Journal of Industrial Safety xxx (xxxx) xxx–xxx

Contents lists available at ScienceDirect

Journal of Industrial Safety

journal homepage: www.sciencedirect.com/journal/jinse

Breaching the mining fatalities plateau: Financial pressure, operational instability, and the case for flow-based safety

Hendrik Lourens, FAusIMM[a,b,*]

[a] Stratflow, Queensland, Australia
[b] Fellow of the Australasian Institute of Mining and Metallurgy (FAusIMM); Member, AusIMM Health & Safety Society Committee, Australia

ARTICLE INFO

Keywords:
Mining safety
Fatality plateau
Theory of Constraints
High Reliability Organizations
Balanced capacity
Production variability

ABSTRACT

Between 2003 and 2013, the Australian mining fatality rates fell by about 80%, but thereafter, in the past decade, they have remained at the same level of approximately seven fatalities per year despite significant investment in safety systems. In Queensland, the Brady Review (2020) revealed a structural fatality cycle linked to the economic and operational pressures. An independent analysis has confirmed that neither increased regulation nor additional administrative controls have improved the situation. This perspective article draws on more than 20 years of practitioner experience across more than 90 interventions in mining operations on three continents, grounded in the Theory of Constraints (TOC). The paper asserts that the current safety plateau is consistent with the predictable behavior of balanced-capacity mining operations—the industry-standard approach of maximizing the utilization of every resource in the production chain. It further argues that balancing capacity tends to create chronic operational instability, which in turn can undercut the cognitive and organizational preconditions for high-reliability organization behavior. In this article, we describe a flow-based production intervention, grounded in the TOC principles, which is associated with improved productivity (a 10–40% increase in output using constant resources). A consistent but initially unanticipated finding was an accompanying improvement in safety outcomes. The study examines how Queensland's Resources Safety and Health Legislation Amendment (RSHLA) Act 2024, while correct in principle, may in some situations amplify instability during balanced-capacity mining operations. The paper proposes a set of testable research hypotheses linking the run-of-mine production variability to safety incident rates and recommends the RSHLA Act 2024 implementation period as a natural experiment to confirm these practitioners' observations.

1. Introduction

Between 2003 and 2013, the Australian mining industry achieved one of the most remarkable safety improvements of any heavy industry globally, with the fatality rates falling by approximately 80%. In 2003, there were 12.4 fatalities per 100,000 workers employed in Australian mines; a decade later, that indicator was 3.4. This improvement was widely attributed to stronger regulation, better equipment, upgraded training, and a cultural shift toward safety consciousness.

However, the desirable trend subsequently stalled, despite continued investment in safety systems, technology, and training. The industry has averaged approximately seven fatalities per year over the past decade. These two regimes, a decade of steep improvement followed by a decade of stasis, are clearly visible in Fig. 1. The figure plots the annual work-related fatality rates in the Australian mining industry, expressed per 100,000 workers, from 2003 to 2022. Its trend is the empirical starting point for this study. The rate falls steeply, from 12.4 per 100,000 workers in 2003 to 3.4 in 2013, and then flattens: for the decade since, it has fluctuated around a mean of roughly three fatalities per 100,000 workers with no sustained further improvement. It is the flat right-hand portion of the curve—the plateau—that this paper sets out to explain and to change. The interpretation of the year-to-year variations within that plateau, including the rise between 2013 and 2015, is presented in Section 2. Most deaths result from well-understood hazards—being struck by objects or caught in machinery, falling from heights, or vehicle collisions—for which well-documented controls exist [1]. The question of why those established controls fail to prevent the recurring events has become one of the most pressing in Australian mining safety.

Peer review under the responsibility of Shandong University of Science and Technology.
* Corresponding author at: Stratflow, Queensland, Australia.
E-mail address: hendrik@stratflow.com.au.

https://doi.org/10.1016/j.jinse.2026.06.002
Received 14 April 2026; Received in revised form 22 June 2026; Accepted 29 June 2026

Please cite this article as: H. Lourens, Breaching the mining fatalities plateau: Financial pressure, operational instability, and the case for flow-based safety, Journal of Industrial Safety, https://doi.org/10.1016/j.jinse.2026.06.002

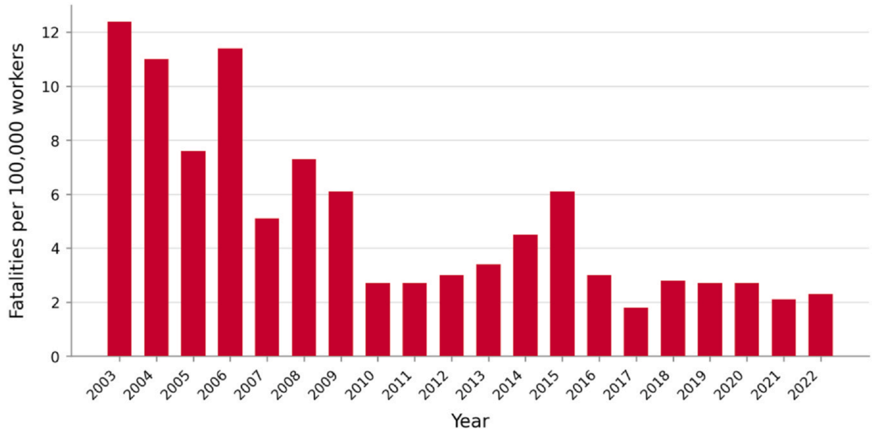

Fig. 1. Annual work-related fatality rate of the Australian mining industry, 2003–2022 (fatalities per 100,000 workers). Source: Safe Work Australia [2]; chart redrawn by the author, after Cliff [3].

Two features of the series warrant comment. First, the steep decline before 2013 is largely due to the fact that the mining workforce expanded roughly 3.5-fold, from about 84,000 in 2003 to about 290,000 by 2013, during the investment boom. Because the number of workers killed each year was approximately stable over that period, of the order of 10, the workforce that expanded naturally depressed a per-capita rate to about a third of its starting value, but with no change whatsoever in the underlying hazard. The transformation visible in the rate, therefore, overstates the improvement actually experienced on site: the number of deaths reached its plateau considerably earlier than the rate did.

In the period 2013–2015, the rate rose noticeably, peaking near six fatalities per 100,000 workers in 2015. Part of this movement reflects a contracting workforce, rather than a change in the underlying hazard. If any part of the upward movement was real, the most plausible systemic contributor was the unwinding of the mining investment boom over precisely these years. The boom peaked in 2012–2013; the construction phase then contracted sharply, and tens of thousands of positions were shed as commodity prices fell. Contractions of this kind bring cost reduction, the loss of experienced personnel, increased casualization, and elevated production pressure on a smaller workforce—the conditions that the Brady Review of two decades of Queensland mining deaths associates with the recurring “fatality cycle”, in which vigilance decays and clusters of deaths follow, with inadequate supervision and inexperience prominent among the causes. Whatever the cause, it seems clear that the conventional safety thinking has carried the industry to a plateau it has not managed to cross in more than a decade.

Queensland staged three safety resets in five years; unfortunately, so far, they have yielded only limited results [1]. Each reset followed the pattern of a serious incident or cluster of incidents triggering heightened regulatory attention and industry commitment, followed by a gradual return to previous behavioral patterns. This is unlikely to be a failure of devotion to safety. It is because the reset addresses symptoms without changing the underlying system dynamics that produce them.

This perspective article proposes that breaking through the persistent safety plateau requires a different way of organizing production. Drawing on more than 20 years of practitioner experience across more than 90 mining interventions on three continents, grounded in established theories from systems thinking, organizational psychology, and High Reliability Organization (HRO) research, this article argues that the industry-standard approach of implementing balanced-capacity operations (which became prevalent from 2000 onwards) [4] creates the chronic instability that makes safety improvements difficult. It describes a counterintuitive alternative, flow-based operations, which the practitioner experience associates with improvements in both production results and safety outcomes, examines the potential unintended consequences of Queensland’s landmark Resources Safety and Health Legislation Amendment (RSHLA) Act 2024 for balanced-capacity mines, and proposes a specific, testable research agenda to corroborate these practitioners’ observations.

2. The safety plateau and the Brady review

2.1. The structural fatality cycle

In 2020, Brady [1] performed a comprehensive review of all fatal accidents in Queensland mines and quarries from 2000 to 2019. Over that 19.5-year period, 47 mining industry fatalities occurred in Queensland alone. Brady identified a pattern termed the ‘fatality cycle’: periods where a significant number of fatalities cluster, followed by periods where there are few to none, then followed again by another cluster. The review estimated that roughly 12 fatalities will occur in any five-year period in Queensland mines if the industry does not fundamentally change its approach. This is not a prediction of randomness; it is a structural forecast based on the way the industry operates.

Critically, Brady found that the industry’s primary response to fatal incidents was to rely on administrative controls—more paperwork, checklists, and procedures. These sit among the least effective tiers in the hierarchy of controls [1]. The industry was responding to system-wide failures with systemic non-solutions.

2.2. The financial pressure link

One of Brady’s most important findings was the correlation between financial and economic stresses and the timing of fatality clusters. The fatality cycle did not act independently of the commercial environment; it tracked it. Periods of intense production pressure—driven by commodity price cycles, cost-cutting, and efficiency drives—corresponded to periods of increased fatal incidents.

The McKinsey MineLens productivity index provides further evidence for this link. From a 2000 baseline of 100, the global mining productivity dropped roughly 40% by 2010. Australia has since shown some recovery, particularly in surface mining, but improvements in underground mining have lagged behind [4]. When overlaid with fatality data, the two series appear to move together; whether this reflects a causal link or merely associated variables is considered below.

Cliff [3] has conducted an independent analysis confirming the plateau. This research demonstrates that despite the adoption of risk management systems, competency training, and a shift away from prescriptive regulation following the Moura disasters, the rate of deaths and serious injuries has not budged over the past decade. Cliff has also identified other contributing factors, such as worsening mental health and rising fatigue. These factors plausibly become more prevalent in a work environment characterized by chronic pressure, unpredictability, and reactive management.

Whether production pressure actually degrades safety is a question Queensland has examined directly. In February 2023, the Queensland Parliament's Transport and Resources Committee reported on its Inquiry into Coal Mining Industry Safety [5]. One section of that report dealt specifically with the impact of coal production rates on safety risk management. Two opposing accounts emerged. Every coal mine operator that made a submission characterized safety as its overriding corporate value, with production metrics "considered subsidiary to safety metrics" and "no production … warranted unless it can be completed safely," and no operator accepted that production rates negatively affect safety. Worker representatives described that the reverse occurred in daily practice. They reported that established procedures were departed from to protect output, and that these departures clustered in the lead-up to the month-end and year-end, when production targets must be met. One coal mine worker put it to the committee plainly, saying, "Production will always come before safety." A 2019 survey of 1010 Queensland coal mine workers found that 80% named companies valuing production over safety as their foremost safety concern.

The committee could not adjudicate between these conflicting accounts. It recorded "the impossible task of evaluating conflicting anecdotal evidence … in the absence of empirical evidence", and recommended that independent research be commissioned into how production rates affect safety risk management [5]. That unresolved issue of whether production pressure degrades safety, and if so, how it could be detected, is the question the present study tackles. The argument developed here is that the disagreement is resolvable, and that the measure everyone has been missing is not the production rate itself. It is production instability: the day-to-day variability that the prevailing balanced-capacity paradigm generates (Section 3), that a flow-based model removes (Section 4), and that can be computed from data mines already collected (Section 7).

3. Conceptual framework

3.1. Balanced capacity as the root cause

The dominant paradigm in mining operations management is balanced capacity: designing and resourcing the production chain so that every unit operates at its maximum utilization. The logic is intuitive—if every resource is fully utilized, the system must be maximally efficient. This logic is wrong.

In any production chain with dependent events and statistical variability—which describes every mining operation—a balanced-capacity operation does not produce stable, high output at the lowest all-in cost per tonne. What it does produce is chronic perturbation [6,7]. When every production unit is resourced to run at its theoretical 100% utilization, there is no buffer to absorb the inevitable disruptions: equipment breakdowns, weather events, geological surprises, absenteeism, and supply chain delays. Each disruption sets off a domino effect throughout the chain, creating a permanent state of recovery: the mine lives in a perpetual cycle of falling behind and pushing to catch up.

Goldratt [6] demonstrated this mathematically and practically in The Goal: a balanced production chain of six units, each rated at 10 tonnes per hour, will not produce 10 tonnes per hour. Due to the interaction of dependent events and statistical variability, the actual throughput can often converge on outputs as low as 5 tonnes per hour, i.e., as little as half the rated capacity. The system appears to have ample capacity, yet delivers dramatically less than expected. This gap between the apparent capacity and the actual output is, this article suggests, a major source of the chronic production pressure that the Brady Review found to be correlated with fatality clusters.

3.1.1. The origin of the production gap in a balanced-capacity production chain

Consider a mining operation as a chain of six sequential stages: drilling, blasting, loading and hauling, crushing, and processing. Material passes through them in order, with each stage receiving work from the preceding stage and handing work to the following stage. Suppose each stage can process, on average, 10 units of material per hour. Most observers would expect the chain as a whole to deliver 10 units per hour, which would imply every stage running at its full efficiency. In reality, it does not achieve that ideal efficiency, and the reason lies in two features that every mining operation shares.

The first is variability. An average of 10 units per hour does not mean 10 units in every hour. A stage might process 12 in one hour, eight in the next, and, when equipment fails, ground conditions change, or a crew is short-handed, none at all for a period. The long-run average is 10, but the hour-to-hour real efficiency fluctuates widely around it.

The second is interdependence. Because the stages are linked in sequence, the effects of a stoppage at any one of them propagate in both directions. The stalled stage blocks the stage feeding it, which now has nowhere to send its output, and starves every stage after it, which now has nothing to receive. A single stoppage therefore idles much of the chain, not only the stage where it occurred.

When variability and interdependence act together, the output of the whole chain at any moment is determined by the slowest stage at that moment. As illustrated in Fig. 2, although every stage is rated at 10 units per hour, this interaction lowers the realized throughput of the whole chain to approximately 5 [8,9]. The gap between the 10 units the chain appears to deliver and the 5 it actually delivers is the structural production deficit at the center of this paper's argument. It generates chronic recovery pressure, and it is present by design in any balanced-capacity operation.

3.2. Rasmussen's dynamic safety model

The Rasmussen [10] dynamic safety model provides a visual scheme for understanding the pathway by which balanced capacity could degrade safety. Rasmussen proposed that organizations operate within a space bordered by three constraints: an acceptable performance boundary (the boundary of safe operation), an economic cost boundary, and a workload boundary. Within this space, organizations are subject to two ongoing pressures: management pressure toward increased efficiency and a natural gradient toward least effort. The combination causes a slow, often invisible migration of the operating point toward the acceptable performance boundary—Rasmussen's 'drift into failure.'.

Rasmussen's model is shown in Fig. 3. The figure depicts the three boundaries that enclose the space within which an operation moves. These are the economic-failure boundary, the unacceptable-workload boundary, and the boundary of acceptable (safe) performance. It also shows the two gradients that act on the operating point: the management push toward efficiency and the individual drift toward least effort. As Fig. 3 makes clear, no single decision drives the operating point across the accident boundary. Under the combined pull of the two gradients, it migrates there gradually, which is what makes the drift so difficult to see from inside the organization.

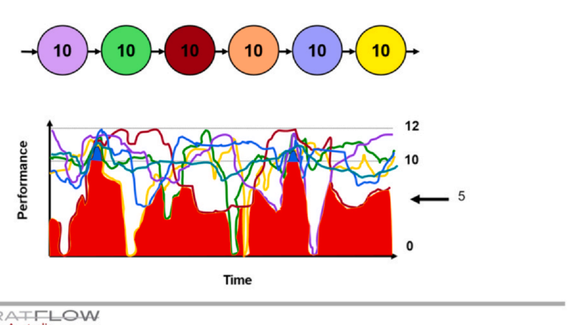

Fig. 2. Balanced-capacity output for an interdependent, highly variable production chain.

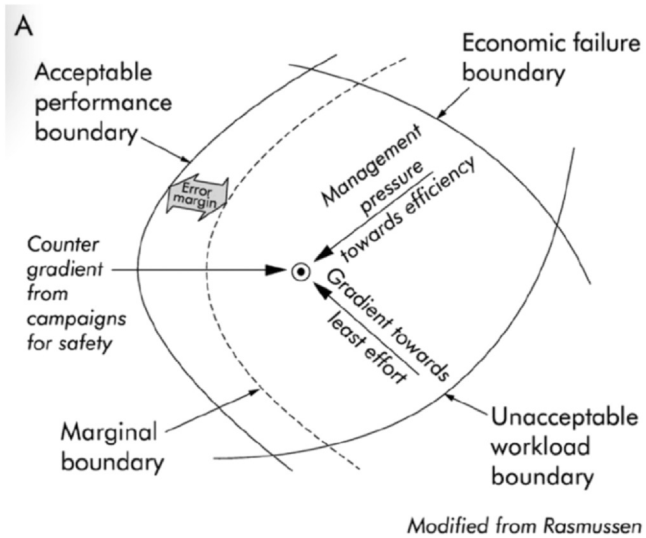

Fig. 3. Rasmussen's dynamic safety model showing the three boundaries (economic failure, workload, and acceptable performance) and the gradients that push the operating point toward the accident boundary. Originally depicted in Rasmussen (1997) [10] and Cook and Rasmussen (2005) [11]. Reproduced from Morrison and Wears (2022) [12] under CC BY 4.0.

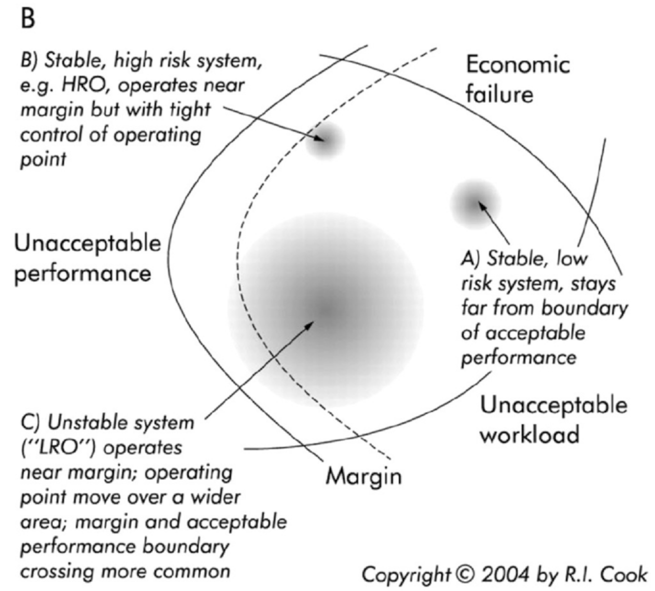

Fig. 4. Rasmussen's dynamic safety model as depicted in Cook and Rasmussen (2005) [11]. (A) Stable, low-risk system—operating point moves in small increments, well away from the boundary; (B) Stable high-risk system—operating point stays near the marginal boundary; (C) Unstable system - large, rapid shifts in operating point. Originally Copyright © 2004 by R.I. Cook. Reproduced from Morrison and Wears (2022) [12] under CC BY 4.0.

This drift is insidious because each step is small and reasonable. The maintenance schedule is extended by a day. A safety check is abbreviated. A buffer is reduced. Each decision makes sense locally, but cumulatively they erode the safety margin until the organization is, as Cook and Rasmussen [11] described it, 'flirting with the margin'.

3.3. High reliability vs. low reliability operating patterns

Cook and Rasmussen [11] extended Rasmussen's model by proposing that HROs can be visualized as a narrowly spaced set of operating points within the safety space clustered well away from the acceptable performance boundary. Low Reliability Organizations (LROs), by contrast, are more diffuse, with operating points scattered widely and frequently approaching or crossing the boundary.

Fig. 4 shows three system states of the model. Cluster A represents a stable, low-risk system, where the operating point moves in small increments and stays well clear of the boundary. Cluster B corresponds to a stable but high-risk system, where the operating point is controlled but sits chronically close to the marginal boundary. Cluster C is an unstable system, where the operating point shifts in large, rapid excursions. In Fig. 4, an HRO corresponds to Cluster A, a narrow cluster of operating points held away from the boundary. A balanced-capacity mine in chronic perturbation corresponds to Cluster C, where the wide scatter of operating points repeatedly carries the operation toward the boundary. An HRO is not simply an organization with good safety systems. It is one whose operating points are consistently and narrowly clustered in the safe zone. This requires two things: the ability to control operational variation (narrowing the cluster) and sufficient margin between the cluster and the boundary (keeping it in the safe zone). Balanced-capacity operations fail on both counts. The high level of production instability produces a wide scatter of operating points, and the financial pressure to maximize utilization pushes the average position toward the boundary.

3.4. The threat-rigidity effect

Staw et al. [13] explain that people subjected to sustained threat narrow their perceptual field, fall back on habitual responses, and push their decision-making upward. This is the opposite of what HRO demands. HRO requires widened attention; threat-rigidity narrows it. Where HRO requires pushing decisions to the frontline, threat-rigidity centralizes them.

A balanced-capacity mine in chronic perturbation is running a permanent, low-grade threat-rigidity response. This is not a temporary state that will pass once the current crisis resolves; it is a structural feature of the system design. On this account, the crises are structural rather than incidental: they are likely to continue for as long as the mine is run as a balanced-capacity system.

3.5. Erosion of the five HRO principles

Weick and Sutcliffe [14] identified five principles that define HRO behavior. Each requires cognitive bandwidth, time, and organizational slack that balanced-capacity operations consume:

Preoccupation with failure requires noticing weak signals—small anomalies that precede serious incidents. In firefighting mode, only major incidents command attention. Supervisors who should scan for early warnings are chasing today's production shortfall instead.

Reluctance to simplify requires time to investigate properly—to resist labelling a problem and moving on. Under chronic production pressure, root causes get labelled and archived. Investigation becomes compliance rather than inquiry.

Sensitivity to operations requires seeing the whole system and understanding cross-functional interactions. Vertical silos and crisis management fragment this view. When every department fights its own fires, no one has the bandwidth to see the bigger picture.

Commitment to resilience explicitly requires buffers and excess resources—what organizational theorists call protective capacity [14]. This is exactly what balanced capacity eliminates by design. Weick and Sutcliffe [14] identified that HROs specifically maintain 'uncommitted resources' as a structural requirement for reliability.

Deference to expertise requires trusting frontline staff and pushing decision-making downward. Under production pressure, decisions are centralized upward. The supervisor who stops a process because something does not look right faces implicit, and sometimes explicit, pressure to keep operations going.

Table 1
Theory of constraints (TOC) definitions.

Term	Plain-language definition
Throughput	The rate at which the whole operation delivers saleable product, distinct from how busy any individual stage is.
Balanced capacity	Resourcing every stage of the production chain to the same nominal rate, leaving no stage with spare capacity. The conventional approach.
Constraint (bottleneck)	The single stage that sets the pace of the whole system. The system cannot produce faster than its constraint.
Buffer	A deliberate, managed reserve—a stock of material before the constraint, and space for output after it—that absorbs variability so the constraint is never starved or blocked.
Protective capacity	Spare capacity is deliberately held at every non-constraint stage so it can recover from disruption and keep the buffers replenished. Not idle waste—the margin that keeps the system stable.
Subordination	Managing every non-constraint stage with one priority: keeping the constraint running. The local efficiency of those stages is secondary.

Weick [15] demonstrated in his landmark study of the Mann Gulch fire disaster that when stress mounts and ambiguity rises, organizations lose the ability to make sense of what is happening. Structure and sensemaking collapse together.

4. The solution: flow-based operations

Table 1 defines the common terms used in flow-based operations. The sections that follow build this idea up from first principles and then describe how it is put into practice.

4.1. The shift from balanced to flow operations

Section 3 defines the problem. A balanced-capacity chain of six stages, each rated at 10 units per hour, delivers only about 5 units, because variability and interdependence interact to lower the realized throughput to the pace of the slowest stage at that moment (Fig. 2). That deficit is present by design in any operation resourced for full utilization. Flow-based operations are the response to this unsatisfactory reality.

At its simplest, it is a different answer to one question: where should a mine concentrate its effort? The conventional answer is everywhere at once, keeping every part of the operation as busy and as fully utilized as possible. This is the balanced-capacity approach Section 3 diagnosed. The spare capacity it purposefully eliminates is the only thing that could have absorbed the disruptions. Stripped of slack, the chain becomes less able to recover, output falls further, and the operation can slide into a cycle in which each disappointing result prompts another round of cost-cutting.

The flow-based operations paradigm offers the opposite answer. It holds that a mine, like any chain of dependent steps, can only move product as fast as its single slowest step, and that the way to lift the output of the whole is not to drive every step harder but to find that one pace-setting step, protect it so that it never stands idle, and organize everything else around keeping it supplied.

4.2. Putting flow into practice

Flow requires that we identify the constraint, protect it with buffers and capacity, and subordinate every other stage to it. In practice, flow is sustained by a single routine: a short daily cross-functional meeting, referred to here as the flow room.

A flow room is a thirty-minute meeting held each day and attended by operations, maintenance, planning, and safety. It addresses one question: what threatens the constraint over the next forty-eight hours? Those present work from the same information, including the state of the buffers, which is coded as green, amber, or red. In this manner, the condition of the system can be read at a glance. Decisions that would once have been escalated up the hierarchy and resolved slowly are settled in the room. While the buffers are healthy, the non-constraint stages use their protective capacity for maintenance, training, and safety-related work, rather than for output that the constraint cannot absorb.

This is best described as a dual operating system [16]. For thirty minutes a day, the flow room supplies cross-functional, horizontal coordination; for the rest of the day, the operation runs through its normal hierarchical structure. The flow room does not replace existing management systems; it adds the function they lack: the agility and cross-departmental coordination needed to keep an interdependent operation aligned. It is also where the operational discipline of flow meets the cognitive requirements of HRO behavior—a standing daily forum in which weak signals are surfaced, cross-functional interactions are made visible, and frontline expertise is brought to the decisions that matter most.

5. Evidence from practice and proposed mechanisms

5.1. The scale of evidence

The evidence underlying this paper is a portfolio of more than 90 production interventions grounded in the TOC and carried out in mining operations since 2000. These interventions span three continents: Africa, South America, and Australia. They cover most major mining methods, including open-cut truck-and-shovel, underground continuous-miner, and conventional bord-and-pillar, longwall, cut-and-fill, and deep-level conventional stopping; and most major commodity groups, including coal, iron ore, gold, platinum, copper, and diamonds.

The portfolio is not evenly distributed. Its largest concentration is in Africa, where the method was first developed and applied from 2000 onward across coal, iron ore, gold, platinum, and copper operations, principally in South Africa and Zambia. The South American interventions were carried out at open-cut copper operations in Chile. The Australian interventions are in coal mining and include a New South Wales longwall mine, examined in detail in Section 5.2.

Two qualifications must be made explicit before this evidence is presented. First, almost all of these interventions were commercial consulting engagements carried out under confidentiality obligations to the operators concerned; the operations in Table 2 are therefore identified by continent, commodity, and mining method rather than by name. This same constraint explains why a practitioner dataset of this size and internal consistency has not previously appeared in the peer-reviewed literature: authorization to publish named operational results is seldom granted, whereas an anonymized synthesis is the form in which the evidence can responsibly be released. Second, as set out in the Limitations (Section 8), this is a practitioner's observation rather than controlled research: the interventions were not randomized, were not designed as studies, and the portfolio carries the selection bias inherent in any consulting engagement portfolio. Table 2 should therefore be read as a structured account of recurring practitioner experience, rather than the result of a controlled trial.

Across this portfolio, the same production response recurs: once an operation is focused on protecting its constraint, output increases, and the day-to-day variability of Run-of-Mine (ROM) output falls. Fig. 5 shows this pattern for the New South Wales longwall mine examined in Section 5.2; the results from the other interventions are similar [8,9].

Table 2 lists representative interventions from across the three continents. The operations are a subset of the portfolio, chosen to span the range of continents, commodities, and mining methods. The results shown are those reported by the operators or recorded in practitioner

Table 2
The case evidence.

Continent	Country	Commodity	Mining method	Reported result (before → after)
Africa	Southern Africa	Diamonds	Open-cut	Processed kimberlite ≈ 2.1 Mt/yr → ≈ 2.9 Mt/yr (+38%) in the first year; operation returned to profit and mine life extended.
Africa	South Africa	Iron ore	Open-cut truck-and-shovel (greenfield ramp-up)	Monthly output ≈ 0.4 Mt → ≈ 1.0 Mt within year 1 (8.5 Mt for year 1); year 2 reached 11.6 Mt, ≈ 29% above original design capacity.
Africa	South Africa	Iron ore	Open-cut truck-and-shovel (mature operation)	Annual output ≈ 23 Mt → > 29 Mt over three years, against an original design capacity of 16 Mt.
Africa	South Africa	Coal	Open-cut truck-and-shovel	All-time production records within four weeks of constraint identification; sustained output moved above budget.
Africa	South Africa	Coal	Underground continuous-miner (bord-and-pillar)	Average daily output ≈ 14,500 t → ≈ 18,000 t; single-day record 27,000 t; ≈ 10% above the global continuous-miner benchmark; *70 consecutive days incident-free during the record period.*
Africa	Zambia	Copper	Underground cut-and-fill	Development rate approximately doubled; production reached record levels within nine months, at the lowest unit cost in the operation's history.
Africa	South Africa	Platinum	Underground conventional stopping	First-phase output gains of 11–33%; best sustained results 31–63% above pre-intervention output (range across operations).
Africa	South Africa	Gold	Deep-level conventional / long-haul	First-phase output gains of 21–30%; best sustained results 56–72% above pre-intervention output (range across operations).
South America	Chile	Copper	Open-cut	Daily throughput ≈ 180,000 t → ≈ 240,000 t (+30%) within three months.
Australia	Australia (NSW)	Coal	Underground longwall	Balanced-capacity average (2018–2023) 2.4 Mt/yr → 3.2 Mt in 2024 (+33%), with no new capital expenditure and no additional personnel; examined in detail in Section 5.2.

documentation, and are subject to the evidential limits stated above and in Section 8.

5.2. *The NSW longwall mine case study*

A good illustration of the impact of flow on production output, variability, and safety involves a longwall coal mine in New South Wales. From 2018–2023, operating in accordance with industry best practice and a conventional balanced-capacity approach, the annual production averaged 2.4 million tonnes. In August 2023, coinciding with the start-up of Longwall S2A, the mine implemented a flow approach. Three changes were made: the Coal Handling and Preparation Plant (CHPP) was identified as the system constraint; the operating model shifted from 'push' to 'pull', pacing production according to what the CHPP could process; and the ROM stockpile was used as a managed buffer with clear triggers.

Production in 2024 reached 3.2 million tonnes, a 33% increase with no new capital expenditure and no additional personnel. The General Manager reported that the mine was running better than ever, and that visitors had commented on the high standards of operation being maintained.

The effect on production stability is shown in Fig. 5, which contrasts the daily ROM output before and after the intervention. Before the intervention, the daily profile is the characteristic saw-tooth of stop-start flow: output fluctuates widely from day to day, frequently far below budget, and only intermittently near the operation's demonstrated capability. After the intervention, the profile is not only higher but also visibly tighter. The daily average rises toward the system capability while the day-to-day scatter around it falls sharply. As Fig. 5 shows, the gain is therefore twofold. The production increase, and central to the safety argument of this paper, has substantially less variability. It is the second of these, the contraction of daily variability, which is linked to safety performance in Section 5.4.

5.2.1. *Safety outcomes: three independent vantage points*

5.2.1.1. Management consensus. By the end of the first full year of flow-room operation, the management team's position was:

"Many of the management team believe that the excellent safety performance is the result of the new way of working."—Management team, recorded in a client-feedback compilation, 2024.

The same compilation captures the Operations Manager's description of the shift in his own working day: from continuous firefighting, to having the headspace to discharge his statutorily mandated responsibilities. Underground operations safety in NSW rests heavily on the Mine Operator and statutory officials being able to attend to their formal duties under the Work Health and Safety (Mines and Petroleum Sites) Act 2013, rather than being absorbed by short-cycle production crises. This is an important positive change.

"Since September [2023] [the Operations Manager] says he can focus on discharging his mandated responsibilities - there is no more firefighting."—client-feedback compilation, 2024.

5.2.1.2. Workforce response. In an underground coal mine, the workforce sees risk earlier than the regulator and earlier than the leadership team. Two separate, independent observations from the workforce are on record:

"For the first time, [we] feel management and workers are working together."—Mining and Energy Union representative, 2024.

"It's great coming to work, everyone's more relaxed, and we're working together."—underground crew member, 2024.

These observations matter because the dominant proximate cause of injuries in underground coal mining is not equipment failure but human-system interaction under pressure: rushed handovers, switching tasks, and decisions taken under time stress. A workforce that describes the operation as "relaxed" and "working together" is, in safety terms, a workforce in which those proximate causes are less prevalent.

Fig. 5. ROM variability before and after (balanced capacity vs flow-based chain).

5.2.1.3. The constraint owner. The CHPP Manager, who owns the identified constraint and runs the flow room, was asked directly how the intervention had changed safety at the mine. His response was deliberately conservative on causation but explicit on the causal pathway:

"It's probably too early to claim credit for safety results, but the feedback from crews and staff has been clear—the flow room has changed how we work. It gives everyone clarity on what matters most over the next 48 h, so we are not constantly chasing spot fires. Because we focus on what affects the bottleneck, everything else can be planned calmly, without the usual switching tasks. That leads to better planning, fewer rushed decisions, and more controlled execution, which is where most safety issues tend to creep in. Regulators often say short-term decisions cause 75% of injuries—so reducing that has real impact."—CHPP Manager, 18 March 2025.

The same interview supplies a concrete illustration. When the longwall hit poor roof conditions during a period of one-piece flow operation, the flow room forward projection of ROM buffer levels signaled the developing risk days in advance. Similar conditions would historically have triggered a rushed escalation. Instead, dozer resourcing, maintenance windows, and CHPP feed strategy were re-aligned across departments before the longwall needed to be pushed. In the constraint owner's words, the situation was "handled without panic-aligned and engaged". This is the operational signature of flow as a safety mechanism: variability is absorbed by the buffer and by planning time, not by people forced to work faster under pressure.

Read together, the production data used in this study and the safety evidence summarized here are consistent with the broader argument that flow-based operations reduce variability not only in throughput but in the demand placed on managerial attention, and that this reduction in attention load is where the safety gains originate. This mine experience does not, on its own, prove the safety claim at the level of a frequency rate, but it provides converging qualitative evidence from three independent vantage points (management, workforce, and the constraint owner), set against a sector backdrop in which injury frequency was deteriorating.

5.3. The safety–production link

Across more than 90 interventions, a consistent pattern has been observed: when a flow state is reached (higher average output with markedly lower daily variability), operators have typically reported concurrent improvements in safety performance. In a number of cases, site management regarded the safety improvement as the more valuable of the two outcomes [8,9,17].

The improved production and process predictability relieves end-of-period pressure on supervisors, resulting in longer-lasting mine plans, more effective planning, and safer delegation. Managerial span of attention expands, and without prompting, more time is spent in daily cross-functional meetings on safety. Management can practice what Hopkins [18] terms 'chronic unease'—a hallmark of HRO behavior.

Across the intervention portfolio, production increases have averaged about 20% (range 10–40%) with the same resources, and operators have reported schedule realization rising from below 50 to above 80%. These figures are operator-reported and uncontrolled (see Section 8) [9,17]. Equipment availability improves because proactive maintenance becomes much easier to schedule.

5.4. The Rasmussen-Cook boundary shift

Drawing on more than 20 years of mining experience, this article proposes that stabilizing flow using the TOC principles not only narrows the operating cluster (moving from LRO to HRO characteristics). It also shifts the acceptable performance boundary, along with the new operating point, further to the left of the balanced-capacity operating point, thereby allowing both safety and production to improve simultaneously [10,11]. The freed-up time, attention, and organizational capacity enable full focus on HRO principles and critical control management, further narrowing the operating cluster.

If correct, this would deliver a qualitatively different outcome from simply adding more controls and should make catastrophic boundary excursions less likely. The threat-rigidity effect [13] predicts that balanced-capacity mines in chronic perturbation operate on a permanent, low-grade threat-rigidity response that structurally prevents HRO behavior. Restoring flow would, on this argument, reverse that pattern.

The pattern across these interventions suggests that the safety gains are not a by-product of the production gains: both follow from the same underlying change—the move from an unstable to a stable operating state.

6. Possible unintended consequences of the 2024 amendments

6.1. Overview of the RSHLA Act 2024

The RSHLA Act 2024 represents the most notable regulatory shift in Australian mining safety in a generation [19]. It introduces mandatory critical control management into the safety and health management system and the principal hazard management plan, facilitates HRO principles, broadens accountability definitions to include critical controls within officers' due diligence obligations, and introduces enforceable undertakings. Full compliance was required by 1 June 2026.

The legislation is correct in what it seeks to achieve. The industry does need proactive critical control management. It does need HRO principles. It does need expanded accountability. The architects of this legislation have correctly identified that the industry must move from reactive to proactive safety management [1,19].

6.2. The compliance paradox

The question this paper considers is not whether the legislation is right—it is—but whether the typical operating mine can effectively support what the legislation demands. Consider what happens when the new critical control management requirements meet an unstable operational platform.

Consider an open-pit operation that has 15–20 documented critical controls across its principal hazard management plans. Suppose a control finding during the morning shift triggers a cascade: a pit supervisor identifies segregation bunds at 40% of the required tire diameter, whereas the standard specifies 50%. The escalation protocol requires operations to stop while corrective action is taken, and where the site senior executive assesses that the deficiency has the potential to cause a significant adverse effect to safety or health, the deficiency is reportable as a High Potential Incident under the existing trigger.

The haul road must now be closed. The trucks are rerouted to an alternative circuit, adding 12 min per cycle. This drops the haulage capacity by 30%. The crusher feed rate drops. The processing plant begins cycling. The ROM stockpile, already low because the operation intentionally runs with minimum buffers, cannot bridge the gap. The operations manager redirects the grader crew from scheduled road maintenance to bund remediation. This defers road maintenance.

By afternoon, the operation is 800 tonnes behind the target, and recovery pressure builds. The afternoon supervisor accelerates truck cycles and decides to defer pre-start inspections on two loaders, each of which is a critical control verification. The following morning, the deferred road maintenance has led to surface deterioration and a fresh near-miss, which on the same trigger (potential to cause significant adverse effect) becomes a reportable HPI. Another stoppage. The production plan needs to be adjusted on the fly, while the stable planning horizon shortens. At some point, last-minute coordination of activities becomes inevitable.

This is the compliance paradox: on an unstable operational platform, increased safety oversight is able to create the conditions that make safety compliance harder. In balanced-capacity mines, more effort could yield worse outcomes—not because the controls are wrong, but because the foundation cannot support them [7,20].

Weick's [15,21] concept of 'enacted sensemaking' describes precisely this trap. Actions intended to manage a situation often intensify it. The supervisor who stops production to address a degraded control is doing the right thing, but the resulting cascade creates the pressure that degrades the next control. The act of compliance itself destabilizes the system.

6.3. The HRO preconditions degraded by balanced capacity

The new legislation requires something that balanced-capacity operations structurally cannot deliver: spare managerial capacity (Goldratt's "managerial span of attention"). Weick and Sutcliffe [14] were explicit that the five HRO principles depend on 'collective mindfulness', a shared quality of organizational attention requiring time, bandwidth, and organizational slack. When those resources are consumed in firefighting, mindfulness collapses into mindlessness, where actions become routine and familiar events are assumed to be understood even when they are not.

Balanced-capacity mines may have difficulty in supporting what the legislation requires. The quickest way to address this is to stabilize production flow.

7. Research proposition

Section 2 left a question open: whether production pressure degrades safety and, if it does, how such an effect could be detected. The Queensland Parliament's Transport and Resources Committee, unable to resolve that question on the evidence before it, recommended that independent research be commissioned "into the impact of coal production rates on safety risk management" [5]. This section sets out a way to conduct such an investigation: a primary hypothesis, a set of secondary hypotheses, and a research design, all built on one proposition that the variable to measure is not the production rate itself but the production variability that balanced-capacity operation generates, and that this variability can be computed from the data that mining operations already collect.

7.1. Primary hypothesis

The primary hypothesis is that ROM production variability correlates with safety incident frequency and severity across mining operations. The article develops a heuristic (based on more than 90 implementations) that treats a chain of operations as balanced-capacity (with moving bottlenecks) when the average ROM over a 12-month period is 25% or more below the best daily ROM (system capability).

Where mining operations exhibit a gap of 25% or greater (implying a balanced-capacity chain), they will show higher incident rates than operations with smaller gaps. The size of the gap, combined with daily ROM variance, is therefore expected to correlate with safety incident rates.

The diagnostic calculation is shown in Fig. 6. From a year of daily ROM data, the best daily ROM is taken as a practical estimate of the operation's true system capability, whereas the 12-month average is its realized output. The gap between the two, expressed as a percentage of capability, is the diagnostic. As Fig. 6 illustrates, a gap of 25% or more is the threshold at which we infer that an operation is being run as a balanced-capacity chain with a moving bottleneck rather than as a flow-based system.

7.2. Secondary hypotheses

First, the correlation between production variability and safety incidents is generated by the mechanisms described in Section 3—specifically, the managerial attention overload (increased coordination loading), reduced situational awareness, and degraded communication bandwidth [13,14,22].

Second, operations that change from balanced capacity to flow-based operations (reducing the +25% gap to less than 20%) will show concurrent improvements in both production stability and safety metrics.

Third, for operations with balanced-capacity characteristics, the new critical control management requirements under the RSHLA Act 2024 will amplify the stoppage spiral, leading to no improvement or even worse safety outcomes. Flow-based operations will, conversely, show safety improvement.

7.3. The Brady Review as starting point

The Brady Review [1] provides data linking economic performance to safety outcomes in Queensland mines. This article proposes

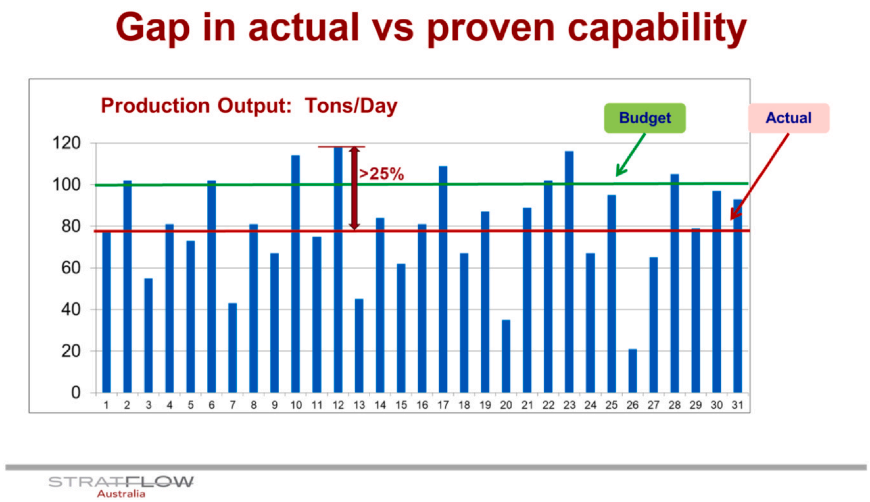

Fig. 6. Calculating the > 25% gap using daily ROM data.

extending that analysis by replacing the broad 'economic performance' metric with the more specific 25% gap diagnostic. This refinement would test whether production variability is a stronger predictor of safety outcomes than economic performance alone, and whether the mechanism described here (mental overload, HRO precondition failure, and stoppage spiral) explains the relationship.

7.4. Proposed research design

7.4.1. Historical analysis of output gap and safety incidents

Using historical mining records to correlate the size of the ROM production gap and the daily ROM variance to safety incident rates could add further support for the Brady Review findings on profitability and mining fatalities.

7.4.2. Observational study during RSHLA Act 2024 implementation in Queensland mining

An observational study is proposed across Queensland mining operations during the RSHLA Act 2024 implementation period (2024–2027). The study would collect daily ROM production data and calculate the gap metric and ROM variability for each participating site, correlate these with safety incident data (HPI reportable incidents, LTIs, and near-miss reports) using trend analysis, adjusting for mining method and site size, and track the impact of RSHLA Act 2024 compliance implementation on both production stability and safety metrics. This observational study is offered as one concrete design for the independent research recommended by the Transport and Resources Committee [5].

8. Limitations

This paper is based on data from productivity interventions across a wide range of mining operations over 20 years. The stoppage spiral mechanism and the HRO implementation gap may exhibit different dynamics across jurisdictions. While the underlying system dynamics principles are universal [7,20,23], specific operational details that trigger the spiral may vary.

The evidence presented is practitioner observation, not controlled research. The more than 90 implementations referenced herein were conducted as commercial consulting engagements, not designed as research studies. They were not randomized, did not follow strict research protocols, and were not independently verified [8,9].

Selection bias exists in the implementation portfolio on which this article draws because organizations that engage external consultants for improvement interventions share certain characteristics, such as being already dissatisfied with their own performance, having a budget for consulting, and a willingness to attempt operational change.

Although the practitioner evidence strongly suggests the hypothesis presented that flow interventions translate into both productivity and safety improvements is correct, the causal direction has not been formally established. The 25% gap diagnostic, while empirically consistent, has not been validated. Prospective studies would be required to establish predictive validity.

These limitations should guide the research to enable the move from practitioner observation to validated knowledge.

9. Conclusions

Between 2003 and 2013, the Australian mining fatality rates fell by approximately 80%, but thereafter the improvement stagnated. In spite of significant ongoing investment in safety systems, regulation, and cultural programs, the industry has averaged approximately seven fatalities per year over the past decade. The Brady Review [1] identified a fatality cycle linked to economic (operational) pressures. This paper has argued that the plateau is not a failure of commitment to regulatory effort, but a predictable consequence of how production itself is organized.

Practitioner experience, albeit without formally confirmed causality, across more than 90 interventions on three continents suggests that the plateau can be breached. When operations move from balanced capacity to flow-based production, a recurring pattern has been observed: output rises 10–40%, ROM variability falls markedly, and operators report an accompanying improvement in safety, even without dedicated safety initiatives [8,9,17]. The NSW longwall mine case study vividly illustrates this pattern: a 33% production increase in the first year of flow-based operation introduction accompanied by a shift from chronic instability to sustained operational predictability.

The mechanisms proposed here, by which flow-based operations would improve safety, are consistent with established theory. Stabilizing the production flow at a higher output removes the chronic pressure to force more work through the mine than it can handle. Brady [1] linked economic pressure to fatality clusters. Miners report that firefighting reduces mine plans' hold, end-of-period surges diminish, and preemptive maintenance becomes feasible. Managerial span of attention expands—supervisors and middle managers report being able to discharge their responsibilities properly, often for the first time. The threat-rigidity response [13] that chronic perturbation induces gives way to the psychological conditions that HRO behavior requires [14]: sustained attention, reflective thinking, situational awareness, and adaptive capacity. In the language of the Rasmussen-Cook boundary framework [10,11], the acceptable performance boundary shifts towards the left while the operating point shifts toward both greater safety and higher productivity.

Queensland's RSHLA Act 2024 represents the most significant advance in Australian mining safety governance in a generation. For operations already in a flow state, the critical control management requirements of the RSHLA Act 2024 should further accelerate safety improvements because compliant mines will possess the operating stability, cognitive bandwidth, and organizational capacity demanded by effective critical control management. However, for the majority of mines that continue to operate as balanced-capacity entities, the expanded compliance requirements risk triggering the compliance paradox described in Section 6. Adding reporting obligations to an already unstable production base may amplify the instability the legislation seeks to address—cascading stoppages across interdependent production units, degrading the cognitive preconditions for HRO behavior, and consuming the managerial attention needed to enact critical controls effectively. The legislation is correct in principle; the question is whether the operational platforms on which it must be applied can support its requirements.

This analysis leads to a testable proposition: if the relationship between production stability and safety performance is as strong as practitioner experience suggests, it should be visible in historical data. Mining safety researchers could use the 25% gap diagnostic described in Section 7 to classify operations by their degree of balanced-capacity behavior and then compare their safety records. The RSHLA Act 2024 implementation period, from 2024 to 2027, presents a time-limited natural experiment: do mines with lower ROM variability respond better to the new critical control requirements, and do mines with high variability experience the compliance paradox? The data that mining operations already collect will be sufficient to answer these questions.

Three recommendations follow.

For regulators: If confirmed by research, production stability should be explicitly recognized as a precondition for effective critical control management, and guidance should be developed to help balanced-capacity operations identify and address the instability that undermines control effectiveness.

For operators: Instead of seeing excess capacity as waste, it should be reframed as protective capacity, i.e., a buffer that enables the system to absorb disruptions and increase productivity while improving safety. Performance measurement should shift from departmental utilization targets to system-level throughput and stability benchmarks.

For researchers: The window from 2024 to 2027 should be treated as an opportunity to establish whether production stability is a leading indicator of safety performance.

None of the above requires regulatory reform, productivity sacrifice, or safety compromise. It requires an operational decision to move from balanced-capacity to flow-based production. Sequencing will help, although it is not a prerequisite: stabilize the production flow first, then compliance systems can function as intended, and HRO principles can take root in an organization freed from chronic firefighting. The Australian mining industry broke through one barrier to enhanced safety between 2003 and 2013 through regulation and cultural change. Breaking through the current safety plateau will require something different: changing the way production itself is organized. If the evidence confirms what practitioners have observed for two decades, then the path to reduced fatalities runs not through more controls, but through more stable operations.

CRediT authorship contribution statement

Hendrik Lourens: Writing – review & editing, Writing – original draft, Methodology, Investigation, Formal analysis.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author used Opus 4.6 in order to perform research on the applicable literature relating to the observations being shared. After using this tool/service, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

Declaration of Competing Interest

The authors declare the following financial interests (e.g., any funding for the research project)/personal relationships (e.g., the author is an employee of a profitable company) which may be considered as potential competing interests: The author is the founder and CEO of Stratflow Australia, a consultancy that implements flow-based operations in mining using Theory of Constraints principles. Stratflow is the primary commercial beneficiary of the intervention approach described in this article. While TOC consulting is practised widely by an international community, the specific Flow Room methodology and implementation approach described here are Stratflow's proprietary practices. Readers should interpret the practitioner evidence presented with this context in mind.

Acknowledgements

The essential practitioner methodology and implementation portfolio described in this article were developed jointly with the late A. van Niekerk (d. 2022), whose role in the theory and practice of flow-based operations in mining is acknowledged throughout.

References

[1] S. Brady, Review of all fatal accidents in Queensland mines and quarries: 2000–2019, Resources Safety and Health Queensland, Brisbane, Australia, 2020.

[2] Safe Work Australia, Mining dashboard, Safe Work Australia, Canberra, Australia, 2024.

[3] University of Queensland Sustainable Minerals Institute (D. Cliff), How safe are Australia's mines? New analysis shows reform has been stalled for a decade, 18 March 2024, Available at: https://smi.uq.edu.au/article/2024/03/how-safe-are-australia%E2%80%99s-mines-new-analysis-shows-reform-has-been-stalled-%C2%A0decade (accessed 11 April 2026).

[4] N. Flesher, M. Moyo, S. Rehbach, E. van Niekerk, Productivity across the global mining sector is starting to improve, McKinsey & Company, June 2018.

[5] Transport and Resources Committee, Inquiry into coal mining industry safety, Report No. 29, 57th Parliament, Queensland Parliament, Brisbane, Australia, 2023.

[6] E.M. Goldratt, The goal: A process of ongoing improvement, North River Press, Great Barrington, MA, USA, 2014.

[7] J.D. Sterman, Business dynamics: Systems thinking and modeling for a complex world, McGraw-Hill, Boston, MA, USA, 2000.

[8] A. Van Niekerk, A TOC practitioner's story of 15 years of TOC applications in the mining industry, TOCICO World Conference, Cape Town, South Africa, 2015https://www.tocico.org/mpage/vanNiekerk9_7IND (accessed 19 July 2026)..

[9] H. Lourens, J. Blakemore, Turning mining performance around: moving from efficiency to effectiveness, Aust. Min. (May 2017) 40–42 https://www.stratflow.com.au/_files/ugd/4e6d7e_521521aa68e9465f93bbd58859e7e4a5.pdf (accessed 19 July 2026).

[10] J. Rasmussen, Risk management in a dynamic society: a modelling problem, Saf. Sci. 27 (2-3) (1997) 183–213, https://doi.org/10.1016/S0925-7535(97)00052-0.

[11] R.I. Cook, J. Rasmussen, Going solid": a model of system dynamics and consequences for patient safety, Qual. Saf. Health Care 14 (2) (2005) 130–134, https://doi.org/10.1136/qshc.2003.009530.

[12] J.B. Morrison, R.L. Wears, Modeling Rasmussen's dynamic modeling problem: Drift towards a boundary of safety, Cogn. Technol. & Work. 24 (2022) 127–145, https://doi.org/10.1007/s10111-021-00668-x.

[13] B.M. Staw, L.E. Sandelands, J.E. Dutton, Threat-rigidity effects in organizational behavior: a multilevel analysis, Adm. Sci. Q. 26 (4) (1981) 501–524, https://doi.org/10.2307/2392337.

[14] K.E. Weick, K.M. Sutcliffe, Managing the Unexpected: Resilient Performance in an Age of Uncertainty, 2nd ed., Jossey-Bass, San Francisco, CA, USA, 2007.
[15] K.E. Weick, The collapse of sensemaking in organizations: The Mann Gulch disaster, Adm. Sci. Q. 38 (4) (1993) 628–652, https://doi.org/10.2307/2393339.
[16] H. Lourens, A dual management operating system to improve digitalisation and automation outcomes in mining, Austmine (30 March 2021) Available at: https://austmine.com.au/Web/News-Resources/Articles/2021/March/A_Dual_Management_Operating_System.aspx.
[17] H. Lourens, P. Moynagh, How mining leaders can take charge to improve safety, productivity and reduce stress, AusIMM Bull. (November 2020) (Available at: https://www.ausimm.com/bulletin/bulletin-articles/how-mining-leaders-can-take-charge-to-improve-safety-productivity-and-reduce-stress/).
[18] A. Hopkins, Disastrous decisions: The human and organisational causes of the Gulf of Mexico blowout, CCH Australia, Sydney, Australia, 2012.
[19] Queensland Government, Resources Safety and Health Legislation Amendment Act 2024, Queensland Government, Brisbane, Australia, 2024 (accessed 19 July 2026)..
[20] P.M. Senge, The fifth discipline: The art and practice of the learning organization, 1st ed., Doubleday, New York, NY, USA, 1990.
[21] K.E. Weick, K.M. Sutcliffe, D. Obstfeld, Organizing for high reliability: Processes of collective mindfulness, Res. Organ. Behav. 21 (1999) 81–123.
[22] D. Kahneman, Thinking, fast and slow, 1st ed., Farrar, Straus and Giroux, New York, NY, USA, 2011.
[23] D.H. Meadows, Thinking in systems: A primer, Chelsea Green Publishing, White River Junction, VT, USA, 2008.