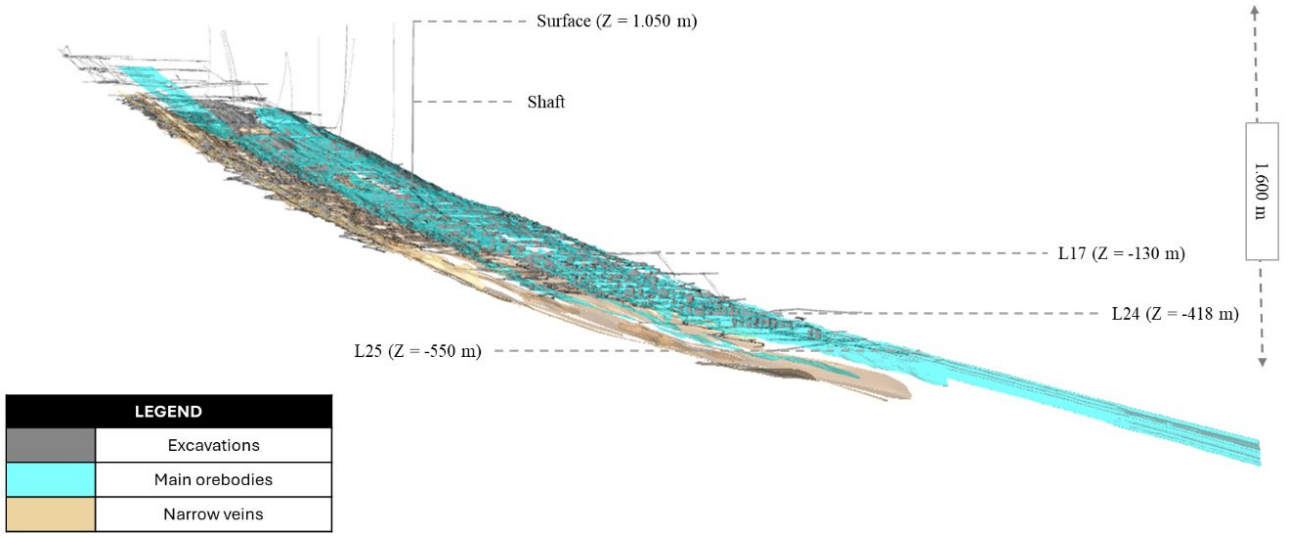


FIG 1 – Location of Cuiabá Mine and narrow-vein zones considered in the study.

Geologically, Cuiabá Mine is located in the eastern extension of Serra do Curral, on the north-eastern margin of the Quadrilátero Ferrífero, a region discussed in regional geological models by Farina *et al* (2016). The case study information, operational context and Scenario comparison presented in this paper are based on the author's MSc thesis (Silva, 2026).

The deposit is hosted in the Nova Lima Group, within the lower portion of the Rio das Velhas Supergroup, and gold mineralisation is predominantly associated with sulfide-rich banded iron formation and subordinate quartz veins hosted in schists. In the deeper areas of the mine, structural records indicate a reduction in the dip of the main mineralised structure, from approximately 42° in shallower levels to approximately 25° below level 18. This change affects key geometric assumptions for mine design and extraction planning, particularly in narrow-vein areas.

The study focused on the narrow-vein zones, where the mineralisation at depth has become thinner and less steep. These geometric changes have increased the exposure of the current mining method to operational dilution and mining recovery losses. For the comparison, a fixed study area comprising ten mining panels was selected. Keeping the same area for all scenarios allowed the assessment to

focus on the relative technical and economic performance of each mining alternative rather than on differences in inventory selection.

The selected panels were designed and evaluated under the same general geological and economic basis. This approach was adopted to provide a consistent comparison between the alternatives and to support a strategic assessment of the trade-off between productivity, selectivity, recovered metal, mine life and economic value.

Mining scenarios

Four mining scenarios were evaluated. The first three were based on sublevel stoping configurations, while the fourth Scenario considered drift and fill as a more selective alternative.

Scenario 1 considered sublevel stoping with two sublevels and no crush pillars. This configuration represents a conventional sublevel stoping approach for the study area, with a relatively simple design and a shorter vertical interval between production levels.

Scenario 2 considered sublevel stoping with three sublevels and no crush pillars. This alternative was designed to assess the impact of increasing the number of sublevels within the same panel configuration, with implications for development requirements, production sequence and stope geometry.

Scenario 3 considered sublevel stoping with three sublevels and crush pillars. This configuration was included to evaluate the effect of temporary pillars on mining recovery, dilution control, geotechnical exposure and economic performance. In this scenario, crush pillars were incorporated as part of the mining layout to support the operational and geotechnical strategy for the panels.

Scenario 4 considered drift and fill. This scenario was included as a more selective mining method for narrow-vein conditions. In comparison with sublevel stoping, drift and fill is expected to provide better adherence to the mineralised geometry and improved dilution control. However, it also requires a different operational cycle, including sequential drift extraction and backfilling, which may reduce productivity and extend the production schedule.

The four mining scenarios evaluated in the study are illustrated in Figure 2.

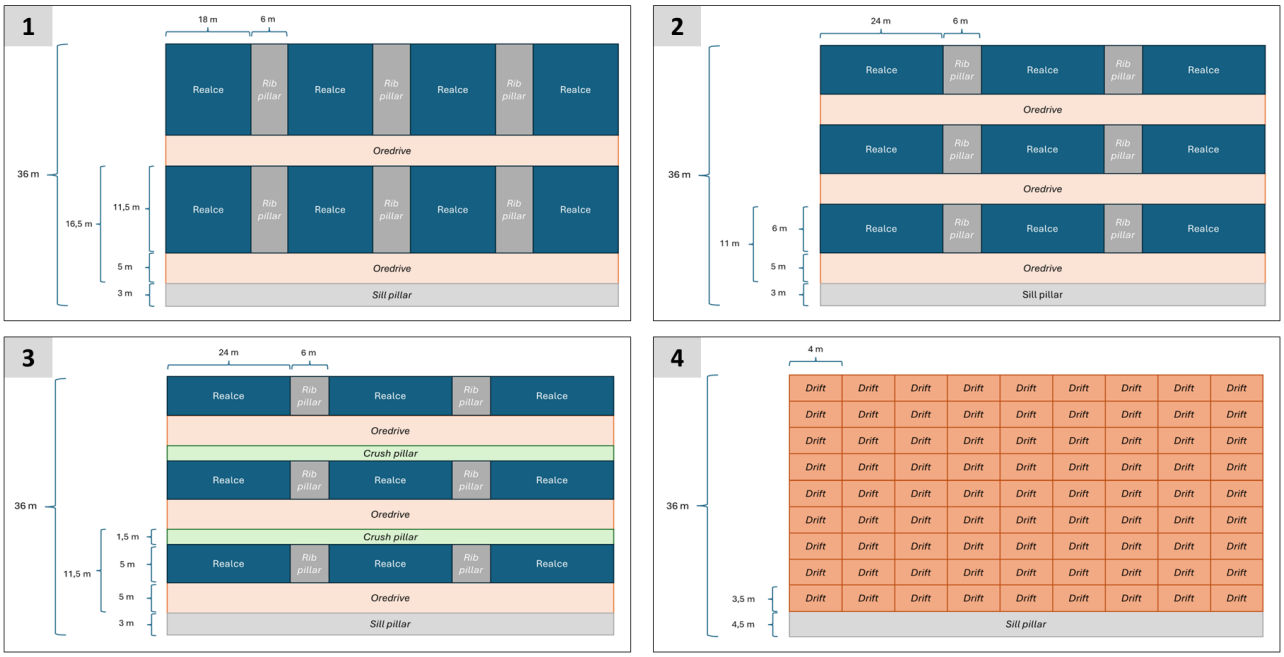


FIG 2 – Mining scenarios evaluated in the study.

For the sublevel stoping scenarios, sill pillars were defined according to geotechnical recommendations. For the drift and fill scenario, a more robust sill pillar was considered because no rib pillars were included between excavations and the voids were assumed to be filled with paste fill after mining.

Modifying factors and economic assumptions

The modifying factors were defined based on historical operating performance and applied consistently across the scenarios according to the selected mining method. The main factors considered were mining recovery, operational dilution and metallurgical recovery. Mining recovery was assumed to be 90 per cent for both sublevel stoping and drift and fill. Operational dilution was assumed to be 45 per cent for sublevel stoping without crush pillars, 30 per cent for sublevel stoping with crush pillars and 15 per cent for drift and fill. Metallurgical recovery was assumed to be 92 per cent for all scenarios.

Operating costs were derived from the cost base used in the production plan. A confidentiality factor was applied to the original costs, but the order of magnitude was maintained so that the comparative results would not be materially affected. The cost base included actual operating performance and forecast costs, providing an updated operational basis for the evaluation. Additional financial assumptions were obtained from internal corporate guidelines.

The cost structure considered mining, development, processing and other relevant economic assumptions. For the sublevel stoping scenarios, mining costs were applied on a US$/t basis for stopes, while primary and secondary development costs were estimated based on development length. For the drift and fill scenario, primary and secondary development were estimated using the same approach, whereas drifts and lateral cuts were costed on a US$/m basis according to the drift and fill unit cost.

After the modifying factors and economic assumptions were consolidated, method-specific cut-off grades were reviewed. This step was necessary because differences in dilution, recovery, unit costs and mining geometry affect the minimum grade required for each method to be economically viable. Cut-off grade is an economic decision variable that connects mining, processing, costs, prices and recovery assumptions, and therefore must be consistent with the mining method being evaluated (Lane, 1988; Rendu, 2013).

Method-specific cut-off grades were calculated using a break-even approach, considering the cost structure, gold price, metallurgical recovery, mining recovery and operational dilution applicable to each mining alternative. Different cut-off grades were required because each method presents a different balance between cost and selectivity. Sublevel stoping provides higher productivity and lower unit mining cost, but is more exposed to dilution, whereas drift and fill is more selective and has lower dilution, but requires higher unit costs and greater development intensity. The cut-off grades adopted therefore reflect the economic threshold required for each method to generate value under its respective technical and operating assumptions.

The main input assumptions used to calculate the cut-off grades are summarised in Table 1.

TABLE 1

Cut-off grade input assumptions.

Item	Unit	Value
Sublevel stoping mining cost	US$/t	98
Ore processing cost	US$/t	78
Primary development cost	US$/m	7518
Secondary development cost	US$/m	6981
Drift and fill cost	US$/m	5203
Gold price	US$/oz	1700
Selling cost	US$/oz	25.5

Note: the cost values were treated using a confidentiality factor and are therefore not disclosed in this table.

The main technical assumptions adopted for each scenario are summarised in Table 2.

TABLE 2

Main technical assumptions used in the study.

Item	Scenario 1	Scenario 2	Scenario 3	Scenario 4
Mining method	Sublevel stoping	Sublevel stoping	Sublevel stoping	Drift and fill
Sublevels	2	3	3	Not applicable
Crush pillars	No	No	Yes	No
Mining recovery (%)	90	90	90	90
Operational dilution (%)	45	45	30	15
Metallurgical rec. (%)	92	92	92	92
Paste fill requirement	No	No	No	Yes
Discount rate (%)	10	10	10	10

Mine planning workflow

The mine planning workflow was developed to translate the geological inventory into mineable and schedulable shapes for each scenario. The first planning step was mineable shape optimisation. The objective of this stage was to define mining envelopes consistent with the geometry of the mineralised zones, the selected mining method and the economic assumptions adopted for the study.

For the sublevel stoping scenarios, optimisation parameters were defined to represent the geometry of stopes and pillars, including minimum and maximum mining widths, stope height, pillar configuration and minimum wall angles. The optimisation process generated preliminary mineable shapes, which were then reviewed and adjusted during the mine design stage. For the drift and fill scenario, the design logic was based on selective extraction by drifts, considering the operational cycle and the need for backfill after mining.

Following optimisation, detailed mine designs were prepared for each scenario using underground mine design tools. The designs included mining shapes, access development, production drives and supporting development required to make each area operationally accessible. The objective was not only to compare the mineral inventory captured by each method, but also to account for the development and operational requirements associated with each mining alternative.

A preliminary economic assessment was carried out during the design stage. This assessment considered the value of the mineable inventory and the associated development requirements. Areas that did not meet the economic criteria were identified and excluded or adjusted, ensuring that the final designs represented technically and economically coherent alternatives.

After the mine designs were completed, logical precedencies were defined between activities. These precedencies represented the required sequence between development, access, production and, where applicable, backfilling activities. The precedence logic was used to create a realistic production schedule for each scenario and to avoid sequencing assumptions that would not be operationally achievable.

Production scheduling was carried out using a mine scheduling platform and at maximum practical capacity for each scenario. This approach was adopted to test the productive potential of each method and to evaluate how quickly each alternative could convert the available inventory into recovered metal and cash flow. Production rates were defined by mining method, reflecting the different operating cycles and productivity constraints of sublevel stoping and drift and fill.

Financial modelling and comparison criteria

The final step consisted of building a discounted cash flow model for each scenario. The financial model incorporated mined tonnes, ROM grade, recovered metal, revenue, mining costs,

development costs, processing costs and other relevant economic assumptions. The comparison focused on both undiscounted and discounted indicators to capture the distinction between total economic margin and value generated over time.

The main comparison criteria were ROM tonnes, ROM grade, recovered metal, mine life, revenue, total cost, undiscounted profit and net present value. ROM tonnes and ROM grade were used to assess the impact of each method on dilution and selectivity. Recovered metal was used to evaluate mineral inventory utilisation. Mine life and annual production profile were used to assess productivity and timing. Undiscounted profit was used to compare total economic margin, while net present value was used to incorporate the effect of time and the discount rate on project value.

This methodology allowed the mining methods to be evaluated not only as operational alternatives, but also as strategic planning options. The comparison therefore considered whether each Scenario was better aligned with short- to medium-term value maximisation or with long-term mineral inventory utilisation.

RESULTS AND DISCUSSION

The comparison between the four mining scenarios provides the technical and economic basis for underground mining method selection in the studied narrow-vein area at Cuiabá Mine. The results indicate different responses in terms of ROM tonnes, ROM grade, recovered metal, development requirements, mine life, undiscounted profit and net present value (NPV). This reinforces that selecting a mining method for narrow veins at depth should not be based on a single indicator, but rather on an integrated assessment of productivity, selectivity, dilution control, mineral inventory utilisation and timing of cash flow.

The consolidated operational and economic results for the scenarios are presented in Table 3.

TABLE 3

Summary of results by scenario.

Item	Scenario 1	Scenario 2	Scenario 3	Scenario 4
ROM tonnes (kt)	416	740	929	709
ROM grade (g/t Au)	6.11	5.82	5.70	8.09
Mined metal (koz)	81.8	138.3	170.3	184.4
Recovered metal (koz)	75.3	127.3	156.7	169.7
Total development (m)	4 742	8 280	9 710	24 687
Gross revenue (US$ million)	128.0	216.3	266.4	288.5
Total cost (US$ million)	-95.8	-162.6	-197.5	-196.9
Undiscounted profit (US$ million)	32.2	53.7	68.8	91.6
NPV (US$ million)	12.9	26.0	33.7	31.4
Mine life (years)	3	4	4	9

The first relevant observation is the difference between the sublevel stoping scenarios and the drift and fill scenario in terms of grade response. The three sublevel stoping scenarios delivered average ROM grades between 5.70 g/t Au and 6.11 g/t Au. Scenario 4, by contrast, achieved a ROM grade of 8.09 g/t Au. This result indicates a higher degree of selectivity and improved dilution control for drift and fill when compared with the sublevel stoping alternatives. In narrow-vein conditions, this is an important technical outcome, as dilution has a direct impact on the grade delivered to the processing plant, the cost per recovered ounce and the economic value generated from a fixed mineral inventory.

However, this improved selectivity is not achieved without operational trade-offs. The drift and fill scenario required the highest development metres, with 24 687 m of total development, compared

with 9710 m for Scenario 3, 8280 m for Scenario 2 and 4742 m for Scenario 1. This increase reflects the more selective and sequential nature of the method, which requires closer access to the orebody, individual drift extraction and backfilling activities. While this design logic improves the ability to follow the mineralised envelope, it also increases operational complexity, development exposure and the duration required to extract the available inventory.

The results also show that the change from two sublevels to three sublevels had a material impact on the performance of the sublevel stoping alternatives. Scenario 2 delivered a significant increase in ROM tonnes, recovered metal and economic value when compared with Scenario 1. This suggests that, under the assumptions adopted in the study, the three-sublevel configuration improved the ability to capture additional mineable inventory and generate value from the fixed study area. The increase in development requirements was compensated by the higher production potential and improved economic return.

The comparison between Scenario 2 and Scenario 3 highlights the effect of introducing crush pillars into the three-sublevel sublevel stoping configuration. Scenario 3 delivered the highest ROM tonnes, at 929 kt, and the highest NPV, at US$33.7 million. Although the average ROM grade in Scenario 3 was the lowest of the four scenarios, at 5.70 g/t Au, the greater tonnage and improved timing of production allowed this alternative to outperform the other scenarios in present value terms. This indicates that, for the assumptions adopted, the combination of increased production volume, operational configuration and shorter production duration outweighed the grade penalty associated with the method.

From a production perspective, Scenario 3 represents the strongest alternative among the sublevel stoping options. It provides the largest ROM tonnage and a substantial recovered metal content within a four-year schedule. This shorter duration is particularly important when discounted cash flow is considered. Earlier production generates earlier revenue, reducing the negative impact of discounting and increasing present value. In contrast, Scenario 4 generated the highest recovered metal but required a nine-year schedule, which materially affected its discounted economic performance.

The NPV comparison between scenarios is shown in Figure 3.

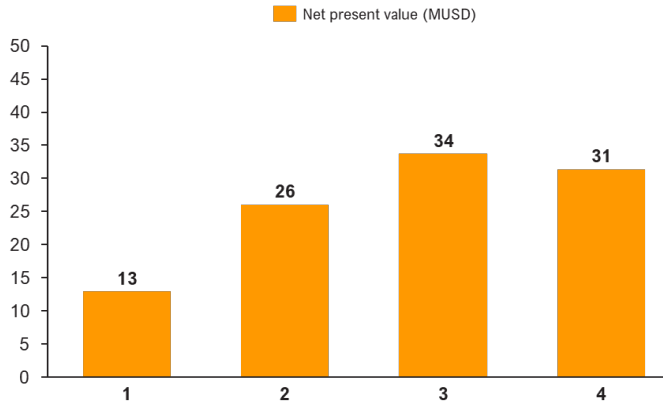


FIG 3 – Net present value comparison by scenario.

The economic results reinforce this interpretation. Gross revenue increased progressively from Scenario 1 to Scenario 4, ranging from US$128.0 million in Scenario 1 to US$288.5 million in Scenario 4. Total cost also increased from US$95.8 million in Scenario 1 to US$196.9 million in Scenario 4. Despite the higher development intensity and longer operating period, Scenario 4 generated the highest undiscounted profit, at US$91.6 million. Scenario 3 generated a lower undiscounted profit, at US$68.8 million, but achieved the highest NPV due to the shorter production schedule and earlier cash flow generation.

This distinction between undiscounted profit and NPV is central to the strategic interpretation of the results. Undiscounted profit measures the total margin generated over the full production horizon, without considering the timing of cash flows. On this basis, Scenario 4 is the best-performing alternative because it maximises grade, recovered metal and total economic margin. NPV, however,

incorporates the time value of money. Under the 10 per cent discount rate used in the study, the longer mine life of Scenario 4 reduces the present value of its future cash flows. As a result, although Scenario 4 generates the highest total profit, its NPV of US$31.4 million is lower than the US$33.7 million achieved by Scenario 3.

The difference between Scenario 3 and Scenario 4 shows that mining method selection is not a simple technical ranking. Scenario 3 is more aligned with short- to medium-term value maximisation because it delivers the highest NPV, a shorter mine life and a production profile that converts the mineral inventory into cash flow more rapidly. Scenario 4, although not the highest NPV alternative, provides stronger selectivity and mineral inventory utilisation, which makes it relevant for longer-term strategic planning.

Scenario 4, on the other hand, is more aligned with long-term mineral inventory utilisation. It delivers the highest ROM grade, the highest recovered metal and the highest undiscounted profit. These characteristics may be preferable when the corporate objective is to maximise total extraction, extend the productive contribution of the narrow-vein areas and improve resource utilisation over a longer time frame. In this context, drift and fill may provide a more sustainable extraction strategy from the perspective of inventory recovery, even if it is not the highest NPV alternative under the current productivity and cost assumptions.

The results also demonstrate the sensitivity of mining method selection to the interaction between geometry and scheduling. In narrow-vein mining, a more selective method may capture a higher-quality inventory and reduce dilution, but it may also slow down the production rate. Conversely, a more productive method may accelerate cash flow generation but may capture the orebody with lower selectivity. The value of the mine plan therefore depends not only on the total amount of metal recovered, but also on when that metal is recovered and at what cost.

The comparison between recovered metal and ROM grade by scenario is shown in Figure 4.

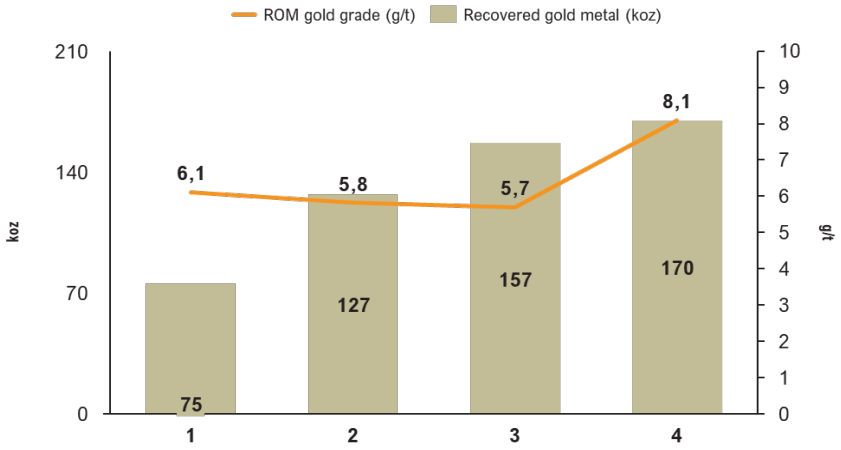


FIG 4 – Recovered metal and ROM grade comparison by scenario.

This is particularly relevant in mature underground operations where geological conditions evolve with depth. A mining method that performed adequately in shallower or more favourable geometries may not remain the optimal choice as the orebody becomes thinner, less steep or more operationally constrained. The results from this study support the need to periodically reassess mining methods using updated geological, operational and economic assumptions. This reassessment should include both technical indicators, such as dilution, recovery, development metres and production rates, and financial indicators, such as undiscounted profit and NPV.

Another important outcome is that the preferred scenario may change depending on corporate strategy. If the business objective is to maximise economic value in the short to medium term, Scenario 3 is the preferred alternative under the assumptions adopted. If the business objective is to maximise recovered metal, extend the productive life of the narrow-vein areas and improve long-term mineral inventory utilisation, Scenario 4 becomes strategically relevant and should remain under evaluation. This distinction is important because the best technical alternative is not

necessarily the best economic alternative, and the best economic alternative may depend on the time horizon considered.

The current results therefore support Scenario 3 as the most attractive option for implementation under the present assumptions, particularly because it provides the highest NPV and a shorter production period. Nevertheless, Scenario 4 should not be discarded. Its superior grade, recovered metal and undiscounted profit indicate that drift and fill may become more competitive if productivity is improved, development intensity is reduced, unit costs are optimised or if the corporate strategy places greater emphasis on long-term mineral inventory recovery. Further work should therefore focus on testing operational improvements for drift and fill, including productivity assumptions, backfill cycle efficiency, development optimisation and sensitivity to dilution and cost parameters.

Overall, the results support the selection of Scenario 3, sublevel stoping with three sublevels and crush pillars, as the preferred mining method under the current assumptions for the studied narrow-vein area at Cuiabá Mine. This recommendation is based on its higher NPV, shorter production duration and stronger alignment with short- to medium-term value maximisation. However, Scenario 4, drift and fill, should remain under evaluation because it maximises ROM grade, recovered metal and undiscounted profit, indicating potential strategic value if productivity, development intensity or cost assumptions can be improved.

CONCLUSIONS

This paper presented an underground mining method selection study for narrow-vein mineralisation at depth at Cuiabá Mine. Three sublevel stoping configurations and one drift and fill scenario were compared within a fixed study area, using an integrated mine planning workflow that included modifying factors, method-specific cut-off grades, mineable shape optimisation, mine design, production scheduling and discounted cash flow modelling.

The results show that mining method selection for narrow veins at depth cannot be based on a single technical or economic indicator. Each scenario presented a different balance between productivity, selectivity, dilution control, recovered metal, mine life and economic value. Under the assumptions adopted in this study, the sublevel stoping scenario with three sublevels and crush pillars generated the highest net present value and is therefore the preferred mining method for the studied area when the objective is to maximise value in the short to medium term.

The drift and fill scenario, however, delivered the highest run-of-mine grade, recovered gold and undiscounted profit. This indicates that a more selective mining method may provide better mineral inventory utilisation and stronger long-term extraction performance in geometrically constrained narrow-vein areas. Its lower net present value was mainly associated with the longer production schedule and the effect of discounting on later cash flows.

The comparison therefore highlights that the most appropriate mining method depends on the strategic objective of the business. If the priority is to accelerate cash flow generation and maximise discounted value, sublevel stoping with three sublevels and crush pillars is the recommended option for the current conditions. If the priority is to maximise recovered metal, improve selectivity and extend the productive contribution of the narrow-vein areas, drift and fill remains a strategically relevant alternative.

Based on the current assumptions, this study supports the selection of Scenario 3 as the preferred mining method for implementation in the studied narrow-vein area at Cuiabá Mine. Nevertheless, Scenario 4 should not be discarded. Further technical work should focus on opportunities to improve drift and fill productivity, reduce development intensity, optimise the backfill cycle and test the sensitivity of the method to cost, dilution and production rate assumptions.

The broader conclusion is that underground mining method selection in mature operations should be treated as a recurring strategic planning exercise rather than a fixed operational decision. As orebody geometry, operating performance, cost structures and corporate objectives evolve, quantitative reassessment provides a robust basis for aligning mining method selection with both operational constraints and business strategy.

ACKNOWLEDGEMENTS

The author acknowledges the Federal University of Minas Gerais for the academic support provided during the development of the MSc research on which this paper is based, and Dr Michel Melo Oliveira for his technical guidance and supervision.

The author also acknowledges AngloGold Ashanti for supporting the development of the work, authorising the publication of this paper and allowing the use of case study information from Cuiabá Mine. The contributions of colleagues involved in mine planning, geology, geotechnics and technical services are also gratefully acknowledged.

ChatGPT was used as an editorial support tool during the preparation of this paper, including language refinement, structure review and consistency checks. The technical content, data interpretation, conclusions and final manuscript are based on the author's MSc research and remain the responsibility of the author.

REFERENCES

Copland, T and Nehring, M, 2016. Integrated optimization of stope boundary selection and scheduling for sublevel stoping operations, *Journal of the Southern African Institute of Mining and Metallurgy*, 116(12):1135–1142. https://doi.org/10.17159/2411-9717/2016/v116n12a7

Farina, F, Albert, C, Lana, C, Stevens, G, Storey, C, Gerdes, A, Dopico, C M, Frei, D and Moreira, H, 2016. The Archean–Paleoproterozoic evolution of the Quadrilátero Ferrífero, Brazil: current models and open questions, *Journal of South American Earth Sciences*, 68:4–21. https://doi.org/10.1016/j.jsames.2015.10.015

Grobler, H, Chatziefstratiou, V, Mousli, O and Yumlu, M, 2019. Design of high-strength backfill for a drift-and-fill mining method at Olympias Mine, Greece, in *Paste 2019: Proceedings of the 22nd International Conference on Paste, Thickened and Filtered Tailings* (eds: A J C Paterson, A B Fourie and D Reid), pp 401–410 (Australian Centre for Geomechanics: Perth). https://doi.org/10.36487/ACG_rep/1910_28_Grobler

Hamrin, H, 1998. Choosing an underground mining method, in *Techniques in Underground Mining* (eds: R E Gertsch and R L Bullock), pp 45–85 (Society for Mining, Metallurgy and Exploration: Littleton).

Hoek, E, 2007. *Practical Rock Engineering* [online]. Available from: <https://www.rocscience.com/learning/hoek/practical-rock-engineering>

Lane, K F, 1988. *The Economic Definition of Ore: Cut-Off Grades in Theory and Practice* (Mining Journal Books: London).

Little, J, Knights, P and Topal, E, 2013. Integrated optimization of underground mine design and scheduling, *Journal of the Southern African Institute of Mining and Metallurgy*, 113(10):775–785. Available from: <https://www.saimm.co.za/Journal/v113n10p775.pdf>

Rendu, J-M, 2013. *An Introduction to Cut-Off Grade Estimation*, 2nd edition (Society for Mining, Metallurgy and Exploration: Littleton).

Silva, R P S, 2026. Definição do método de lavra subterrânea para veios estreitos em profundidade na Mina Cuiabá (Definition of the underground mining method for deep narrow veins at the Cuiabá mine), MSc thesis, Universidade Federal de Minas Gerais, Belo Horizonte.

Stewart, P, 2005. Minimising dilution in narrow-vein mines, PhD thesis, The University of Queensland, Brisbane. Available from: <https://core.ac.uk/download/pdf/15022456.pdf>

Underground planning – caving and stope optimisation

Integration of machine learning and mixed-integer programming in underground mine planning

P Chimunhu[1], E Topal[2], M W A Asad[3], R S Faradonbeh[4] and A D Ajak[5]

1. Mining Engineer, Sessional Academic – WA School of Mines, Curtin University, Bentley WA 6102. Email: prosper.chimunhu@curtin.edu.au
2. Professor, WA School of Mines, Curtin University, Bentley WA 6102. Email: e.topal@curtin.edu.au
3. Professor, WA School of Mines, Curtin University, Bentley WA 6102. Email: waqar.asad@curtin.edu.au
4. Lecturer, WA School of Mines, Curtin University, Bentley WA 6102. Email: roohollah.shiranifaradonbeh@curtin.edu.au
5. Mining Consultant, Dufico Consulting, Perth WA 6000. Email: ajak.ajak@duficoconsulting.com

INTRODUCTION

Persistent inconsistencies between schedule forecasts and actual production in mining operations, driven by imprecise parametric inputs such as mining dilution, remain a significant risk in mine planning and production scheduling (Amoako, Chowdu and Brickey, 2023; MacLean, 2017). Unplanned dilution affects profitability through additional resources and costs of handling sub-economic material (Bagde, 2021; Cadenillas, 2023). Despite their inherent limitations in addressing uncertainty, deterministic models such as mixed integer programming (MIP) remain dominant in production scheduling optimisation due to their simplicity, mathematical rigour and practical applicability. Concurrently, a growing body of evidence points to Machine Learning (ML)'s ability to model complex, non-linear relationships between input variables and outcomes with high computational efficiency and predictive performance, positioning it as an effective tool for improving the estimation of uncertain planning inputs (Chimunhu *et al,* 2024; Sahoo and Tripathy, 2025). This study harnesses ML's predictive capabilities by integrating ML-predicted dilution inputs into a deterministic MIP scheduling framework to evaluate their impact on traditional MIP scheduling solutions and improve alignment between scheduled and actual production outcomes.

METHODOLOGY

The optimisation methodology integrates ML-predicted dilution inputs with a classical MIP scheduling framework for a sub-level open stoping operation in Western Australia, with explicit provenance from Trout (1995), Topal (2003) and Little, Knights and Topal (2013). Mining is undertaken on 35–75° dipping reefs, 2–4 m wide, using a staggered top-down sequence. Typical stope sizes range from 1500–15 000 t, with dilution averaging 40 per cent. For the dilution prediction framework, a standalone gene expression programming (GEP) model was trained and deployed upstream of the MIP optimisation framework as a companion dilution prediction model. Model choice is informed by the simplicity of the GEP model's tree-structured output, which, combined with its high performance, speed, predictive reliability and minimal data transformation requirements (Faradonbeh and Monjezi, 2017; Ferreira *et al,* 2002), makes it preferable to more complex black-box models.

Dilution prediction phase – enhancing input robustness through ML

Two dilution factors were used to estimate mined tonnes input for schedule optimisation: an empirical factor (Emp), derived from histogram data at the P75 confidence level in accordance with current mine practice, and an ML-predicted factor (GEP) generated for each stope. The GEP model predicted dilution at stope level using geotechnical and design parameters (dip, hydraulic radius, modified stability number, and stope size), achieving an R^2 value of 0.74 against observed dilution values (Chimunhu *et al,* 2025a). A comparative overview of the two dilution factors against actual reconciled dilution (Act) for 23 independent stopes is presented in Figure 1.

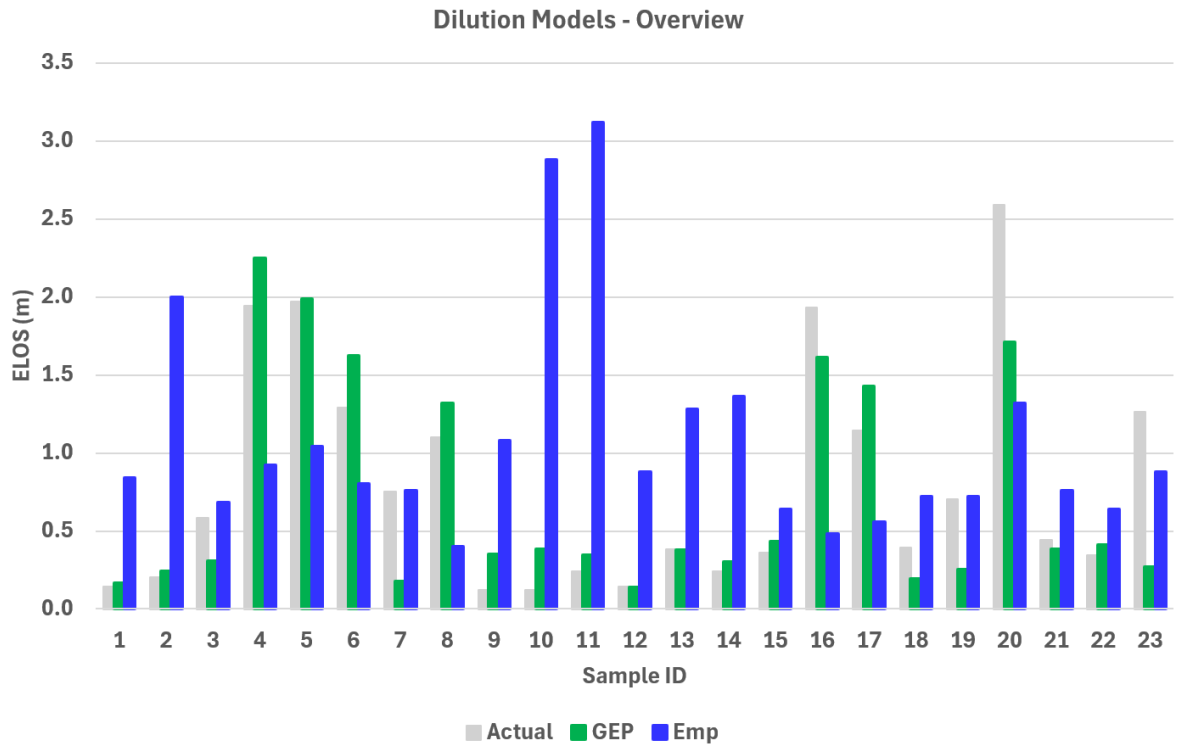


FIG 1 – Comparative overview of dilution factors across 23 independent stopes.

ML's per-stope predictions provide greater visibility of localised variations and associated schedule implications (Figure 1), enabling proactive mitigation of potential schedule slippage and performance deviations.

MIP optimisation phase – leveraging mathematical rigour for scheduling

The classical MIP model was developed in IBM CPLEX and programmed in Python using tonnage inputs derived from different dilution factors to generate schedules. A 61-stope sample from a sub-level open stoping operation was used to generate schedules (Table 1), which were optimised over a 16 month horizon under typical operational constraints (Chimunhu *et al,* 2025b). The modular architecture supports integration of auxiliary ML prediction models into established optimisation frameworks while preserving their mathematical rigour. Accordingly, tonnage estimates derived from the Emp, Act and GEP dilution factors were used to generate the MIP-Emp, MIP-Act and MIP-GEP schedules, respectively. For benchmarking completeness, a manual Emp-based schedule was also developed to compare mathematical optimisation against conventional manual scheduling (Table 1).

TABLE 1

Schedules' summary description.

Model	Optimisation criteria	Dilution factor	Model comparisons	Comparison objectives
MAN-Emp	Manual	Empirical		
MIP-Emp	MIP	Empirical	MAN-Emp versus MIP-Emp	Assess impact of the scheduling optimisation method
MIP-GEP	MIP	ML-GEP	MIP-Emp versus MIP-GEP	Assess overall impact of dilution input per stope versus the flat factor for all stopes
MIP-Act	MIP	Actual	All	Benchmark for all schedules

RESULTS

Preliminary results are summarised in Table 2.

TABLE 2

Schedules – summary description and results.

Parameter	MIP-Emp versus MIP-Act				MIP-GEP versus MIP-Act				Net Var
Models	MIP-Emp	MIP-Act	Var	Var%	MIP-GEP	MIP-Act	Var	Var%	
Dilution factor (%)	40	Act			variable	Act			
Mined tonnes (t)	209 095	198 282	10 813	5.5	195 928	198 282	-2354	-1.2	4.3
NPV$(000)	40 540	41 594	-1054	-2.5	41 832	41 594	237	0.6	-2.0

Var – variance.

Summary

ML-enhanced dilution inputs improve schedule predictive performance, reducing benchmark deviations by 4.3 per cent and 2.0 per cent for mined tonnes and NPV, respectively. Schedules optimised using per-stope ML-predicted dilution best matched the benchmark schedule (Table 3), suggesting improved deterministic MIP robustness.

TABLE 3

Schedules – overview of schedules' performance on key metrics.

Model	Number of stopes mined	Total mined tonnes (tonnes)	Recovered metal (ounces)	Net present value ($)
MAN-Emp	58▼	200 439▲	38 280▼	39 297 543▼
MIP-Emp	61 —	209 095▲	39 584 —	40 540 242▼
MIP-GEP	61 —	195 928▲	39 584 —	41 831 771▲
MIP-Act	61 —	198 282 —	39 584 —	41 594 298 —

Discussion

Key schedule insights over the 16 month planning horizon are summarised further, while Figure 2 illustrates the temporal granularity of NPV performance:

- Number of stopes mined – MIP-based schedules outperform manual scheduling.
- Tonnes mined – ML-enhanced dilution inputs improve tonnage forecast reliability.
- Metal ounces – ML-enhanced tonnage forecasts enhance schedule reliability.
- NPV – ML-enhanced inputs improve the robustness of cash flow forecasts and profitability.

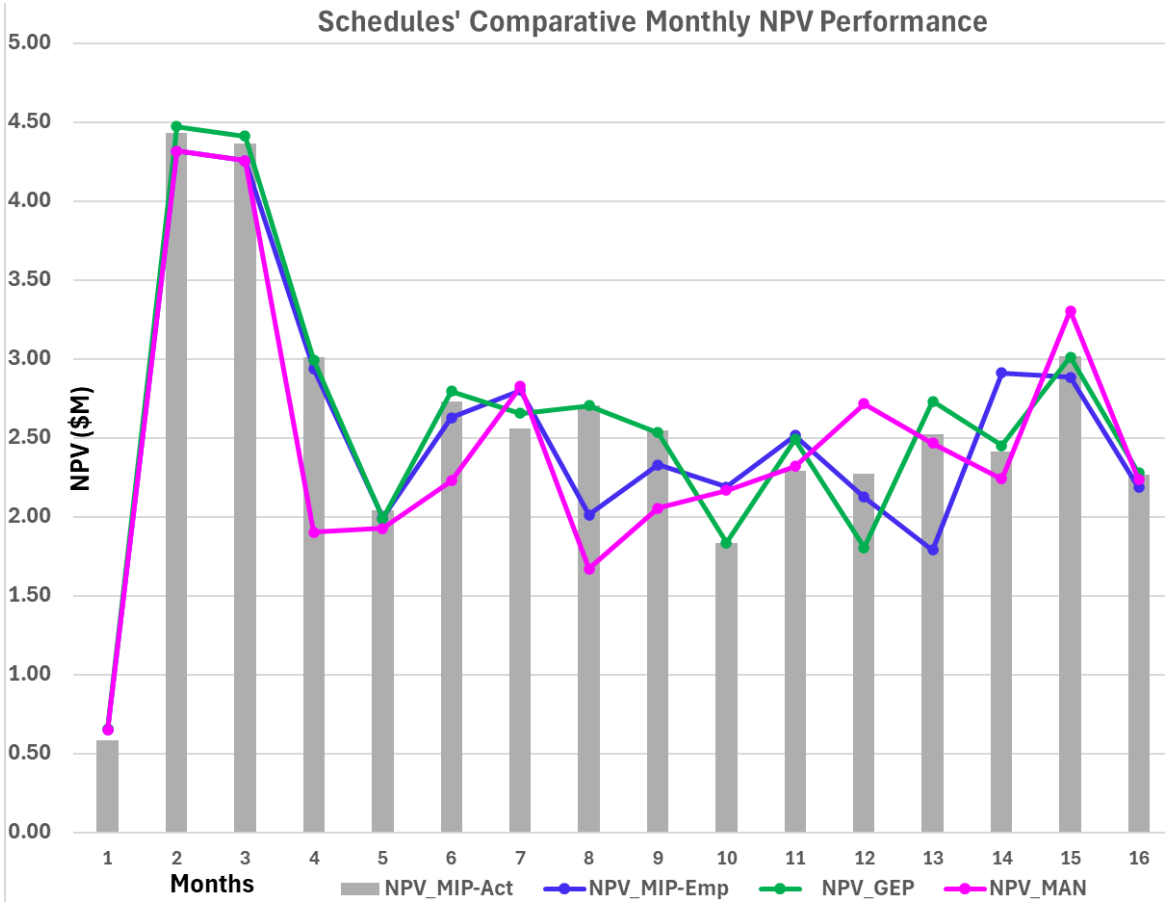


FIG 2 – Schedules' monthly performance based on NPVs.

Overall, the MIP-GEP architecture increases optimisation search space, enabling more feasible monthly stope combinations to be evaluated without violating key scheduling constraints. This flexibility, enabled by granular dilution prediction, explains the superior performance of the MIP-GEP model over the MIP-Emp, despite both completing all stopes and achieving the same ounces within 16 months.

Collectively, these results suggest that ML enhances solution robustness, reliability, and overall scheduling performance.

CONCLUSIONS

ML improves deterministic production scheduling by providing higher-quality stope-level dilution estimates for MIP optimisation, reducing planned-versus-actual discrepancies. Future work should extend the framework to other canonical mine planning inputs by integrating multiple ML models within a unified optimisation process. Additionally, probabilistic ML-derived dilution distributions should be explored using stochastic integer programming (SIP) to explicitly account for input uncertainty and assess schedule robustness.

ACKNOWLEDGEMENTS

The authors received no financial support for the research, authorship, and/or publication of this article.

REFERENCES

Amoako, R, Chowdu, A and Brickey, A, 2023. A Deviation-Minimization Approach to Short-Term Underground Mine Schedule Optimization, *Mining, Metallurgy and Exploration*, 40(5):1749–1765. https://doi.org/10.1007/s42461-023-00802-5

Bagde, M N, 2021. Ore and Backfill Dilution in Underground Hard Rock Mining, *Journal of Mining Science*, 57(6):995–1005. https://doi.org/10.1134/S1062739121060120

Cadenillas, C T, 2023. Prediction of Unplanned Dilution in Underground Mines Through Machine Learning Techniques, MSc thesis, McGill University, Canada. https://doi.org/10.82308/19434

Chimunhu, P, Faradonbeh, R S, Topal, E, Asad, M W A and Ajak, A D, 2025a. Dilution prediction in underground open stope mining using gene expression programming and backpropagation artificial neural network algorithms, *Mining Technology*, 134(2):105–120. https://doi.org/10.1177/25726668251348707

Chimunhu, P, Topal, E, Asad, M W A, Faradonbeh, R S and Ajak, A D, 2024. The future of underground mine planning in the era of machine learning: Opportunities for engineering robustness and flexibility, *Mining Technology*, 133(4):331–347. https://doi.org/10.1177/25726668241281875

Chimunhu, P, Topal, E, Asad, M W A, Faradonbeh, R S and Ajak, A D, 2025b. Production scheduling optimisation using mixed integer programming with machine learning dilution prediction capabilities for underground open stoping operations, *Machine Learning with Applications*, 22:100776. https://doi.org/10.1016/j.mlwa.2025.100776

Faradonbeh, R S and Monjezi, M, 2017. Prediction and minimization of blast-induced ground vibration using two robust meta-heuristic algorithms, *Engineering with Computers*, 33(4):835–851. https://doi.org/10.1007/s00366-017-0501-6

Ferreira, C, Furuhashi, T, Köppen, M, Hoffmann, F, Ovaska, S and Roy, R, 2002. Gene Expression Programming in Problem Solving, in *Soft Computing and Industry*, pp 635–653 (Springer: London). https://doi.org/10.1007/978-1-4471-0123-9_54

Little, J, Knights, P and Topal, E, 2013. Integrated optimization of underground mine design and scheduling, *Journal of the Southern African Institute of Mining and Metallurgy*, 113(10):775–785.

MacLean, J, 2017. Biggest risks for mining companies shift, yet challenges remain, *Canadian Mining Journal*, 138(1):6. Available from: <https://www.canadianminingjournal.com/digital-edition/january-2017/>

Sahoo, A K and Tripathy, D P, 2025. Applications of AI and machine learning in mining: digitization and future directions, *Safety in Extreme Environments*, 7(1):4. https://doi.org/10.1007/s42797-025-00118-1

Topal, E, 2003. Advanced Underground Mine Scheduling Using Mixed Integer Programming, PhD thesis, Colorado School of Mines.

Trout, L, 1995. Underground mine production scheduling using mixed integer programming, 25th International APCOM Symposium Proceedings, Melbourne.

Geotechnical simulation for block caving – quantifying rock mass uncertainty to inform long-term planning

S Light[1], M Musolino[2], R Barnes[3], C Chester[4] and I Minniakhmetov[5]

1. Principal Geologist, BHP Copper SA, Adelaide SA 5000. Email: shaun.light@bhp.com
2. Senior Mining Engineer, BHP Copper SA, Adelaide SA 5000. Email: matthew.musolino@bhp.com
3. Principal Data Scientist, BHP, Brisbane Qld 4000. Email: rachael.barnes@bhp.com
4. Principal Geotechnical Engineer, BHP Copper SA, Adelaide SA 5000. Email: chris.chester@bhp.com
5. Principal Geoscientist, BHP, Perth WA 6010. Email: ilnur.minniakhmetov@bhp.com

ABSTRACT

Robust geotechnical characterisation is fundamental to long-term planning in block caving, where cave initiation, propagation, and draw performance are highly sensitive to rock mass conditions. Conventional estimation approaches smooth spatial variability, masking critical heterogeneities that influence caveability, fragmentation, and the risk of stalled propagation or airgaps. This paper presents a cross-disciplinary workflow that explicitly models geotechnical uncertainty using advanced conditional simulation to generate multiple three-dimensional realisations of key rock mass parameters. Simulated properties are transformed into Rock Mass Rating (RMR) and propagated through a cave simulator to quantify ranges of outcomes in cave-back progression, airgap risk, and cave morphology. Results demonstrate how local contrasts in rock mass properties materially affect propagation rates and the likelihood of airgaps, providing probabilistic rather than deterministic inputs for strategic decisions including draw strategy, reinforcement priorities, and production ramp-up pacing.

INTRODUCTION

Block cave mining relies on the progressive failure and flow of rock under gravity. The rate and geometry of cave propagation are governed by rock mass properties that vary spatially in ways poorly captured by conventional kriged estimates. Smoothing inherent in estimation suppresses short-scale variability and local contrasts that control whether a cave stalls, develops airgaps, or propagates asymmetrically. Fine-scale geological structures such as faults and dyke contacts exert a disproportionate influence on caveability and fragmentation (Brown, 2003), yet their spatial uncertainty is rarely quantified explicitly. These outcomes have direct consequences for production ramp-up, drawpoint availability, and personnel safety.

The purpose of this work is not to produce multiple cave designs for execution, but to stress-test the robustness of the current cave plan against plausible rock mass variability. By quantifying the range of cave behaviour, the framework identifies which risks could prevent plan implementation and what design buffers—in draw strategy, support systems, or ramp-up pacing—are needed to protect the critical path. The approach required geologists, geostatisticians, data scientists, and geotechnical engineers to collaborate end-to-end, linking data acquisition, domain interpretation, spatial modelling, and cave performance assessment in a single decision-focused framework.

Critically, the strategic decisions informed by this analysis—cave footprint extent, extraction-level layout, ventilation infrastructure, and production ramp-up rate require commitments years before cave initiation. Probabilistic geotechnical characterisation must therefore be conducted proactively in the strategic planning space to support these long-lead-time decisions with quantified confidence. The case study is from a copper block cave operation.

METHODOLOGY

Categorical domain simulation

The geological setting comprises multiple lithological units including host rock, dyke intrusions, and cover sequences. Hierarchical truncated pluri-Gaussian simulation (HTPGS) is used to simulate the

spatial distribution of mineralisation domains, dyke boundaries, and geotechnical domains (Figure 1). HTPGS preserves the spatial relationships and contact patterns between rock types without imposing deterministic boundaries, generating multiple plausible realisations of the geological architecture that controls rock mass quality distribution. Local proportions and spatial trends are incorporated to handle the strong non-stationarity typical of complex intrusive settings.



FIG 1 – West–east cross-section showing one realisation of categorical domain simulation (HTPGS). Colours represent distinct lithological and geotechnical domains including host rock, dyke intrusions, and cover sequences. Grey semi-transparent body is the deterministic interpretation of breccia body. All coordinates removed due to commercial sensitivity.

Continuous geotechnical variable simulation

Four key geotechnical variables are simulated within the categorical domain framework: uniaxial compressive strength (UCS), rock quality designation (RQD), discontinuity spacing, and discontinuity separation. These are the primary inputs to the Rock Mass Rating (RMR) classification system that governs cave behaviour in the simulator. Sample data are transformed using declustering, normal score transformation, and non-stationary detrending to achieve stationarity (Williams and Minniakhmetov, 2025). The projection pursuit multivariate transform (PPMT) is applied to decorrelate the variables while preserving complex multivariate relationships (Barnett, Manchuk and Deutsch, 2014). Turning bands simulation (Journel, 1974) generates spatially correlated realisations for each decorrelated variable, followed by back-transformation in reverse order. Twenty realisations of the complete geotechnical block model are produced, each honouring conditioning data and reproducing observed spatial continuity and multivariate structure.

RMR calculation and cave simulation

Simulated geotechnical properties are combined to calculate RMR for each block in each realisation, providing a spatially variable input to the cave simulator. Each realisation is run through the simulator independently, producing an ensemble of cave propagation outcomes. The simulator tracks cave-back progression, drawpoint muck pile thickness, airgap development, and cave morphology through time. A reference case using uniform RMR (set at the deposit average) is included to isolate the effect of spatial variability from overall average rock mass quality.

RESULTS

Figure 2 shows the cave morphology at an intermediate draw step, with each colour representing a different geotechnical scenario (realisation). The overlaid cave shapes demonstrate clear differences in propagation geometry driven by local contrasts in rock mass properties across realisations. Some scenarios show more asymmetric cave development, which influences airgap distribution and drawpoint loading patterns.

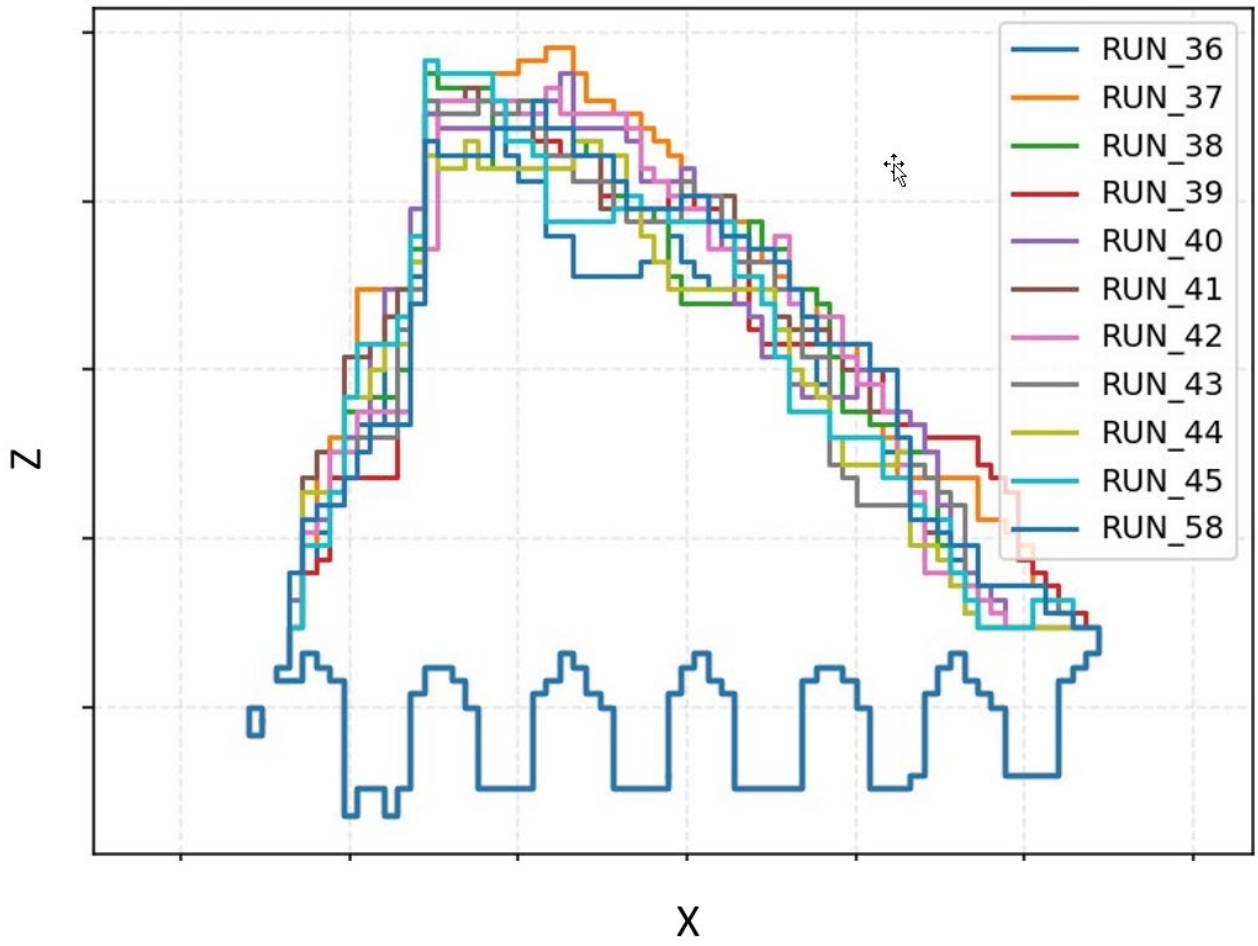


FIG 2 – Cave morphology at an intermediate draw step showing overlaid profiles from multiple geotechnical realisations (scenarios). Differences in cave shape are driven by spatial variability in simulated rock mass quality. Coordinates removed due to commercial sensitivity.

Airgap risk analysis reveals that while peak cave-back levels converge across realisations, the growth paths to reach that point differ substantially. The number of drawpoints experiencing high-risk airgap conditions varies materially between realisations demonstrating that deterministic averaging of geotechnical inputs masks operationally significant variability in cave behaviour. Cave morphology at intermediate steps shows clear differences in shape driven by local contrasts in UCS, RQD, and fracture parameters. These results confirm that the simulated block model, with its preserved short-scale heterogeneity, provides materially different cave behaviour predictions compared to a smoothed deterministic input.

STRATEGIC IMPLICATIONS

The ensemble-based results directly inform strategic planning decisions that must be committed years before cave initiation. The framework enables:

- draw strategy adjustments and extraction sequencing to manage airgap risk in zones of weaker rock mass.
- reinforcement and support priorities based on probabilistic identification of high-risk areas.
- production ramp-up pacing calibrated to the range of plausible propagation rates.
- targeted drilling campaigns to reduce the most consequential uncertainties identified by the ensemble.

The framework provides a transparent linkage from data uncertainty through to operational risk, replacing reliance on downstream risk controls with proactive, quantified inputs to strategic decisions. The methodology is applicable to any block cave operation where geotechnical variability influences cave performance and offers a consistent approach across the full planning continuum.

ACKNOWLEDGEMENTS

The authors acknowledge the collaboration between the Geotechnical, Geology, Mining Engineering, and Data Science teams at BHP, whose cross-disciplinary partnership enabled the end-to-end integration of geotechnical uncertainty into cave planning.

REFERENCES

Barnett, R M, Manchuk, J G and Deutsch, C V, 2014. Projection Pursuit Multivariate Transform, *Mathematical Geosciences*, 46:337–359.

Brown, E T, 2003. *Block Caving Geomechanics*, second edition (Julius Kruttschnitt Mineral Research Centre: Brisbane).

Journel, A G, 1974. Geostatistics for Conditional Simulation of Ore Bodies, *Economic Geology*, 69:673–687.

Williams, C and Minniakhmetov, I, 2025. In situ uncertainty quantification and implications for risk mitigation – a case study from an iron ore operation, in *Proceedings of the Mineral Resource Estimation Conference 2025*, pp 333–356 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Simultaneous stochastic optimisation of underground mining complexes – production and development scheduling with multiple design alternatives

A MacRae[1] and L Petro[2]

1. MSc Candidate, COSMO Stochastic Mine Planning Laboratory, McGill University, Montreal QC H3A 0G4, Canada. Email: angus.macrae@mail.mcgill.ca
2. Principal, Resource Development Plan, Vale Base Metals, Toronto ON M5J 2K2, Canada. Email: luis.petro@vale.com

ABSTRACT

Simultaneous stochastic optimisation of mining complexes (SSOMC) has been shown to generate mine plans with greater value than conventional methods while effectively managing risk. However, the methodology has previously only been applied in open pit operations. In underground mines, production and development scheduling/design decisions influence one another in an interconnected manner, along with the destinations and downstream routing of material. These decisions should therefore be optimised in a simultaneous and stochastic framework to capitalise on such nonlinear interactions while managing uncertainty.

The proposed stochastic integer program (SIP) extends SSOMC to sublevel open stoping operations, jointly optimising the mine design and schedule. The SIP selects an optimal set of feasible (non-conflicting) design elements from a set of candidates, along with their production schedule, under a unified objective that discounts incremental development costs on an as-needed basis. The method offers flexibility to optimise decisions between multiple design alternatives jointly in a simultaneous framework with the production schedule. The resulting combinatorial optimisation is not linear or differentiable, requiring the use of black box optimisation strategies to solve on a realistic scale. A custom implementation of the Greedy Randomised Adaptive Search Procedure (GRASP) metaheuristic solves the formulation in a real-world sublevel stoping operation with multiple mines. A case study demonstrates the algorithm's ability to outperform conventional, sequential and deterministic methods.

INTRODUCTION

Underground mine design and planning are conventionally approached in a sequential framework in which the stope design/layout is first optimised and used to design the development layout, which are then both used as inputs for long-term production scheduling (Alford, Brazil and Lee, 2016). This sequential framework neglects to account for interactions between scheduling and design decisions, which arise due to time value of money and the incremental nature of development costs (Furtado e Faria, Dimitrakopoulos and Lopes Pinto, 2022). Additionally, common stope optimisation approaches require the initial assumption of an economic cut-off grade, which must be assumed upstream of the design (Alford, Brazil and Lee, 2016), and oversimplify the destination policy (Goodfellow and Dimitrakopoulos, 2016). Finally, it is typically performed deterministically, with the orebody model and all other input parameters provided as a single mean estimate, failing to account for uncertainty (Carelos Andrade, Dimitrakopoulos and Conway, 2024; Dimitrakopoulos and Grieco, 2009). The proposed research addresses these shortcomings by jointly optimising the mine design, stope layout and production schedule for sublevel open stoping operations within the simultaneous stochastic optimisation of mining complexes (SSOMC) model proposed by Goodfellow and Dimitrakopoulos (2016).

Stochastic mine planning approaches are well established in open pit mine planning literature for creating value and managing risk (Leite and Dimitrakopoulos, 2007; Albor Consuegra and Dimitrakopoulos, 2009; Ramazan and Dimitrakopoulos, 2013, 2018; Benndorf and Dimitrakopoulos, 2013). Although stochastic integer programming models can generalise to capture many sources of uncertainty, such as market and recovery uncertainties (Jiang and Dimitrakopoulos, 2025), most previous work has focused on quantifying orebody uncertainty through geostatistical simulation (Dimitrakopoulos, 2011). Simulation produces many equiprobable realisations of an orebody,

quantifying variability and uncertainty, which would be neglected by estimation approaches such as kriging (Goovaerts, 1997; Rossi and Deutsch, 2014). Stochastic integer programming offers the ability to optimise for uncertain inputs, resulting in greater value under uncertainty (Birge and Louveaux, 2011). In addition to being analytically proven, the explicit positivity of the value of the stochastic solution (VSS) has been demonstrated in a mine planning context by Dimitrakopoulos and Ramazan (2008).

Simultaneous optimisation principles have been applied to unify decision-making at different stages of the planning process. Furtado e Faria, Dimitrakopoulos and Lopes Pinto (2022) proposed a stochastic mixed integer linear program to simultaneously optimise mine design, stope layout and production scheduling in sublevel open stoping operations by selecting and sequencing a non-overlapping set of stopes from a densely overlapping set of preprocessed candidate designs. The stope candidates' base elevations are dictated by a higher-level mine design decision, which selects a non-conflicting set of level designs. Carelos Andrade, Dimitrakopoulos and Conway *(*2024) proposed a similar model, which also scheduled delayed backfill activities and respected a specific operation's precedence model. This model included design candidates dictated by the vertical sublevel spacing and maximum stope width and included a single stockpile in a linearised framework.

Goodfellow and Dimitrakopoulos (2016) addressed the shortcomings of precomputed economic block values in an open pit setting through the SSOMC. A mining complex is an integrated business unit consisting of multiple interconnected mines, stockpiles, waste dumps and processing resources. It can be modelled as a connected graph in which nodes represent these destinations, and edges represent the ability to route material from one destination to another. Economic/recovery factors are then applied to the material at each destination in each period. Conventional methods linearise the problem by applying these factors to each individual mining block/stope, thereby assuming its downstream destination *a priori*. SSOMC instead optimises the destination policy and downstream routing decisions jointly with the extraction sequence, more closely modelling operational flexibility, and capturing nonlinear interactions between the two sets of decisions. Introducing such nonlinearity eliminates the possibility of an exact solution and necessitates the use of heuristic optimisation methods (Leite and Dimitrakopoulos, 2007; Lamghari, Dimitrakopoulos and Ferland, 2014; Montiel and Dimitrakopoulos, 2017).

Sublevel open stoping is an underground mining method in which stopes are excavated between sublevels and subsequently backfilled to maintain stability for future stopes. Depending on the competence of the rock, pillars may be left between stopes for stability. Drilling may be performed from the sublevel above or below the stopes, after which stopes are blasted and excavated from the level below (Villaescusa, 2014). Many optimisation approaches have been proposed towards the planning and design of such operations, though most have optimised either the stope layout, mine design or production schedule independently rather than jointly (Topal and Sens, 2010; Alford, Brazil and Lee, 2016; Sotoudeh *et al*, 2020).

Stope design optimisation is generally approached in the literature either as a geometric optimisation (ie with continuous decision variables that determine 2D or 3D stope boundaries), or as a combinatorial, discrete selection from candidate designs. The proposed research falls into the latter category and places the generation of candidate stope design boundaries outside the scope. However, many algorithms have been proposed and/or commercialised to optimise stope boundaries (Alford, Brazil and Lee, 2016).

An early approach is the Floating Stope Algorithm (Alford, 1995) which is analogous to the floating cone method in open pit mine planning. This algorithm floats a minimum stope shape through the orebody and evaluates the aggregated stope grade at each possible stope position. A maximum envelope is generated containing the union of all economic stope positions. A minimum envelope contains the union of the highest-grade stope outlines for each block in the orebody model. The envelopes provide inner and outer boundaries; however, engineers still must manually design stopes between the envelopes. Alford and Hall (2009) proposed a method to improve on the floating stope algorithm by optimising the final stope/pillar shape while accounting for dilution and narrow or steep-dipping orebodies. It takes a seed stope shape along with geometric/pillar constraints and anneals the shape to maximise recovered value.

Ataee-Pour (2004) proposed the maximum value neighbourhood (MVN) algorithm, a heuristic that steps through the block model, locating for each block the highest-value feasible neighbourhood of blocks satisfying a minimum stope size and flagging it if its marginal contribution is positive. Topal and Sens (2010) later proposed a structurally similar heuristic: candidate stopes are ranked by profit (or profit per area or time), the top-ranked stope is selected and all overlapping candidates removed, with an 'envelope-average' check used to justify each selection before repeating until no candidates remain. Both methods select stopes through a sequential, greedy process, ranking or evaluating candidates one at a time and committing to each selection immediately, without searching over the full combinatorial space of possible non-overlapping configurations.

Sari and Kumral (2021) formalised the problem as a weighted independent set, in which candidate stope designs are represented as vertices weighted by their economic value, and overlapping stopes are connected by edges, so that a maximum-weight independent set corresponds to a non-overlapping, profit-maximising stope layout. They demonstrate the computational challenges of solving the problem exactly at scale. To address the computational challenges, they proposed a greedy heuristic that limits the number of retained combinations to obtain comparable results much faster than the exact method.

All the aforementioned methods are deterministic, considering a single smooth estimate of the orebody and thereby failing to account for uncertainty. Dimitrakopoulos and Grieco (2009) address this by applying conditional simulation to quantify geological risk for a stope design generated with the floating stope algorithm. They found that conventional deterministic methods neglect to account for significant upside potential and downside risk. Furthermore, they proposed a probabilistic alternative that allows the user to define an acceptable level of risk that a stope design will fall below cut-off. As a probabilistic formulation, this model designs stopes to meet a certain confidence threshold but does not explicitly minimise uncertainty as a stochastic optimisation approach would.

All aforementioned methods optimise the stope design and layout upstream of production scheduling, without accounting for time value of money. Production scheduling has previously been approached with mixed integer programming (MIP) based on a fixed mine design (Alford, Brazil and Lee, 2016; Sotoudeh *et al*, 2020). By neglecting production scheduling, the stope optimisations also fail to explicitly account for the incremental, time-discounted nature of development costs (Furtado e Faria, Dimitrakopoulos and Lopes Pinto, 2022). In fact, the only way they can account for development costs is to apply a precomputed life-of-mine (LOM) average to the stope's economic values (Alford, Brazil and Lee, 2016) which cannot be known with any confidence prior to stope optimisation.

Unlike any of the other aforementioned models, Furtado e Faria, Dimitrakopoulos and Lopes Pinto (2022) and Carelos Andrade, Dimitrakopoulos and Conway (2024) discounted horizontal development costs in the objective function on an as-needed basis. This accounts for time value of money and incremental development costs. However, the problem is linearised through the precomputation of each stope candidate's economic value for each orebody scenario. It therefore still requires the initial assumption of an economic cut-off grade, implicitly assuming the destination of all mined material based on a low-confidence assumption.

The proposed SIP builds on the mine design, stope layout and production scheduling decisions and constraints proposed by Furtado e Faria, Dimitrakopoulos and Lopes Pinto (2022) by integrating them into a mining complex model (Goodfellow and Dimitrakopoulos, 2016). This allows for the simultaneous optimisation of multiple underground mines, considering final products and nonlinear interactions in shared stockpiling/processing resources. While previous and conventional approaches to underground mine design and planning require the initial assumption of an economic cut-off grade, the proposed SIP generates the destination policy as an output of the optimisation, capturing synergies between the design, schedule and material routing.

METHOD

Problem structure

The selection of a non-conflicting set of design alternatives dictates the set of candidate stope designs. The conflicts and linkage constraints between mine design and stope decisions are shown in Figure 1.

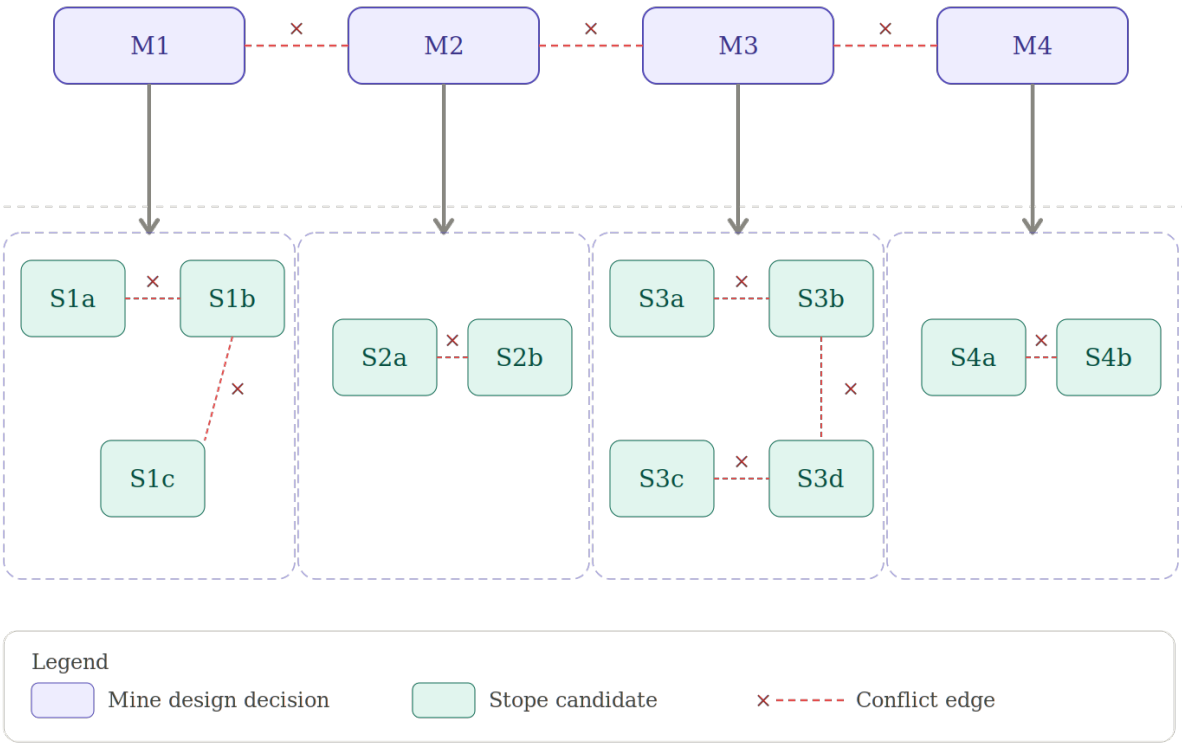


FIG 1 – Linkage relationship between mine design decisions and stope design candidates.

The proposed SIP optimises downstream destination and routing decisions simultaneously by modelling the mining complex as a connected graph where nodes represent mines, stockpiles and processing destinations while edges represent the ability to route material from one destination to another. Material flowing through the mining complex is modelled by primary attributes. Economic and recovery factors are applied to primary attributes at the appropriate destination to derive hereditary attributes which are included in the objective function. Material from mines is directed with a cluster-based destination policy (discussed in a subsequent section), and further downstream flows are decided fractionally due to the scenario dependent nature of primary attributes, such as material tonnages and contained metal. Simultaneously optimising a mining complex captures nonlinear interactions between extraction decisions and downstream routing decisions (Goodfellow and Dimitrakopoulos, 2016).

Mine design decisions

The selection of a feasible set of mine design decisions dictates the set of candidate stope designs and their associated development costs. Mine design options that cannot feasibly coexist for geometric or geomechanical reasons are assigned as conflicts. In the case study presented, mine design decisions represent the level base elevation. Vertically overlapping level designs are proposed as shown in Figure 2, and the model enforces the selection of a non-overlapping set, as proposed by Furtado e Faria, Dimitrakopoulos and Lopes Pinto (2022).

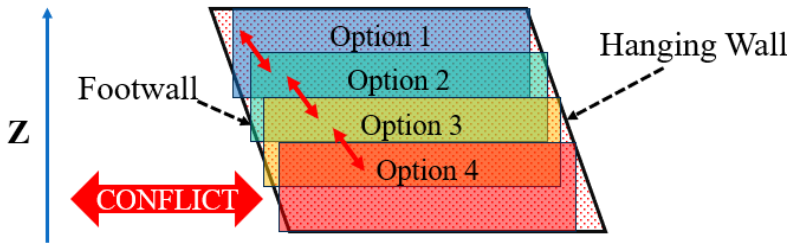


FIG 2 – Vertically overlapping level design candidates.

In addition to representing level base elevations, mine design decisions are also used to select from candidate drift layouts. At each candidate base elevation, multiple alternative drift layouts are proposed as shown in Figure 3.

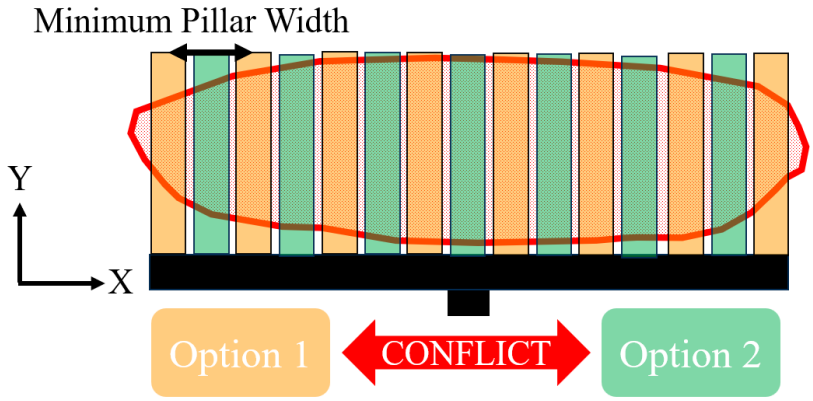


FIG 3 – Two conflicting drift layouts at the same base elevation.

Stope candidates

The set of candidate stope designs depends on the selected level base elevation and drift positions. For each mine design option, multiple stope shapes are floated at overlapping strides along each cross-cut drift to create a densely overlapping set of candidate stope designs as shown in Figure 4 (Furtado e Faria, Dimitrakopoulos and Lopes Pinto, 2022).

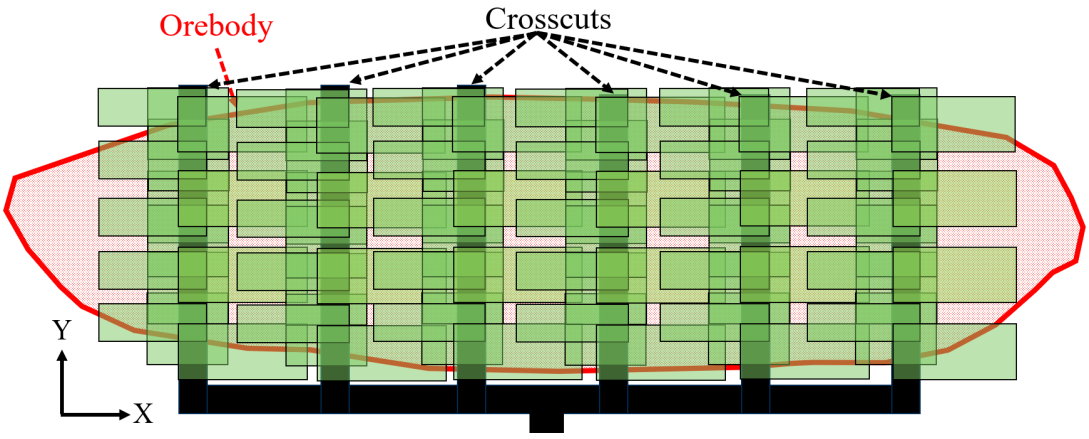


FIG 4 – Dense overlapping set of stope design candidates.

For every candidate, a pillar perimeter is defined, and any other candidates that fall within the perimeter are assigned a conflict. This enforces that only one of the conflicting pair can be selected (Furtado e Faria, Dimitrakopoulos and Lopes Pinto, 2022). Figure 5 shows the pillar perimeter for one stope candidate A. Other nearby candidates are shown, those conflicting with candidate A being shown in red and those that can be selected with stope A are shown in green.

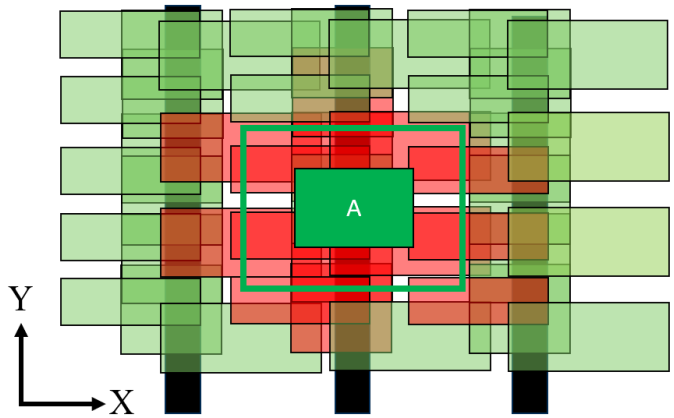


FIG 5 – Conflicting stope candidates for stope A marked in red based on pillar perimeter (green).

Cluster based destination policy

The proposed method implements the cluster-based destination policy as described by Goodfellow and Dimitrakopoulos (2016). This entails performing k-means clustering on the simulated block attributes and imposing the constraint that material belonging to the same cluster must be sent to the same destination in each period. Since the clustering attributes are scenario dependent, material can belong to different clusters and therefore be sent to different destinations in each scenario. This style of destination policy generalises well to any number of valuable/deleterious elements (Figure 6b), applies well in complicated blending situations and offers more flexibility than a simple cut-off grade while still offering a course of action when the ground truth inevitably deviates from the model. When a single element, this policy is equivalent to using grade bins (Figure 6a).

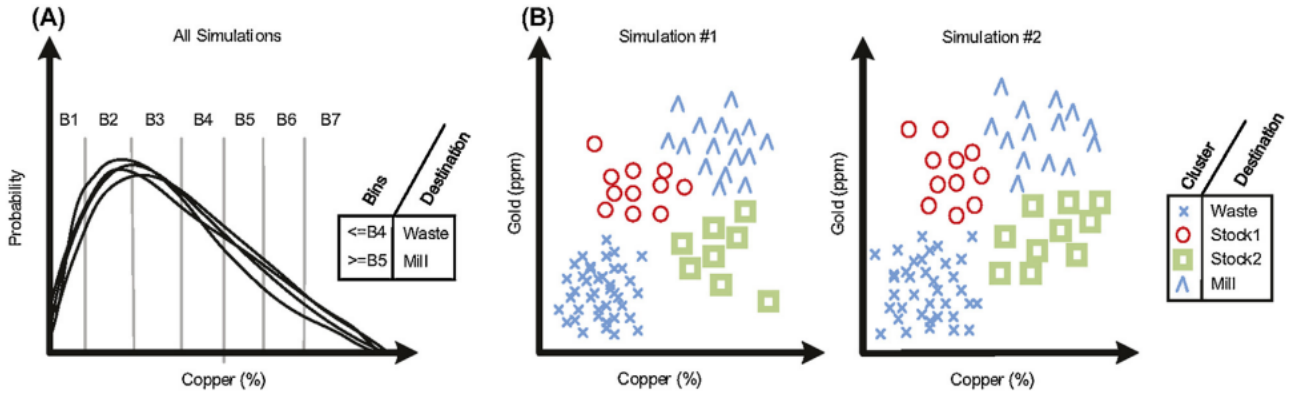


FIG 6 – (a) grade bin policy for one element; (b) cluster-based policy for two elements (Goodfellow and Dimitrakopoulos, 2016).

Most previous methods have aggregated blocks within the stope boundary and assumed material within a stope would be routed to the same destination. The proposed method routes material at a per-block level, assuming that blocks in the same stope can be routed to different destinations, representing a more selective paradigm and a closer representation of operational flexibility.

Solver

The nonlinearity introduced by considering final products and downstream decisions instead of a precomputed economic block/stope value makes the problem unsolvable with exact methods. Additionally, the non-convex, non-differentiable and combinatorial nature of the problem prevents the use of gradient-based optimisation methods. Metaheuristic optimisation algorithms efficiently explore a solution space to find very good, but not necessarily globally optimal solutions for such complex problems where finding an exact solution is unfeasible. A metaheuristic is a high-level strategy for executing heuristics (or ‘moves’) on a solution, generally exploiting those which lead to better solutions while also exploring many different regions of the solution space. Metaheuristics have been shown to perform well on underground and open pit mine design and planning problems

(Lamghari and Dimitrakopoulos, 2012; Lamghari, Dimitrakopoulos and Ferland, 2014; Sari and Kumral, 2021).

The proposed solver is based on the metaheuristic framework known as GRASP (Greedy Randomised Adaptive Search Procedure). This algorithm seeks to identify many local optima through a repeated sequence of greedy/random solution construction followed by a local search procedure to identify the nearest local optimum. Repeating this sequence many times gives a set of many locally optimal solutions, the best of which is stored as the incumbent best solution (Feo and Resende, 1995).

CASE STUDY AT AN UNDERGROUND GOLD MINING COMPLEX

The proposed method is applied at a gold mining complex consisting of two mines feeding a shared processing plant, stockpile and waste dump, as shown in Figure 7. To show the value of both simultaneous and stochastic optimisation, the proposed SIP is compared to both sequential and deterministic experiments. In the sequential experiments, the two mines are optimised separately from one another, each with a tonnage target to produce half the plant feed target. The sequential experiments also used a precomputed block economic value which was calculated by applying the recovery curve and metal price to the grade and contained metal in each block. In the simultaneous experiments, the mines are optimised jointly, applying these factors at the processing plant and giving them the ability to supplement one another so long as each mine obeys its shaft capacity. Deterministic experiments consider only a single orebody scenario, represented by the e-type (mean) grade across 25 simulations. Stochastic experiments account for geological uncertainty and variability by considering all simulations and penalising the objective for failing to meet production targets. The experiments are compared based on a risk analysis across 25 simulations.

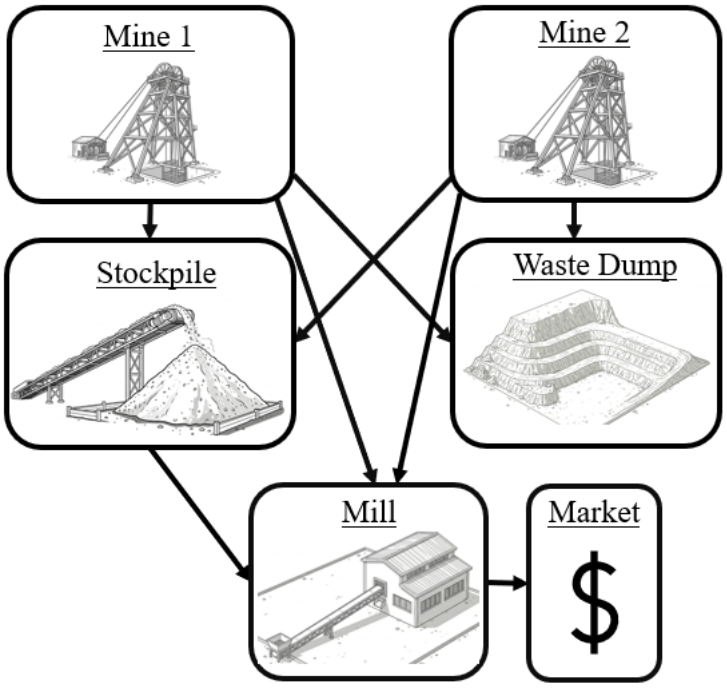


FIG 7 – Mining complex consisting of two mines, a stockpile, a waste dump and a processor.

A steeply dipping mineralised vein was conditionally simulated to generate many equiprobable realisations of the orebody. Conditional simulation of the orebody is placed outside the scope of this research. A cross-section of the steeply dipping gold vein is shown in Figure 8 with three realisations from conditional simulation, along with the average across 25 realisations. Table 1 shows economic inputs to the case study.

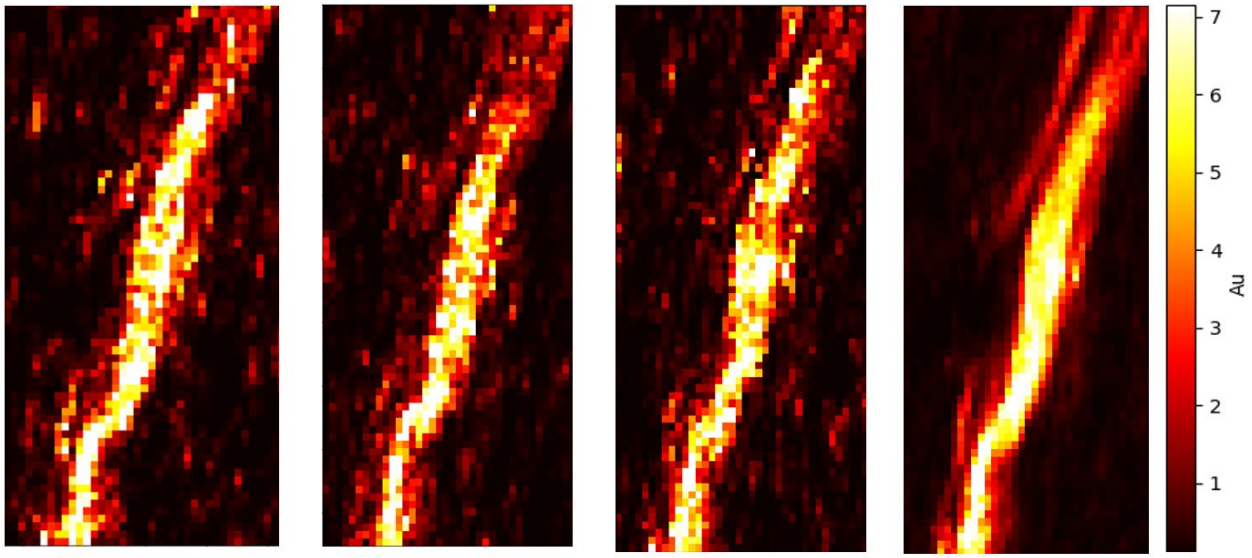


FIG 8 – Three realisations (left) from conditional simulation and e-type (right) across 25 realisations.

TABLE 1

Economic parameters for case study.

Parameter	Value
Discount rate	10%
Risk discount rate	10%
Gold price ($/oz)	$3000
Mining cost ($/t)	$50
Processing cost ($/t)	$20
Vertical development cost ($/m)	$50k
Horizontal development cost ($/m)	$20k
Service cost ($/open level/a)	$200k

Each level requires 36 m in height, but candidate level bases are placed every 18 m (Figure 2). At each candidate level base elevation, two alternative drift layouts are proposed (Figure 3) with a 12 m offset between them. Table 2 shows input parameters which dictate the size of the optimisation problem.

TABLE 2

Optimisation inputs.

Parameter	Value
# of Mines	2
# of Candidate design alternatives	96
# of Candidate stopes per alternative	355
# of Periods	10
# of Scenarios	25
# of Clusters	20 (10 per mine)

Targets and capacities are included in the model as soft constraints, meaning they can be violated but the objective is penalised. Targets and capacities for the case study are summarised in Table 3.

TABLE 3

Targets and capacities for case study.

Parameter	Value
Shaft capacity per mine (t/mine/a)	200k
Shaft capacity per level (t/a)	75k
Horizontal development capacity (m/mine/a)	700
Plant feed grade (g/t)	4–5
Plant feed (t/a)	350k

Results

A risk analysis of the cumulative net present value (NPV) found by each experiment is shown in Figure 9. The shaded regions represent the P_{10}–P_{90} range of values, while the line highlights the median (P_{50}). The best results were found in the simultaneous and stochastic case, with a 13 per cent NPV improvement over the sequential stochastic case, showing the value of optimising both mines simultaneously while accounting for uncertainty. A significant reduction in value (39–40 per cent) is seen when the mine plans are optimised deterministically, neglecting uncertainty.

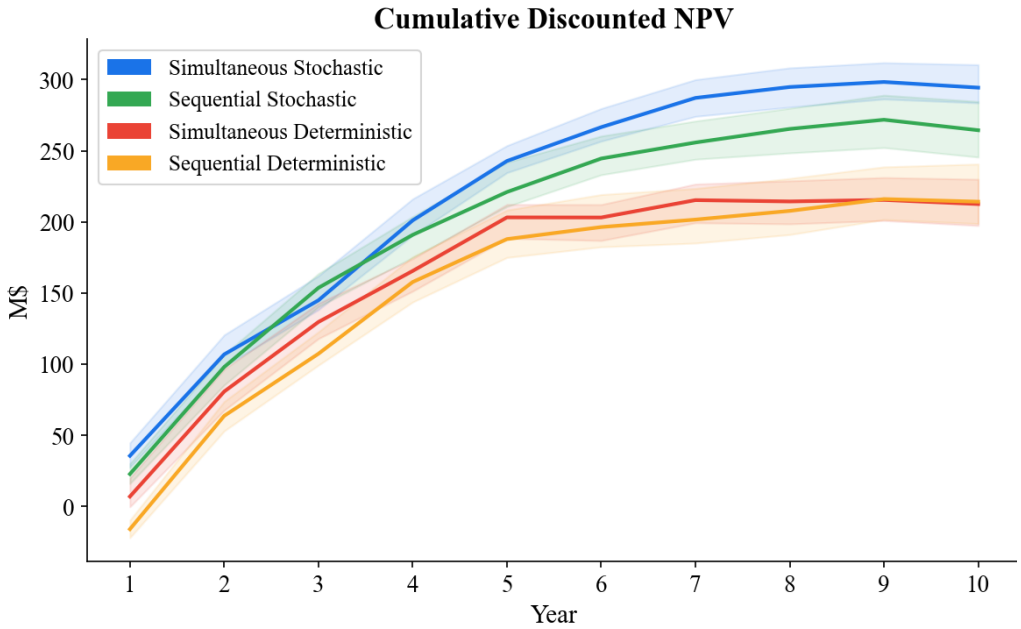


FIG 9 – NPV comparison.

Figure 10 shows similar results for cumulative recovered metal, with simultaneous and stochastic optimisation increasing total metal recovery independently and in combination.

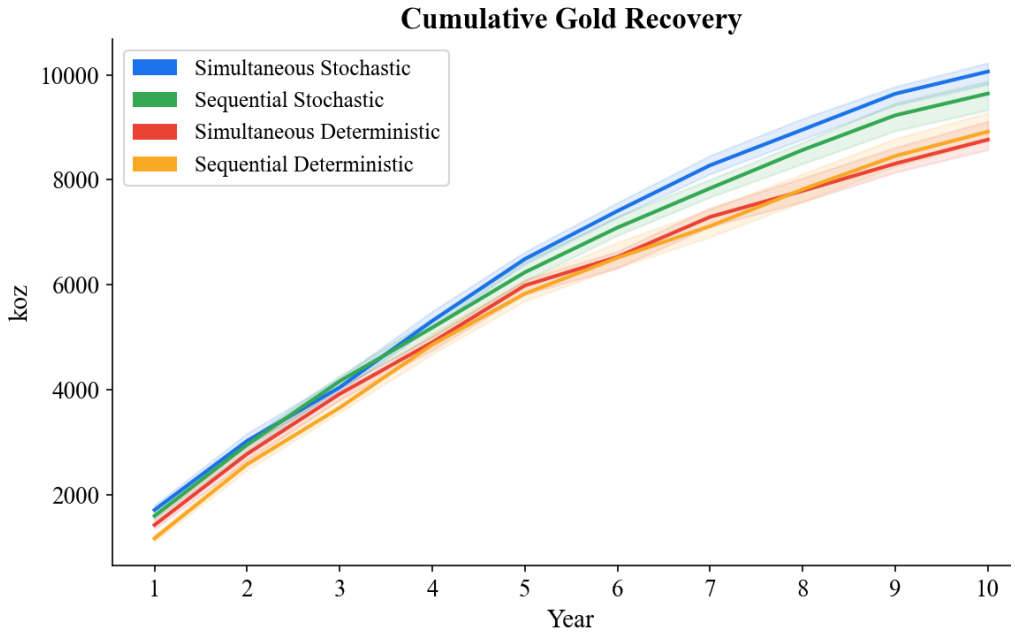


FIG 10 – Cumulative gold recovery comparison.

Figure 11 highlights the value of accounting for geological uncertainty with stochastic optimisation. In the deterministic experiments, a single orebody estimate is considered when routing material to the processor. This results in high-variance material being mischaracterised, forcing the operation to reroute material in a reactive fashion. Considering this uncertainty and variability when making extraction and routing decisions results in a much more consistent plant feed, as shown by both stochastic experiments.

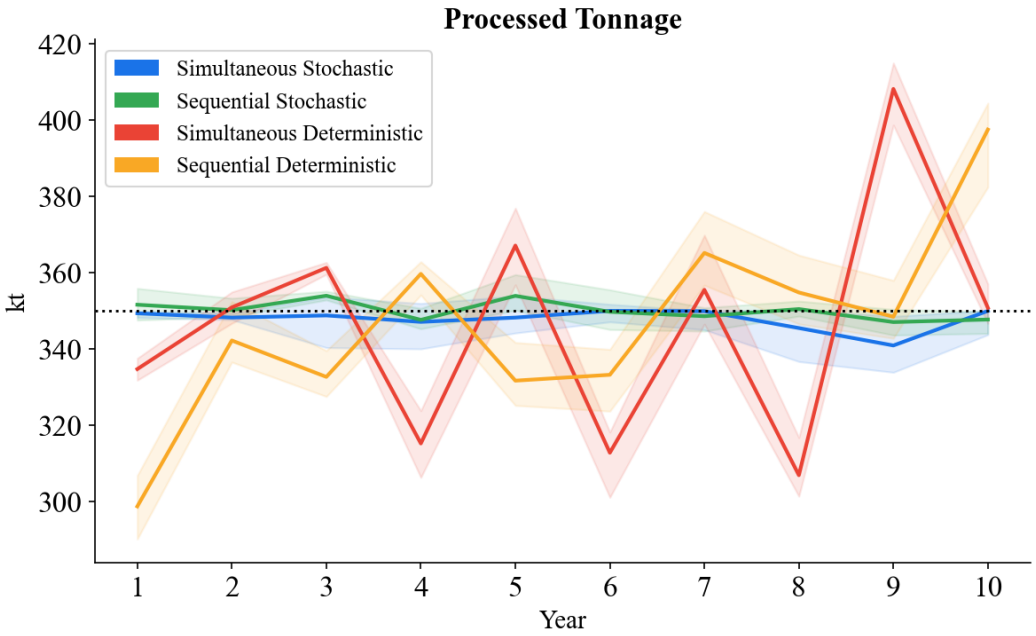


FIG 11 – Processed tonnage comparison.

Figure 12 shows the plant feed grade throughout the LOM for all experiments. Although a target range of 4–5 g/t was set, its penalty cost was quite low resulting in the model prioritising the maximisation of NPV over the plant feed target. Considering geological uncertainty results in a lower variance feed grade, while optimising the mines jointly results in a higher average grade throughout the life-of-mine.

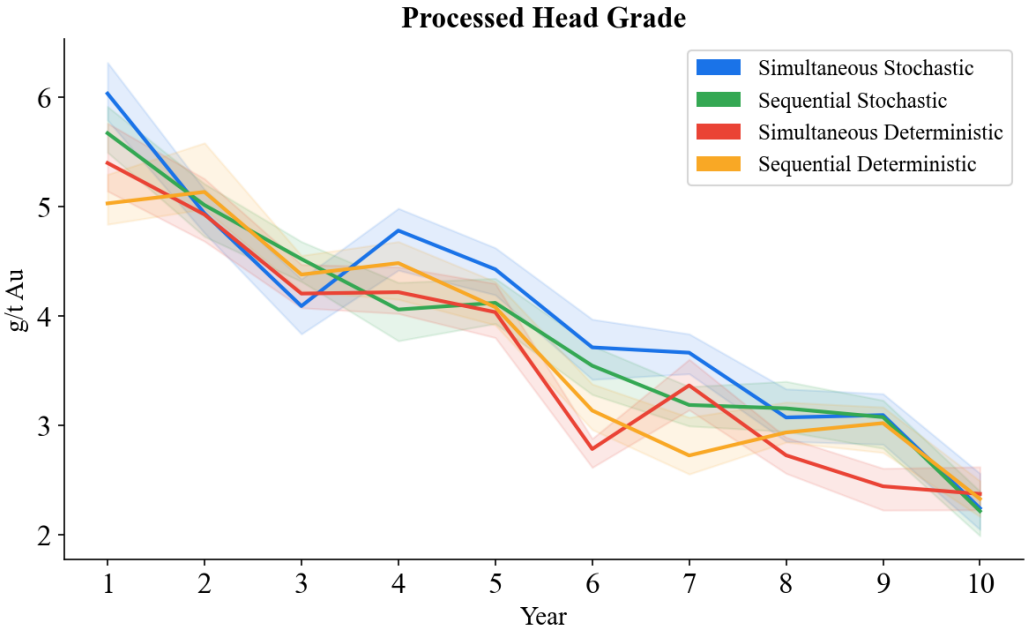


FIG 12 – processed head grade comparison.

The design and schedule determined by the simultaneous stochastic model are shown in Figure 13. The model prioritises high-grade stopes for early periods but balances this with prioritising stopes closer to the shaft to defer development costs to later periods.

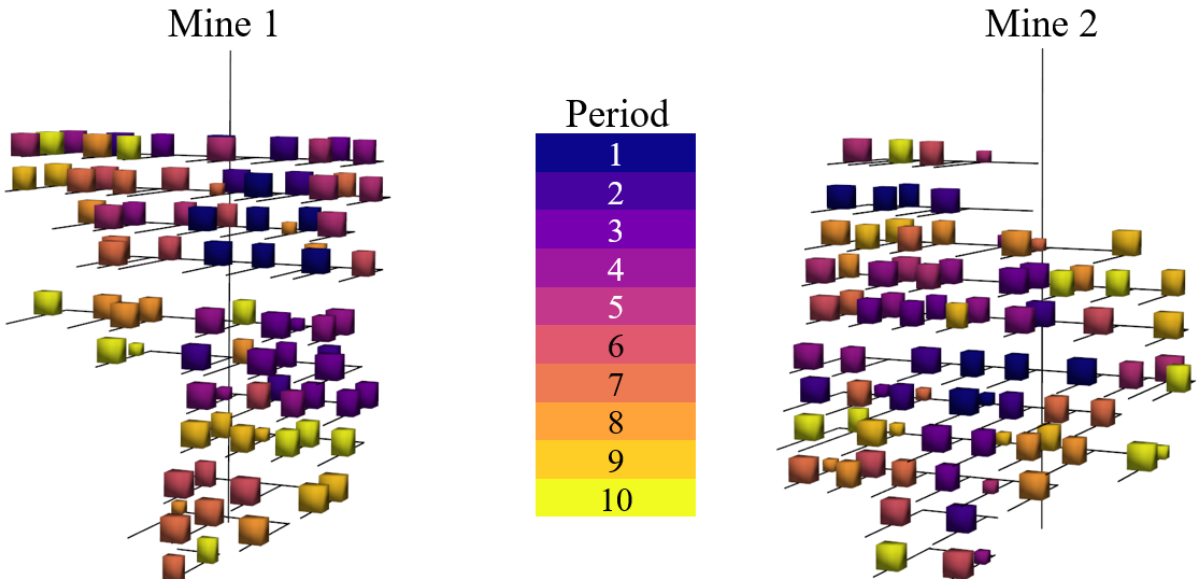


FIG 13 – Mining sequence and design for the simultaneous and stochastic approach.

CONCLUSIONS

The proposed method extends the simultaneous stochastic optimisation of mining complexes (Goodfellow and Dimitrakopoulos, 2016) to underground sublevel open stoping operations using the joint design and scheduling methodology proposed by Furtado e Faria, Dimitrakopoulos and Lopes Pinto (2022). The benefits of optimising an underground mining complex while accounting for uncertainty are seen clearly when compared to the results of sequential and deterministic optimisation. Stochastic optimisation across multiple conditionally simulated orebody realisations results in a design and schedule that is more resilient to geological uncertainty and variability (Furtado e Faria, Dimitrakopoulos and Lopes Pinto, 2022). Integrating this formulation into a mining complex model offers the ability to optimise multiple mines with shared/alternative processing streams simultaneously. This captures synergies between the operations and optimises the business unit in its entirety.

The proposed method has some shortcomings that motivate future research. Firstly, the SIP lacks stope sequencing constraints. Generalising the model to different geomechanical sequences requires a precedence model, which should be added in future research. The generation of candidate stope boundaries could also be improved by using existing algorithms to generate candidates, or by implementing continuous stope design decisions directly into the formulation. The construction of the formulation as an addition to the SSOMC model proposed by Goodfellow and Dimitrakopoulos (2016) allows for generalisation to mixed open pit and underground mining complexes. Future research should simultaneously optimise open pit and underground operations with shared processing resources. An additional extension could jointly optimise an open pit to underground transition. Another direction of research should extend the simultaneous stochastic optimisation of mining complexes to underground mines with different mining methods. Finally, different sources of uncertainty could be quantified, such as market, recovery and geomechanical uncertainties to jointly manage different sources of risk.

ACKNOWLEDGEMENTS

The work in this paper was funded by The Natural Sciences and Engineering Research Council (NSERC), the mining industry consortium members of McGill University's COSMO Stochastic Mine Planning Laboratory [AngloGold Ashanti, Agnico Eagle, BHP, De Beers, AngloAmerican, IAMGOLD, Kinross Gold, Newmont Mining, Vale Base Metals and Vale S.A.]; and the Canada Research Chairs Program.

REFERENCES

Albor Consuegra, F R and Dimitrakopoulos, R, 2009. Stochastic mine design optimisation based on simulated annealing: Pit limits, production schedules, multiple orebody scenarios and sensitivity analysis, *Transactions of the Institutions of Mining and Metallurgy*, *Section A: Mining Technology,* 118:79–90. https://doi.org/10.1179/037178409X12541250836860

Alford, C and Hall, B, 2009. Stope optimisation tools for selection of optimum cut-off grade in underground mine design, *Australasian Institute of Mining and Metallurgy Publication Series*, pp 137–144.

Alford, C, 1995. Optimization in underground mine design, in *Proceedings of APCOM XXV: Application of Computers and Operations Research in the Minerals Industries,* pp 213–218 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Alford, C, Brazil, M and Lee, D H, 2016. Optimisation in underground mining, *International Series in Operations Research and Management Science*, 99:561–577. https://doi.org/10.1007/978-0-387-71815-6_30

Ataee-Pour, M, 2004. Optimisation of stope limits using a heuristic approach, *Mining Technology,* 113:123–128. https://doi.org/10.1179/037178404225004959

Benndorf, J and Dimitrakopoulos, R, 2013. Stochastic long-term production scheduling of iron ore deposits: Integrating joint multi-element geological uncertainty, *Journal of Mining Science,* 68–81. https://doi.org/10.1134/S1062739149010097

Birge, J R and Louveaux, F, 2011. *Introduction to Stochastic Programming* (Springer: New York).

Carelos Andrade, L, Dimitrakopoulos, R and Conway, P, 2024. Integrated stochastic optimisation of stope design and long-term production scheduling at an operating underground copper mine, *International Journal of Mining*, *Reclamation and Environment,* 38:619–641. https://doi.org/10.1080/17480930.2024.2337499

Dimitrakopoulos, R and Grieco, N, 2009. Stope design and geological uncertainty: Quantification of risk in conventional designs and a probabilistic alternative, *Journal of Mining Science,* 45:152–163. https://doi.org/10.1007/S10913-009-0020-Y

Dimitrakopoulos, R and Ramazan, S, 2008. Stochastic integer programming for optimising long term production schedules of open pit mines: methods, application and value of stochastic solutions, *Mining Technology,* 117:155–160. https://doi.org/10.1179/174328609X417279

Dimitrakopoulos, R, 2011. Stochastic optimization for strategic mine planning: A decade of developments, *Journal of Mining Science*, 47:138–150. https://doi.org/10.1134/S1062739147020018

Feo, T A and Resende, M G C, 1995. Greedy Randomized Adaptive Search Procedures, *Journal of Global Optimization*, 6:109–133. https://doi.org/10.1007/BF01096763

Furtado e Faria, M, Dimitrakopoulos, R and Lopes Pinto, C L, 2022. Integrated stochastic optimization of stope design and long-term underground mine production scheduling, *Resources Policy,* 78:102918. https://doi.org/10.1016/JR.ESOURPOL.2022.102918

Goodfellow, R C and Dimitrakopoulos, R, 2016. Global optimization of open pit mining complexes with uncertainty, *Applied Soft Computing,* 40:292–304. https://doi.org/10.1016/JA.SOC.2015.11.038

Goovaerts, P, 1997. *Geostatistics for Natural Resources Evaluation* (Oxford University Press).

Jiang, Y and Dimitrakopoulos, R, 2025. An application of simultaneous stochastic optimisation on an open-pit copper mining complex with supply, recovery and market uncertainties, *International Journal of Mining*, *Reclamation and Environment*, 39:74–92. https://doi.org/10.1080/17480930.2024.2381904

Lamghari, A and Dimitrakopoulos, R, 2012. A diversified Tabu search approach for the open-pit mine production scheduling problem with metal uncertainty, *European Journal of Operational Research*, 222:642–652. https://doi.org/10.1016/je.jor.2012.05.029

Lamghari, A, Dimitrakopoulos, R and Ferland, J A, 2014. A variable neighbourhood descent algorithm for the open-pit mine production scheduling problem with metal uncertainty, *Journal of the Operational Research Society,* 65:1305–1314. https://doi.org/10.1057/jors.2013.81

Leite, A and Dimitrakopoulos, R, 2007. Stochastic optimisation model for open pit mine planning: application and risk analysis at copper deposit, *Mining Technology,* 116:109–118. https://doi.org/10.1179/174328607X228848

Montiel, L and Dimitrakopoulos, R, 2017. A heuristic approach for the stochastic optimization of mine production schedules, *Journal of Heuristics*, pp 397–415. https://doi.org/10.1007/S10732-017-9349-6

Ramazan, S and Dimitrakopoulos, R, 2013. Production scheduling with uncertain supply: A new solution to the open pit mining problem, *Optimization and Engineering*, 14:361–380. https://doi.org/10.1007/s11081-012-9186-2

Ramazan, S and Dimitrakopoulos, R, 2018. Stochastic optimisation of long-term production scheduling for open pit mines with a new integer programming formulation, *Advances in Applied Strategic Mine Planning*, pp 139–153. https://doi.org/10.1007/978-3-319-69320-0_11

Rossi, M E and Deutsch, C V, 2014. Recoverable Resources: Simulation, in *Mineral Resource Estimation* (eds: M E Rossi and C V Deutsch), pp 167–192 (Springer Netherlands: Dordrecht).

Sari, Y A and Kumral, M, 2021. Sublevel stope layout planning through a greedy heuristic approach based on dynamic programming, *Journal of the Operational Research Society,* 72:554–563. https://doi.org/10.1080/01605682.2019.1700179

Sotoudeh, F, Nehring, M, Kizil, M, Knights, P and Mousavi, A, 2020. Production scheduling optimisation for sublevel stoping mines using mathematical programming: A review of literature and future directions, *Resources Policy,* 68:101809. https://doi.org/10.1016/j.resourpol.2020.101809

Topal, E and Sens, J, 2010. A new algorithm for stope boundary optimization, *J Coal Sci Eng China,* 16:113–119. https://doi.org/10.1007/s12404-010-0201-y

Villaescusa, E, 2014. *Geotechnical Design for Sublevel Open Stoping*, 1st ed (CRC Press). https://doi.org/10.1201/b16702

A stochastic economic block-model framework for optimising underground stope layouts under grade uncertainty for a platinum-group-element deposit

A S Nhleko[1], M I Mabala[2] and C Musingwini[3]

1. Director, Centre for Sustainable Mining, North-West University, Potchefstroom, North West 2520, South Africa. Email: sihe.nhleko@nwu.ac.za
2. Lecturer, School of Mining Engineering, University of the Witwatersrand, Johannesburg, Gauteng 2001, South Africa. Email: isaac.mabala@wits.ac.za
3. Professor, School of Mining Engineering, University of the Witwatersrand, Johannesburg, Gauteng 2001, South Africa. Email: cuthbert.musingwini@wits.ac.za

ABSTRACT

Underground stope-layout optimisation is usually carried out on a single, smoothed estimate of grade, typically an ordinary-kriging block model. Grade is, however, uncertain, and that uncertainty can be represented by an ensemble of equally likely conditional simulations. This paper asks a direct question: when, and by how much, does optimising the stope layout on such an ensemble create more economic value than optimising on one kriged estimate? We build a known synthetic reference panel calibrated to the Platreef, a platinum-group-element deposit, sample it with drill holes, and from those samples build both an ordinary-kriging model and a 50-realisation sequential-Gaussian-simulation ensemble. Every candidate stope layout is valued on the known reference, so the economic value captured is measured rather than assumed. Using economic parameters from the Platreef 2025 Integrated Development Plan with a break-even grade of 1.64 g/t 3PGE+Au (platinum, palladium and rhodium plus gold), three results emerge. Firstly, for a linear mining decision in which the entire stope is mined, the simulation mean and the kriged estimates are close enough, indicating no material advantage. For the linear objective, the optimal decision depends only on the conditional expected value, which kriging already estimates. Secondly, for a non-linear decision in which the stope is developed and sent to either the mill or waste dump according to its bulk grade, optimising on the ensemble recovers up to 22.5 per cent more realised value at selective cut-offs under a develop-and-destination decision, because it credits marginal stopes with a real chance of being ore; the gain is positive across the ensemble at cut-offs up to 4.0 g/t, not only on the reference. Lastly, block (selective mining unit) support matters: coarsening the model lowers grade dispersion and the achievable value through the volume-variance effect, and the economically optimal stope size is governed by development cost rather than by geology alone. When the risk of the entire stope layout is optimised globally rather than stope-by-stope, the risk-return frontier is modest, because stope-level risks diversify across the layout. This paper extends a previous deterministic dual-interchange-algorithm study by incorporating grade uncertainty and complements a stochastic stope-design framework by incorporating development costs.

INTRODUCTION

Converting a geological block model into a mine plan requires a sequence of decisions, and one of the most consequential decisions in underground mining is the stope layout. Since a stope is defined by many blocks and many geometrical and geotechnical constraints, many feasible layouts can be generated, and the problem is known to be difficult to solve to true optimality (Nhleko, Tholana and Neingo, 2018). Several algorithms that include the floating-stope method embodied in some commercial mine planning software and metaheuristics-based methods, have been developed to solve the problem. For example, Nhleko (2022), Nhleko and Musingwini (2022) and Nhleko, Musingwini and Mabala (2023) presented a dual-interchange-algorithm (DIA) that combines particle-swarm and genetic-algorithm operators to optimise three-dimensional stope layouts. These studies showed that by allowing the stope width to vary, mining can follow irregular ore boundaries and higher stope-layout economic value (SLEV) can be generated than in a fixed stope width scenario. However, the DIA approach optimises on a single, deterministic block model and does not represent the reality that the grade of every block is uncertain.

Due to grade uncertainty, a single average grade per block can misclassify ore and waste, and that misclassification propagates through block economic value into the stope layout. Conditional simulation can be used to overcome this challenge because it produces many equally likely models that together describe the uncertainty instead of one smoothed grade model (Goovaerts, 1997; Deutsch and Journel, 1998). Optimising a decision over such an ensemble, stochastic optimisation, has repeatedly been shown to add value in strategic mine planning (Dimitrakopoulos, 2011). For stope design specifically, Faria, Dimitrakopoulos and Pinto (2022) optimised stope boundaries under grade uncertainty with development costs and demonstrated increase in value over conventional mine design and planning practices.

This paper is not presenting a new optimiser but seeks to answer a narrower and, we argue, more useful question: when, and by how much, does carrying grade uncertainty into the stope layout actually create value, and how does that depend on the reporting cut-off, on the block (selective mining unit, SMU) support, and on the stope geometry? We answer it on a known synthetic reference, so that the value each method truly captures can be measured rather than inferred, and we are careful to separate the cases in which uncertainty helps from those in which it does not. In doing so, the paper extends the deterministic DIA stope layout work by incorporating grade uncertainty and complements the work of Faria, Dimitrakopoulos and Pinto (2022) on a stochastic stope-design framework by incorporating development costs.

METHODOLOGY

A known reference stope panel

The analysis is based on a synthetic stope panel whose true grades are known. One unconditional sequential-Gaussian simulation (SGS), mapped through a lognormal anamorphosis, serves as the reference, the synthetic ‘ground truth’. It reproduces the grade tenor and spatial continuity of the Platreef, a platinum-group-element (PGE) deposit and is discretised into 128 × 35 × 11 = 49 280 blocks of 5 m (a 640 × 175 × 55 m panel), which is the same test bed used in the DIA stope layout work. Since the reference embeds exactly the continuity model used downstream, any departure in an estimator is attributable to estimation, not to model misspecification. The realised reference has a mean grade of 3.74 g/t 3PGE+Au (platinum, palladium and rhodium plus gold) and a coefficient of variation of 0.55; the anamorphosis uses lognormal parameters $\mu = 1.184$ and $\sigma = 0.55$ in log space, which give a model mean of 3.80 g/t and a coefficient of variation of 0.59, so 0.55 is the realised value.

Two grade models from the same samples

The deposit analogue is the Flatreef, which is the thick, flat-lying, down-dip portion of the Platreef on the Northern Limb of the Bushveld Complex, that is planned to be mined using mechanised longhole (sublevel) open stoping (Grobler *et al,* 2019; Maier *et al,* 2023). A geological block model of 640 m × 175 m × 55 m is discretised into 49 280 blocks of 5 m × 5 m × 5 m on a 128 × 35 × 11 block grid (Figure 1). The reference grade field is generated by unconditional SGS of a standard Gaussian variable with a normal-score variogram comprising a nugget of 0.25 and a spherical structure of sill 0.75, with ranges of 150 m, 80 m and 12 m along the major, minor and vertical axes respectively, and the major axis oriented at an azimuth of 137° consistent with the published continuity of the deposit (OreWin, 2017). The Gaussian field is transformed to grade by lognormal anamorphosis calibrated to the published Flatreef grade, giving a mean of 3.74 g/t 3PGE+Au and a coefficient of variation of 0.55 (Kekana *et al,* 2017). Nickel and copper grades are carried as fixed ratios to the 3PGE+Au grade (0.078 and 0.038, respectively), consistent with the published base-metal tenor and the adopted economic model (Kekana *et al,* 2017; Nhleko and Musingwini, 2022). The lognormal form is appropriate because precious-metal grades are characteristically positively skewed (Journel and Huijbregts, 1978), and the monotone anamorphosis preserves the spatial structure imposed in Gaussian space, so the comparison between estimation methods is conducted on a structurally consistent reference.

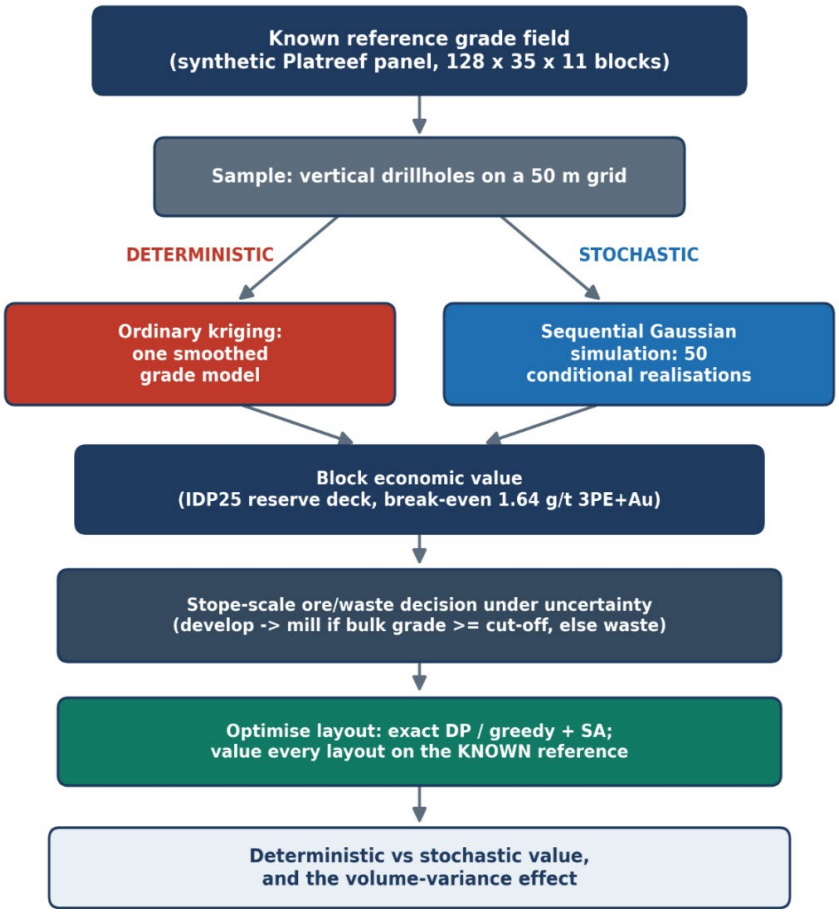


FIG 1 – The known-reference framework. A synthetic reference grade field is sampled by drill holes; from the samples an ordinary-kriging (OK) model and a 50 member simulation ensemble are built. Both feed a stope-scale ore/waste decision, and every resulting layout is valued on the known reference. Layouts are optimised by exact dynamic programming (DP) or a greedy rule, with simulated annealing (SA) for non-separable objectives.

The reference is sampled by vertical drill holes on a 50 m × 50 m plan grid, 13 holes along the 640 m panel and four across it (52 holes), each intersecting the full 55 m height and yielding 11 × 5 m composites, 572 in all. A composite takes the reference value of the block it occupies, so sampling error and support change are excluded and every difference between the grade models arises from estimation alone; the pattern is dense relative to the roughly 100 m spacing that supports Indicated Resources at Platreef (Kekana *et al,* 2017). From the composites an ordinary-kriging model is estimated in grade units at every 5 m block centre, using the reference variogram (nugget 0.25, spherical structure 0.75, ranges 150 m, 80 m and 12 m, major axis at 137°) rescaled to the composite grade variance, a moving neighbourhood of the 40 nearest composites in anisotropy-scaled space and the ordinary-kriging unbiasedness constraint. The 50-realisation ensemble is generated by sequential Gaussian simulation on the same block grid. Composites are transformed to normal scores through an empirical table with the GSLIB tail models of Deutsch and Journel (1998) (power lower tail, exponent 1.5, to 0 g/t; hyperbolic upper tail, exponent 1.5); each realisation visits the 48 708 unsampled nodes along a random path, solves at each node a simple-kriging system with zero mean and the reference normal-score variogram over up to 40 informed nodes drawn from the 160 nearest in anisotropy-scaled space, draws from the resulting conditional Gaussian distribution, honours the composites exactly, and is back-transformed through the same table. The reference and the 50 realisations use independent random streams spawned from the master seed in Table 1. The E-type (realisation-mean) grade is used only as the diagnostic of Table 2; every stochastic decision uses the full ensemble.

TABLE 1

Economic parameters from the IDP25 Mineral Reserve price deck and site operating-cost basis; the 2.0 g/t 3PGE+Au figure is used only as a Mineral Resource reporting/selectivity cut-off.

Parameter	Value
In situ density	3.1 t/m^3
Prill split (Pt/Pd/Rh/Au)	0.444 / 0.456 / 0.031 / 0.069
Metal prices (Pt/Pd/Rh/Au)	US$1600 / 815 / 1500 / 1300 per oz
Blended 3PGE+Au basket	US$1218 per oz
3PGE+Au recovery to payable	0.70 (0.85 concentrator × 0.82 payable)
Nickel credit	0.078% Ni per g/t; US$19 621/t; 0.50 payable
Copper credit	0.038% Cu per g/t; US$6614/t; 0.64 payable
Unit operating cost	US$60.09/t (site); mining component US$34.27/t
Break-even grade	1.64 g/t 3PGE+Au
Base-case reporting cut-off	2.0 g/t 3PGE+Au
Realisations / master seed	50 / 20260629

TABLE 2

Closeness of the ordinary-kriging model and the simulation E-type (924 candidate stopes).

Diagnostic	Value
Grade correlation (OK versus E-type)	0.87
Grade root-mean-square difference	0.63 g/t
Stope SEV correlation	0.92
Stope SEV root-mean-square difference	US$0.1 M
Stope eligibility mismatch at 4.0 g/t	105 of 924

Economic block value

Each block is assigned a block economic value (BEV) from its grade, its by-product nickel and copper credits, and the unit costs and recoveries, as in Equation 1. To keep the study current and defensible, the economic parameters are taken from the Platreef 2025 Integrated Development Plan (IDP25; Ivanhoe Mines, 2025) as listed in Table 1. The resulting break-even grade is 1.64 g/t 3PGE+Au. The 2.0 g/t 3PGE+Au figure used as the base-case reporting cut-off is the published base-case Mineral Resource cut-off of the plan; it sits above break-even, consistent with the plan's description of an elevated cut-off above the marginal economic cut-off. Mineral Reserves in the plan are reported on a declining net-smelter-return cut-off (US$155/t to US$80/t), not a grade cut-off, so the grade cut-off used here is a reporting and selectivity device, not a reserve statement.

$$\mathrm{BEV}(i,r) = t_i \left[g(i,r) \left(\frac{P_b R_b}{31.1035} + \frac{R_{Ni} P_{Ni} f_{Ni}}{100} + \frac{R_{Cu} P_{Cu} f_{Cu}}{100} \right) - C_{op} \right] \quad (1)$$

where:

- t_i = tonnes in block i (block volume × *in situ* density)
- g(i, r) = 3PGE+Au grade of block i in realisation r (g/t)
- Pb = blended 3PGE+Au basket price (US$/oz)
- Rb = 3PGE+Au recovery to payable
- 31.1035 = grams per troy ounce

r_{Ni}, r_{Cu} = nickel and copper grade ratios (% metal per g/t 3PGE+Au)

P_{Ni}, P_{Cu} = metal prices (US$/t)

f_{Ni}, f_{Cu} = payable fractions

C_{op} = unit site operating cost (US$/t)

With the Table 1 values the revenue term is US$36.6/t per g/t, so the break-even grade is 60.09/36.6 = 1.64 g/t 3PGE+Au; BEV is linear in grade and may be negative, consistent with the entire-envelope accounting of the next section.

The stope decision under uncertainty

How a stope is valued is decisive, and it is where the DIA stope layout study has a shortcoming. A sublevel open stope is mined in bulk: the whole envelope is blasted, so internal sub-cut-off material is mined and paid for and cannot be left behind selectively. We therefore value a mined stope as the sum of the economic value of every block in it, dilution included. Formally, for a candidate stope, s, and realisation, r, its stope economic value (SEV) and mean grade are given by Equations 2 and 3:

$$\text{SEV}(s,r) = \sum_{i \in s} \text{BEV}\,(i,r) \tag{2}$$

$$\bar{g}(s,r) = \frac{1}{n_s} \sum_{i \in s} g\,(i,r) \tag{3}$$

Under this entire-envelope accounting e stope value is linear in grade. The natural non-linearity for bulk stoping is the ore/waste destination: a stope is developed and mined; at production its bulk grade decides its destination. Its value under realisation, r, is given by Equations 4 and 5:

$$V(s,r;c) = \text{SEV}(s,r) - D_s \quad \text{if } \bar{g}(s,r) \geq c \quad \text{(milled)} \tag{4}$$

$$V(s,r;c) = -C_{\text{mine}}(s) - D_s \quad \text{if } \bar{g}(s,r) < c \quad \text{(waste)} \tag{5}$$

where:

c = cut-off grade (g/t 3PGE+Au)

D_s = per-stope development cost (US$)

$C_{mine}(s)$ = mining cost of stope s, incurred whether it is milled or classified as waste (US$)

n_s = number of blocks in stope s

The mining cost component of the operating cost in Table 1 is US$34.27/t; the processing cost component is spent only when the stope is milled and is already contained in the SEV. This pay-off is kinked, so the expected value of developing a stope is not the value based on its mean grade. The deterministic layout develops a stope when, V, evaluated on the kriged estimate is positive; the stochastic layout develops it when its expected value across the 50 realisations, $(1/50)\Sigma_r$ V(s, r; c), is positive. Every layout is finally valued on the known reference.

Solvers, and why not the dual interchange algorithm

Two meanings of ‘stochastic’ must be kept apart. Particle-swarm, genetic and simulated-annealing algorithms are all stochastic in the sense that they search using randomness; the DIA is such an algorithm, and so is simulated annealing (Kirkpatrick, Gelatt and Vecchi, 1983). That is different from stochastic optimisation in the mining sense, which means optimising under geological uncertainty. The published DIA optimises a single deterministic model and so is not a stochastic optimiser in that second sense, but neither is simulated annealing on its own. What makes the present study stochastic is that the objective is evaluated over the ensemble; the choice of solver is not the object of the study.

Therefore, we do not use the DIA here, not because it is inferior, but because the work presented in this paper is on finding the difference in value between the deterministic and stochastic layouts, and that difference can only be trusted if the optimum can be trusted. We solve exactly wherever the structure allows: the separable expected-value problem by a greedy rule and the variable-geometry problem by dynamic programming. Simulated annealing is used for the non-separable objectives

that have no closed form. In this paper these are the global percentile objectives of the risk analysis (see ‘Risk of the whole layout’), in which the 10th or 90th percentile of the entire-layout value is maximised; the annealing solution is checked against the exact optimum wherever one exists. A metaheuristic method cannot guarantee optimality and returns different answers from iterative runs, which is acceptable for a design tool but a liability when the research question is a value gap of a few per cent.

The optimisation workflow is as follows. Candidate stopes are generated by a disjoint tiling of the panel with a base-case stope of 3 × 3 × 4 blocks (15 m along strike, 15 m wide and 20 m high) in two 20 m sublevels above a 15 m sill at the base of the panel, giving 42 × 11 × 2 = 924 non-overlapping candidates. Because the candidates do not overlap, the develop/do-not-develop decision is separable and the expected-value layout is obtained exactly by a greedy rule: every candidate whose ensemble-mean value $(1/50)\Sigma_r$ V(s, r; c) is positive is developed, and the deterministic layout applies the same rule to V evaluated on the kriged model. For the variable-geometry problem the width (15 m) and height (20 m) are fixed and the strike length is free: each strike line is solved by dynamic programming over block positions along strike, in which the state is the position reached, the decision is to leave the next block column unmined or to end an ore-eligible stope of 1 to 12 blocks (5 to 60 m) at that position, and the recursion maximises cumulative whole-envelope value, which guarantees the optimal non-overlapping packing of every line. In the stope-size analysis the width is varied from two to seven blocks (10 to 35 m) at fixed strike length and height, and results are averaged over three strike offsets and two width offsets of the tiling. For the non-separable percentile objectives, simulated annealing starts from an empty layout and proposes single-stope flips within the ore-eligible pool, always accepting an improving move and accepting a worsening move with probability exp(Δ/T); the initial temperature is US$5.0 M, the temperature is multiplied by 0.95 every 400 moves, the run stops after 40 000 moves, the best layout visited is retained, and the random stream is seeded from the master seed. The annealing solution is checked against the exact greedy or dynamic-programming optimum wherever the objective admits one.

RESULTS

Block-level validation

Before any layout is optimised, it is worth seeing what the two grade models do at the block scale (Figure 2). Kriging compresses the grade distribution: its standard deviation is 1.26 g/t against the reference’s 2.07 g/t, while the simulations restore it at 2.44 g/t. At the 2.0 g/t cut-off this smoothing over-calls ore, kriging labels 95.8 per cent of the panel as ore against the truth’s 82.5 per cent, and 17.4 per cent of blocks are misclassified as false ore (ie 7575 blocks) rather than false waste (ie 1006 blocks). However, the recoverable *in situ* value at that cut-off is barely affected: kriging gives US$1456 M against the truth of US$1482 M, and the ensemble encloses the truth (10th to 90th percentiles, P10–P90, of US$1386–1522 M). On this dense, high-grade reef almost all values lie well above cut-off, so misclassifying near-threshold blocks increases tonnes but generates little value.

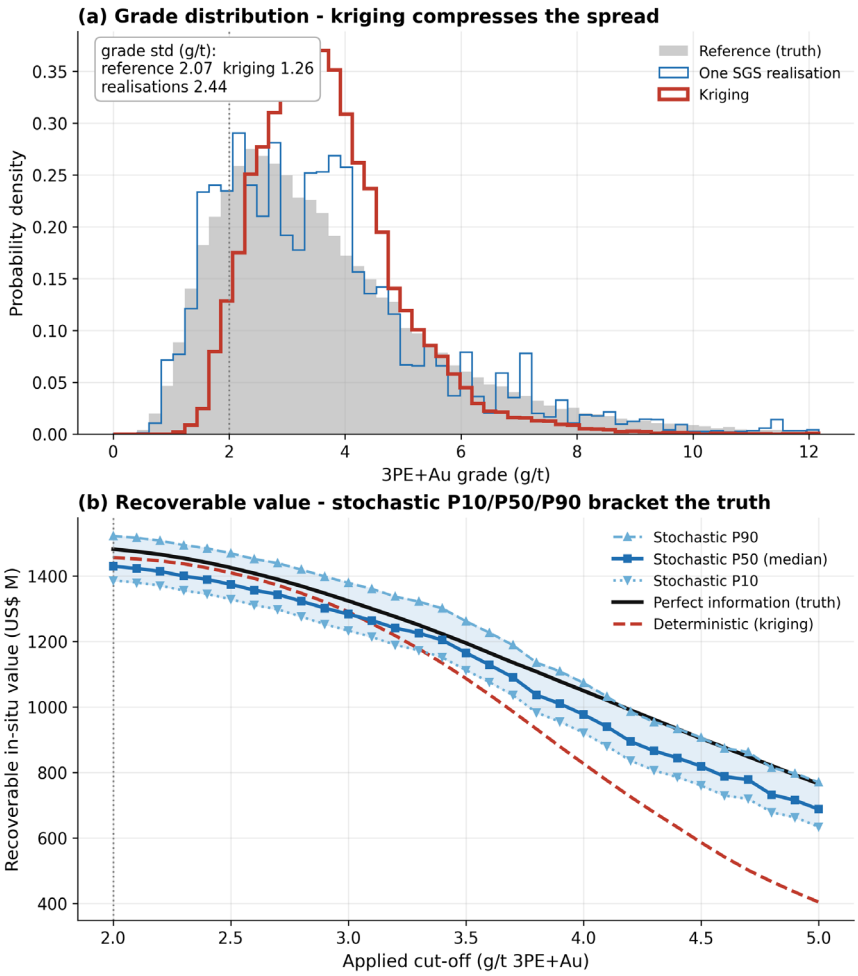


FIG 2 – Block-level validation: (a) Kriging compresses the grade distribution while the simulations reproduce it; (b) Recoverable *in situ* value against cut-off: the kriged value tracks the reference at the base cut-off and the stochastic P10, P50 and P90 lines enclose the truth.

The value of stochastic optimisation

Figure 3 is the central result. Figure 3a plots, against cut-off raised from the base 2.0 g/t to a selective 5.0 g/t, the value that the perfect-information, deterministic and stochastic layouts realise on the known reference. Two different quantities are shown for the stochastic layout and must not be confused: its realised value on the reference (a single line, directly comparable with the deterministic line), and the P10, P50 and P90 of the value it assigns to itself across the 50 equally likely realisations (the shaded band), which shows the uncertainty the layout carries rather than its performance. Figure 3b is derived from Figure 3a by subtraction: for each cut-off it plots the distribution across the 50 realisations of (stochastic layout value − deterministic layout value), together with the same difference realised on the reference, which is the vertical gap between the two realised lines in Figure 3a. At the base cut-off the three layouts coincide: the reef is mostly ore, almost every stope is developed, and there is nothing for a better grade model to decide. As the cut-off rises and the ore/waste destination becomes contested, the stochastic layout is better, reaching a peak advantage of 22.5 per cent at 4.0 g/t. Figure 3b shows that this advantage is not an accident of the particular reference: across the 50 realisations the stochastic gain over the deterministic layout is positive at every cut-off from 3.0 to 4.0 g/t (P10 of +US$10 M at 4.0 g/t, median +US$41 M), and the realised-on-reference gain of +US$77 M (US$420 M against US$343 M) sits inside that range. At 4.5 g/t and above, where few stopes remain economic, the gain collapses towards zero and its P10 turns slightly negative.

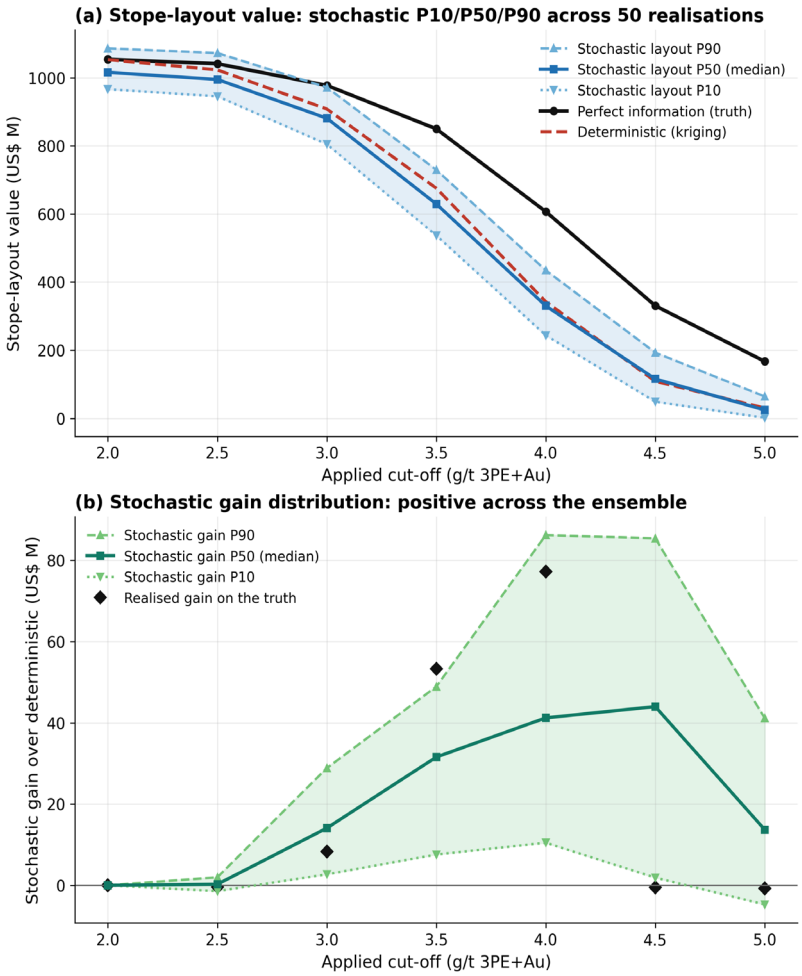


FIG 3 – Stochastic layout value and uncertainty: (a) Value realised on the reference by the perfect-information, deterministic and stochastic layouts, with the stochastic layout's self-assessed P10, P50 and P90 across the 50 realisations shown as a band; (b) Distribution of the stochastic gain over the deterministic layout across the 50 realisations (P10, P50 and P90 lines), with the realised-on-reference gain overlaid; the gain is positive across the ensemble between 3.0 and 4.0 g/t.

It is important to be more explicit about the boundary of this result, because it corrects a natural intuition. Under the linear entire-envelope accounting alone, with no ore/waste destination, the appropriate input to an expected-value decision is the conditional expected block value, which ordinary kriging already estimates. The simulation Expected-value type (E-type) is close enough to the kriged benchmark that no material advantage appears: over the whole panel the two grade models correlate at 0.87 (root-mean-square difference 0.63 g/t), and at the stope scale their economic values correlate at 0.92 with a root-mean-square difference of only about US$0.1 M (Table 2). This is not a claim that the simulation mean equals ordinary kriging after a non-linear back-transformation; it is the weaker and correct statement that, for a decision depending only on the conditional mean, the two are sufficiently close that carrying the ensemble brings no additional value. What matters is that stochastic optimisation is beneficial only when the objective is non-linear, as it is in the ore/waste destination decision.

Block support and the volume-variance effect

The block is the SMU at which grade is modelled and classified, and its size changes what the model can see (Figure 4). To isolate the support effect from any boundary effect, all supports are re-blocked from a single common reference volume of 128 × 32 × 8 fine blocks (640 × 160 × 40 m). Averaging coarser supports lowers the grade dispersion from 2.09 g/t at 5 m to 0.85 g/t at 20 m, thus, demonstrating the classical volume-variance effect. That smoothing has an economic price: the recoverable value at a selective 4.0 g/t cut-off decreases from about US$693 M at 5 m support to US$485 M at 20 m support, because a coarse model can no longer separate high-grade material

from its surroundings. The stochastic P10–P90 band lies below the perfect-information value at every support and its median sits below the kriged estimate (about US$435 M against US$503 M at 5 m), so the ensemble does not recover the value that a single smoothed estimate recovers at the block scale. At block classification, then, the single smoothed estimate outperforms any individual realisation, because each realisation carries the full conditional variance as noise. The value of carrying uncertainty is therefore a stope-scale effect, arising from the ore/waste decision of 'The value of stochastic optimisation', not a block-classification one.

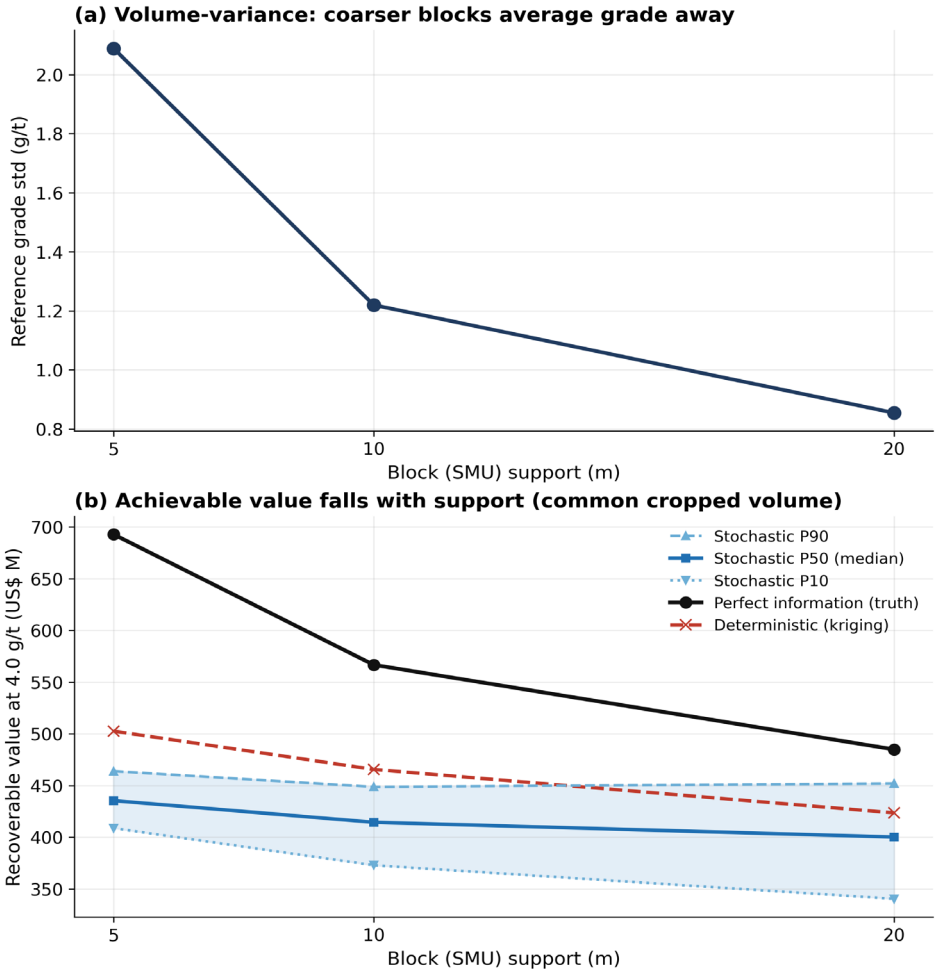


FIG 4 – Block (SMU) support, on a common cropped reference volume: (a) Grade dispersion decreases as blocks coarsen (volume-variance effect); (b) Recoverable value at 4.0 g/t decreases with support; the stochastic P10, P50 and P90 lines lie below the perfect-information value, and the kriged estimate lies above the band at 5 m and 10 m and within it at 20 m; at block classification the smoothed estimate outperforms any single realisation.

Stope size and development cost

If a stope can be any size, which size is best? Figure 5 answers this at the selective 4.0 g/t cut-off, averaged over tiling positions so that the result reflects geometry rather than how a particular tiling falls on the ore. In this idealised tiling analysis with fixed stope height and a simplified per-stope development cost, stope width is nearly value-neutral before development cost, wider and narrower stopes capture essentially the same value, because in bulk mining the same ore is extracted, and shifts toward larger stopes once a fixed per-stope cost is applied, as larger stopes spread that cost over more tonnes. The same conclusion is reached by Faria, Dimitrakopoulos and Pinto (2022): the economically optimal stope size is set by the development-cost structure rather than by geology alone. This is an isolated lever, not a general rule; real stope-size selection also depends on dilution and extraction ratio, geotechnical stability, slot-raise and drilling layout, sequencing, equipment productivity, recovery and backfill exposure, and local grade continuity. Across all stope widths the

stochastic layout realises more on the reference than the deterministic one: US$403–420 M against US$321–350 M without development cost, and US$268–372 M against US$248–316 M with it.

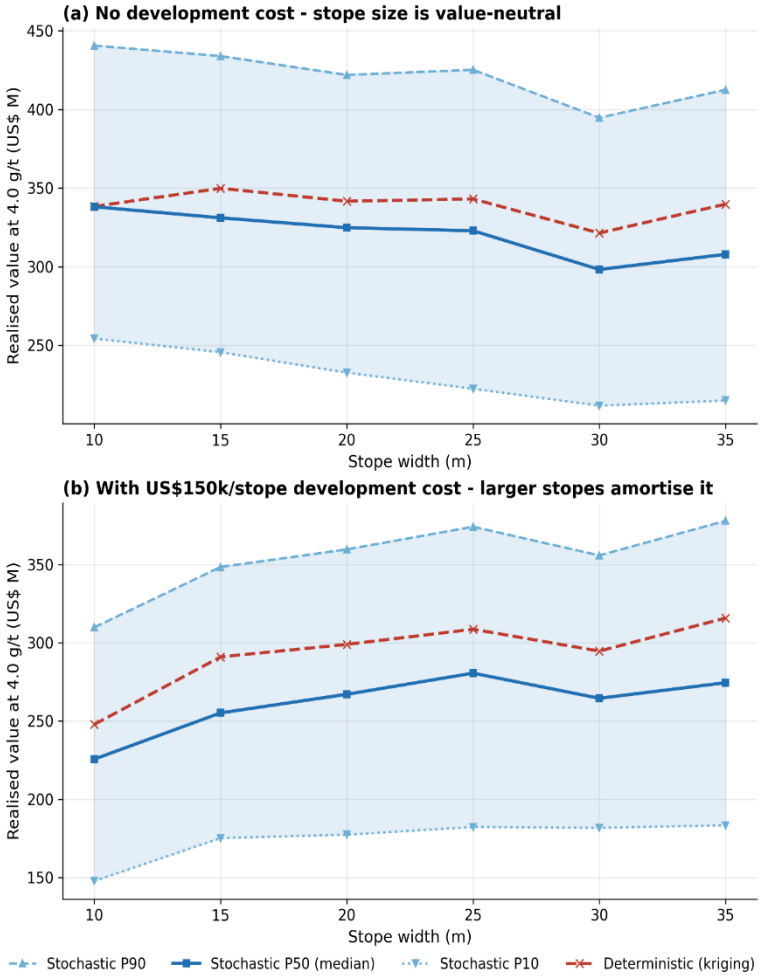


FIG 5 – Stope size and development cost, at 4.0 g/t (averaged over tiling positions). Each panel shows the value realised on the reference by the deterministic (kriging) and stochastic layouts, with the stochastic layout's self-assessed P10, P50 and P90 across the 50 realisations shown as a band. (a) Without development cost, stope size is value-neutral. (b) With US$150 000 per stope for access development, larger bulk stopes amortise the fixed cost.

Risk of the whole layout

Since the ore/waste decision is non-linear, the ensemble can in principle support different risk appetites. The right way to explore this is to optimise the risk of the whole layout, not to screen stopes one at a time: we maximise the 10th percentile (downside-averse), the mean (expected value) or the 90th percentile (upside-seeking) of the total layout value across the 50 realisations by simulated annealing. Optimised globally, the frontier is modest (Figure 6, Table 3). At 4.0 g/t the downside-averse layout realises US$368 M against the expected-value layout's US$420 M, and the three layouts share most of their stopes. Overlap is measured by the Jaccard similarity, the number of stopes developed in both layouts divided by the number developed in either, so that 1 denotes identical layouts and 0 no stope in common; it is 0.63 between the downside-averse and upside-seeking layouts. The reason is diversification: stope-level risks partly cancel across a layout of several hundred stopes, so the downside of the whole layout is far less sensitive to risk appetite than the downside of any single stope. A per-stope screening rule, which would develop a stope only if it were positive even in pessimistic realisations, suggests a much wider frontier, but that is an artefact of ignoring diversification. The genuine message is that carrying uncertainty improves the expected-value decision ('The value of stochastic optimisation') more than it opens a wide risk-return trade-off on this dense, well-diversified reef.

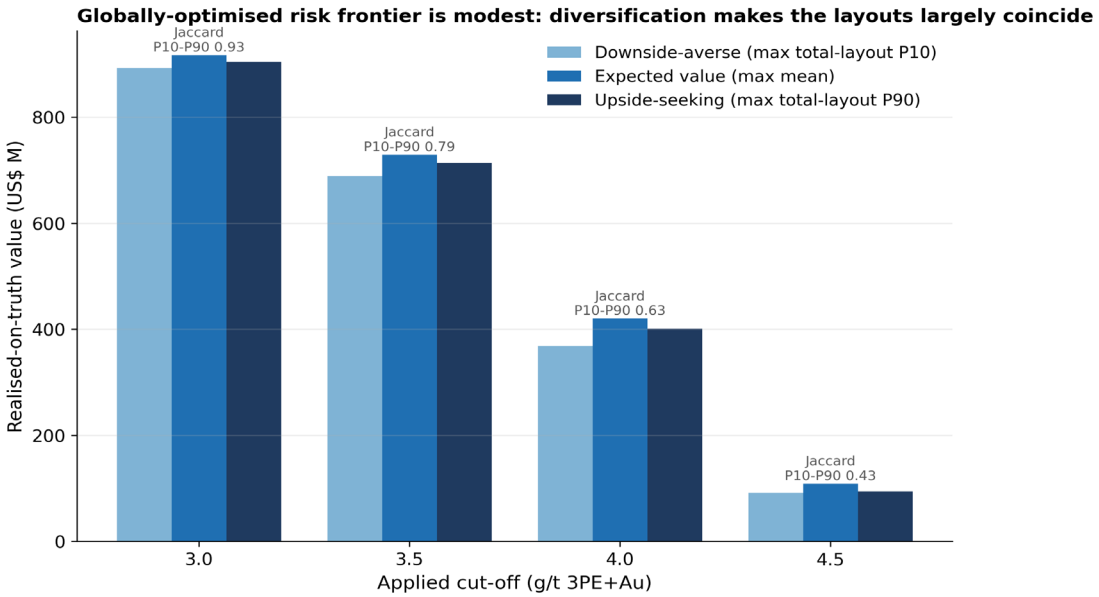


FIG 6 – Globally-optimised risk frontier. Realised-on-reference value for downside-averse (maximise total-layout P10), expected-value and upside-seeking (maximise total-layout P90) objectives. The frontier is modest and the layouts largely coincide (Jaccard shown), because stope-level risks diversify across the layout.

TABLE 3

Overlap (Jaccard similarity) of the globally risk-optimised layouts.

Cut-off (g/t)	P10-mean	mean-P90	P10-P90
3.0	0.94	0.96	0.93
3.5	0.84	0.87	0.79
4.0	0.77	0.75	0.63
4.5	0.55	0.64	0.43

DISCUSSION

The results draw a clear line around when stochastic optimisation helps. It does not help a decision that depends only on the average grade, there a single well-estimated conditional mean is sufficient, and the simulation E-type is close enough to kriging to change nothing. It helps a decision that is non-linear in grade, and the more contested the decision, the more it helps. For bulk stope mining the relevant non-linearity is the ore/waste destination, and the effect is small at a permissive cut-off and large at a selective one, positive across the ensemble at cut-offs up to 4.0 g/t rather than only on the reference. This reconciles our findings with the broader literature: the value of stochastic optimisation reported by Dimitrakopoulos (2011) and by Faria, Dimitrakopoulos and Pinto (2022), comes precisely from such non-linearities, and where those are absent, so is the value.

The block-support result adds a practical warning often overlooked in stope studies: the size of the block model is not a neutral modelling choice, because a coarse model smooths grade and blurs the ore/waste boundary. The stope-size result complements it, geometry alone does not select a best size; the development-cost structure does. The risk analysis carries its own caution: a genuine risk-return frontier must be evaluated on the whole layout, because diversification can make individually risky stopes collectively safe, and a per-stope rule will otherwise overstate the room for risk-tailoring.

LIMITATIONS AND FUTURE WORK

The study is deliberately bounded. Value is measured as an undiscounted stope-layout economic value consistent with the DIA stope layout framework, not a net present value; scheduling, discounting and mill capacity are not modelled, and full stochastic mine scheduling, where the value

of uncertainty is typically greatest, is left to future work. The analysis uses a single synthetic reference panel, so the magnitudes are illustrative of the mechanisms rather than predictive for any specific mine. Since the reference is drawn from the same multi-Gaussian model framework used for simulations, the comparison is an in-model validation rather than an out-of-model stress test. The development-cost figure used to illustrate the stope-size trade-off is representative and would need project-specific calibration. The reference is not a Mineral Resource nor a Mineral Reserve estimate for the Platreef deposit and is not a valuation of any asset. Natural extensions of the work presented in this paper include application to an erratic, sparsely drilled deposit, on which a wider risk-return frontier may open, and the integration of the stope decision into a stochastic production schedule.

CONCLUSIONS

Carrying grade uncertainty into underground stope-layout optimisation is worthwhile, but not universally. On a known synthetic Platreef reference, optimising the stope layout on a 50-member simulation ensemble rather than on a single kriged model adds no material value when the mining decision is linear, and up to 22.5 per cent of realised value when the decision is a non-linear ore/waste destination decision at a selective cut-off, a gain that is positive across the ensemble at cut-offs up to 4.0 g/t. The block-model support and the development-cost structure are first-order controls: coarser blocks lose value through the volume-variance effect, and the optimal stope size is set by development cost, not geology. When the risk of the whole layout is optimised globally, the risk-return frontier is modest, because stope-level risks diversify. The work extends the authors' deterministic DIA stope-optimisation research to grade uncertainty, and its conclusions align with the stochastic stope-design results of Faria, Dimitrakopoulos and Pinto (2022).

DATA AND CODE AVAILABILITY

The grade-simulation, economic and optimisation engine is implemented in pure NumPy/SciPy and is seeded for reproducibility. The deposited package includes the code, the economic parameter file of Table 1, and the result files (results_final.json and the recomputed review data) used to produce every figure, together with the fixed master seed and an expected-runtime note for regenerating the grade fields. The authoritative code and parameter file are available from the corresponding author.

REFERENCES

Deutsch, C V and Journel, A G, 1998. *GSLIB*: *Geostatistical Software Library and User's Guide*, second edition (Oxford University Press: New York).

Dimitrakopoulos, R, 2011. Stochastic optimization for strategic mine planning: a decade of developments, *Journal of Mining Science*, 47(2):138–150.

Faria, M F E, Dimitrakopoulos, R and Pinto, C L L, 2022. Stochastic stope design optimisation under grade uncertainty and dynamic development costs, *International Journal of Mining, Reclamation and Environment*, 36(2):81–103.

Goovaerts, P, 1997. *Geostatistics for Natural Resources Evaluation* (Oxford University Press: New York).

Grobler, D F, Brits, J, Maier, W D and Crossingham, A, 2019. Litho- and chemostratigraphy of the Flatreef PGE deposit, northern Bushveld Complex, *Mineralium Deposita*, 54:3–28.

Ivanhoe Mines, 2025. Platreef Integrated Development Plan 2025 (IDP25), NI 43-101 Technical Report, effective 15 February 2025 [online]. Available from: <https://www.ivanhoemines.com/wp-content/uploads/250329-Platreef-IDP25_RevF.pdf> [Accessed: 4 September 2026].

Journel, A G and Huijbregts, C J, 1978. *Mining Geostatistics* (Academic Press: London).

Kekana, S, Broughton, D, de Wet, B, Williams, N, Nielsen, S and Grobler, D, 2017. Discovery of the deep Flatreef PGE-Ni-Cu-Au deposit on Turfspruit Farm, Northern Limb, Bushveld Complex, South Africa, in *Proceedings of Exploration* 17: *Sixth Decennial International Conference on Mineral Exploration* (eds: V Tschirhart and M D Thomas), pp 671–678.

Kirkpatrick, S, Gelatt, C D and Vecchi, M P, 1983. Optimization by simulated annealing, *Science*, 220(4598):671–680.

Maier, W D, Barnes, S-J, Godel, B M, Grobler, D and Smith, W D, 2023. Petrogenesis of thick, high-grade PGE mineralisation in the Flatreef, northern Bushveld Complex, *Mineralium Deposita*, 58:881–902.

Nhleko, A S and Musingwini, C, 2022. Optimisation of three-dimensional stope layouts using a dual interchange algorithm for improved value creation, *Minerals*, 12(5):501.

Nhleko, A S, 2022. A dual interchange algorithm for three-dimensional stope boundary optimisation for underground mines, PhD thesis (unpublished), University of the Witwatersrand, Johannesburg.

Nhleko, A S, Musingwini, C and Mabala, M I, 2023. Using a dual interchange algorithm to evaluate the effect of stope size on economic value in planning an underground sublevel open-stoping layout, in *Proceedings 41st APCOM Symposium* (Society for Mining, Metallurgy and Exploration: Englewood) [online]. Available from: <https://onemine.org/documents/using-a-dual-interchange-algorithm-to-evaluate-the-effect-of-stope-size-on-economic-value-in-planning-an-underground-sublevel-open-stoping-layout> [Accessed: 4 September 2026].

Nhleko, A S, Tholana, T and Neingo, P N, 2018. A review of underground stope boundary optimization algorithms, *Resources Policy*, 56:59–69.

OreWin, 2017. Platreef 2017 Feasibility Study, NI 43-101 Technical Report prepared for Ivanhoe Mines Ltd, effective 4 September 2017 [online]. Available from: <https://www.ivanhoemines.com/wp-content/uploads/platreef-feasibility-study-september-2017.pdf> [Accessed: 4 September 2026].

Dynamic optimisation of integrated mine value chains – a tractable MILP framework for adaptive decision-making

J Qu[1], C Xu[2] and P Dowd[3]

1. PhD student, The University of Adelaide, Adelaide SA 5005. Email: junlin.qu@adelaide.edu.au
2. Associate Professor, The University of Adelaide, Adelaide SA 5005. Email: chaoshui.xu@adelaide.edu.au
3. Professor, The University of Adelaide, Adelaide SA 5005. Email: peter.dowd@adelaide.edu.au

ABSTRACT

Modern mine production planning increasingly requires decision frameworks that extend beyond the generation of a single static extraction schedule. Mining, stockpiling and processing decisions are physically and economically interconnected: extraction determines the timing and characteristics of material entering the value chain, stockpiles provide inter-period flexibility, and downstream processing requirements influence the value of upstream routing and sequencing decisions (Goodfellow and Dimitrakopoulos, 2017; Carelos Andrade and Dimitrakopoulos, 2024). Treating these components sequentially may simplify the planning problem, but can overlook interactions that determine system-wide performance. The limitation becomes more significant during execution, when realised production can depart from the assumptions under which the original schedule was generated. Changes in material extraction, routing and inventory conditions alter the physical state available for subsequent decisions, such that a schedule that was optimal, or near-optimal, when generated may no longer represent the best attainable plan for the remaining operation (Goodwin *et al*, 2006; Armstrong *et al*, 2021). The resulting challenge is therefore twofold: to retain an integrated mine–stockpile–plant decision scope while making the optimisation sufficiently tractable for repeated re-planning as the value-chain state evolves.

This study develops a Dynamic Mixed-Integer Linear Programming framework, referred to as DYMILP, for adaptive production planning in an integrated underground mine value chain. The work builds on a reference mixed-integer quadratically constrained programming (MIQCP) formulation that jointly represents ring-level sublevel-caving extraction, direct mine-to-plant delivery, run-of-mine stockpile allocation, stockpile reclaiming and multi-element plant-feed control. Although the reference MIQCP provides the required integrated decision structure, stockpile-related nonlinearities make repeated optimisation computationally demanding (Moreno *et al*, 2017; Carelos Andrade and Dimitrakopoulos, 2024). In the three-stockpile reference benchmark, the non-convex MIQCP required approximately 100 000 sec under the adopted high-performance-computing setting and terminated with an optimality gap of approximately 10 per cent. Such computational behaviour may be acceptable for an occasional strategic optimisation exercise, but is poorly suited to an adaptive framework in which the remaining planning problem must be solved repeatedly as production progresses. The present study therefore develops a MILP-compatible representation of the computationally restrictive stockpile relationships while retaining the principal physical and economic decision scope of the integrated value chain. The detailed mathematical transformation is outside the scope of this extended abstract; the emphasis here is on the resulting ability to use the integrated formulation as a recurrent optimisation engine.

The significance of this tractability improvement extends beyond computational speed. DYMILP embeds the resulting MILP engine within a rolling-horizon, state-feedback decision architecture. At each dynamic stage, the integrated mine–stockpile–plant problem is solved over the remaining planning horizon conditional on the currently realised physical state. Only the decisions associated with the immediate production period are implemented. Realised production outcomes are subsequently observed, the remaining mine and inventory states are reconciled, and all unimplemented future decisions are reconsidered through a new remaining-horizon optimisation. The procedure is then repeated as the planning horizon advances. Consequently, future extraction, routing, stockpile utilisation and processing decisions are not treated as irrevocable commitments made at the beginning of the horizon; they remain conditional on how the operation actually develops. This distinction converts the role of the optimisation model from generating a fixed

production schedule to supporting a sequential plan–implement–observe–update–re-optimise decision process (Goodwin *et al*, 2006; Armstrong *et al*, 2021).

The proposed framework should not be interpreted as a globally optimal multistage stochastic policy. DYMILP does not enumerate all possible future production deviations in a scenario tree, nor does it solve a dynamic programming problem over every possible future state. Instead, it provides stage-wise conditional re-optimisation: at each decision stage, the remaining value chain is optimised, subject to the current realised state and the deterministic assumptions of the underlying MILP. This interpretation is important both computationally and operationally. A full multistage stochastic formulation combining detailed sublevel-caving sequencing, material routing, stockpile inventories, reclaim decisions and multi-element plant-feed restrictions would generate a substantially more difficult optimisation problem (Martinez and Newman, 2011; Carelos Andrade and Dimitrakopoulos, 2024). The rolling-horizon approach instead concentrates computational effort on the state that has actually been realised, while preserving the flexibility to revise subsequent decisions when previous production outcomes differ from plan.

The framework is evaluated using a multi-element iron-oxide copper-gold sublevel-caving case study containing approximately 4.7 Mt of mineralised material. The production domain is represented by 60 stopes and 300 ring-level production units over a 12-period monthly planning horizon. Copper and gold represent payable elements, while uranium is treated as a deleterious plant-feed attribute, requiring the optimisation to coordinate economic value, material availability, stockpile flexibility and multi-element feed control. The numerical evaluation is structured in two stages. First, the static MILP engine is assessed against the integrated MIQCP reference to establish whether the reformulated optimisation retains the principal economic, material-flow and plant-feed behaviour required for subsequent dynamic use while providing a more suitable computational structure for repeated solution. Second, the rolling-horizon framework is compared with static MILP schedule execution under the same realised production-deviation setting. In the static policy, the original optimised future decisions remain unchanged after execution deviations occur; in DYMILP, the resulting physical state is propagated through the integrated value chain before the remaining horizon is re-optimised. This controlled comparison isolates the value of adaptive feedback rather than differences in geological, economic or disturbance assumptions.

Performance is evaluated from economic, operational and computational perspectives. Economic assessment considers realised value and the decomposition of downstream penalty and feed-control effects. Operational performance is examined through plant-feed behaviour, stockpile utilisation, material-flow trajectories and reserve depletion, while material-balance diagnostics verify physical consistency throughout sequential state updates. Computational reliability is assessed at each optimisation stage because the practical usefulness of a rolling-horizon framework depends not on a single successful solve, but on the ability of the optimisation engine to repeatedly generate credible remaining-horizon solutions as the state changes. The purpose is therefore not to assume that re-optimisation must outperform static execution in every realised trajectory, but to determine under controlled conditions when and how state feedback can improve the allocation of remaining flexibility across the mine value chain.

The principal contribution of DYMILP lies in connecting integrated value-chain optimisation, computational tractability and adaptive decision-making within a single planning architecture. The study does not claim novelty from stockpiling, MILP scheduling or rolling-horizon optimisation individually. Rather, it demonstrates how an integrated sublevel-caving mine–stockpile–plant optimisation structure can be made suitable for recurrent remaining-horizon solution and subsequently used to respond to realised production states without decomposing the value chain into independent scheduling, inventory-management and plant-feed correction problems. More broadly, the framework reframes mine production planning from the search for one fixed optimal schedule toward the repeated optimisation of the remaining system as operational reality unfolds. This provides a pathway toward more responsive strategic and tactical mine planning in which stockpiles function not merely as storage facilities, but as sources of temporal flexibility that allow upstream extraction and downstream processing decisions to be coordinated throughout execution.

REFERENCES

Armstrong, M, Lagos, T, Emery, X, Homem-de-Mello, T, Lagos, G and Sauré, D, 2021. Adaptive open pit mining planning under geological uncertainty, *Resources Policy*, 72:102086.

Carelos Andrade, L and Dimitrakopoulos, R, 2024. Integrated stochastic underground mine planning with long-term stockpiling: method and impacts of using high-order sequential simulations, *Minerals*, 14:123.

Goodfellow, R and Dimitrakopoulos, R, 2017. Simultaneous stochastic optimization of mining complexes and mineral value chains, *Mathematical Geosciences*, 49:341–360.

Goodwin, G C, Seron, M M, Middleton, R H, Zhang, M, Hennessy, B F, Stone, P M and Menabde, M, 2006. Receding horizon control applied to optimal mine planning, *Automatica*, 42:1337–1342.

Martinez, M A and Newman, A M, 2011. A solution approach for optimizing long- and short-term production scheduling at LKAB's Kiruna mine, *European Journal of Operational Research*, 211:184–197.

Moreno, E, Rezakhah, M, Newman, A M and Ferreira, F, 2017. Linear models for stockpiling in open pit mine production scheduling problems, *European Journal of Operational Research*, 260:212–221.

Author index

Abuntori, C 193
Acheampong, A 193
Ajak, A D 579
Aji, C 441
Alegre, D 311, 431
Almeida, C 155
Amaral, A 405
Andrade, L C 311, 431
Andre, J 339
Ansah, E O 193
Artica, C 327
Asad, M W A 579
Aydar, A 227
Balboa-Espinoza, V 49
Barbosa, D 405
Barnes, R 585
Benzaazoua, M 401
Both, C 333
Bouhlali, N 401
Brady, T 137
Burgers, H 441
Câmara, T 373, 379, 519
Camero, H 17
Carswell, J 177
Carvalho, J 405
Chauhan, S S 73
Chauke, E 79
Chester, C 585
Chimunhu, P 579
Chung, J 119
Dadzie, L K 193
Dagdelen, K 137, 227
Damasceno, F 437
de Villiers, D J 495
Delage, E 73
Dempster, K 321
Deucker, E 119
Deutsch, C 343
Deutsch, J 343
Deutsch, M 343
Dimitrakopoulos, R 49, 69, 73, 259, 367
Ditlef, A 405
Dowd, P 617
Egger, F 97
Elghali, A 401
Elkington, T 245
Faradonbeh, R S 579
Faramarzi, F 293
Faria, M 311, 417
Ferrie, F 281
Findlay, L 367
Fletcher, Z 427
Flett, A 453
Fosdike, E 339
Fouché, L 97
Fox, N 327
Fox-Martin, M J 495
Gadelha, B 21, 373, 379, 519
Galloway, D 379
Gariba, G 311
Godoy, M 3
Goodfellow, R 321
Hannett, A 453
Herrero, J 259
Heyes, J 107
Hill, S 481
Hoerger, S F 137
Irwin, R 107
Jani, A 279
Jewbali, A 45
Jordan, L 177
Kaykov, D 197

King, B 33
Knights, P 441
Kumar, A 321
Kumral, M 197
Lawrence, L 173
Leiva, O 43
Light, S 585
Lioret, A 417
Lison, C 107
Luo, J 69
Mabala, M I 603
Mackenzie, J 527
MacRae, A 589
Mallett, R 379
Mariz, J L V 311
Marriott, B 391
Mayrink, A 311, 405, 417
Minniakhmetov, I 427, 585
Mochubele, M 79
Montiel, L 373, 379, 519
Moradi, A 401
Mu, Y 281
Musingwini, C 603
Musolino, M 585
Nhleko, A S 603
Noppé, M 327
O'Neill, T 21
Oliveira, M M 565
Patton, W 391
Peroni, R 155, 311
Petro, L 589
Place, D 481
Poblete, C 481
Ponce, E 453
Prawasono, A 17
Price, B 457
Qu, J 617
Raval, S 213
Rispoli, A 431
Roberts, M 339
Rowe, L 213
Rudge, I 481
Santos, G 405
Santos, R 245
Saxena, A 547
Sega, J 437
Senécal, R 417
Shchichko, L 73
Silva, A 405
Silva, L 417
Silva, M 21, 519
Silva, P 437
Silva, R P S 565
Suchak, A 465
Supriatna, R 213
Taha, Y 401
Talbot, B 177
Todd, J 21, 373, 379
Topal, E 579
Town, J R 547
Victor, L 527
Ward, R 107
Whitchurch, K 173
Whittle, G 11
Wilogo, A R 347
Xu, C 617
Yaqub, B 539
Yeates, G A 13