

Risk 4 Managers

Groundbreaking Risk Management for the Mining Industry

From hindsight to foresight: what four decades underground and in the boardroom teach us about managing risk before it manages us

Albert C van der Vyver

© Copyright 2026 by Albert C van der Vyver, SheRisk. All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior written permission of the author.

The case studies in this book are drawn from real events in the South African and international mining, industrial and construction sectors. Identifying details have been withheld where appropriate. The views expressed are those of the author and are offered to advance the practice of risk management. They do not constitute legal advice.

To the men and women of the mining industry who go to work every day in some of the most demanding environments on earth — and to every leader, engineer, supervisor and risk practitioner who makes sure they come home.

Contents

Introduction

PART I: Foundations

Chapter 1: The Long Road of Risk

Chapter 2: Metamorphosis — How Risk Reinvents Itself

Chapter 3: The Changing Nature of Risk — Ten Shifts for the Modern Mine

PART II: The Mining Arena

Chapter 4: Lessons Forged in Tragedy

Chapter 5: Governance, Law and the Rainbow Nation

Chapter 6: Building on Solid Ground — Industrial, Construction and Project Risk

PART III: The Human Factor

Chapter 7: The Enigma of Human Conduct

Chapter 8: Forensic Investigation — From Blame to Understanding

PART IV: Mastering Controls

Chapter 9: Clarity of Control

Chapter 10: Key Risk Indicators — Measuring What Matters

Chapter 11: Bow Tie Analysis Done Right

Chapter 12: Critical Control Management — Restoring the Window

Chapter 13: Risk Response Plans — Ready Before the Event

PART V: Culture, Framework and the Future

Chapter 14: Appetite for Risk Culture

Chapter 15: A Robust Risk Management Framework

Chapter 16: The Future Risk Professional

Conclusion: The Groundbreaking Agenda

References

Acknowledgements

About the Author

Introduction

> *Imagine a mine where the next serious event has already been investigated — before it happens. Where every control is measured, not assumed. Where the risk matrix is a tool for priorities rather than a scoreboard, and where the people who face the hazard and the people who fund its control share one language. That mine is not a fantasy. It is the subject of this book.*

Did you know that in 2025 the South African mining industry recorded 41 fatalities, the lowest number in its history, and that by mid-September 2026 the year-to-date figure had already passed that total? A record low followed within nine months by a reversal is not a statistical accident. It is the clearest possible demonstration of the central message of this book: **a declining accident rate tells us about the past; it says nothing about our readiness for the future.** History does not predict. It repeats — unless we change the way we look ahead.

The mining industry deserves enormous credit. Deep-level mines that were once regarded as among the most hazardous workplaces in the world have achieved safety improvements that few would have believed possible a generation ago. Fatality rates have fallen dramatically since the early 1990s, when hundreds of mineworkers lost their lives every year. Engineers, production managers, safety practitioners and organised labour have built systems that save lives every single shift. That achievement is the foundation on which this book stands.

Yet every practitioner knows the frustration of the plateau. The same events recur. Conveyances still leave their guides. Tailings structures still fail. Vehicles still meet pedestrians. After each event we revise the procedure, retrain the operator, update the risk assessment and add a control — and the cycle begins again. The opportunity in front of the industry is to break that cycle by moving the whole of risk management **one step to the left on the event timeline**: away from proof after the event, and towards understanding vulnerability, pre-conditions and the release mechanisms of energy before anything goes wrong.

Why This Book, and Why Now

This book is written by a practitioner, adviser and risk owner, not by an academic. Its perspective has been shaped over many years as an employee, supervisor, manager and consultant in the mining industry, in the space between obligation and accountability where real risk decisions are made. It offers a practical interpretation of risk management, tested in shafts, open pits, smelters, plants and boardrooms.

Three observations prompted it.

- **Resilience plans are often written after the event.** Many organisations have developed impressive contingency and recovery plans. A dam wall fails, and a plan is in

place to recover. The leap forward is to build plans that address *susceptibility* to the event, so that resilience begins with prevention and is proven through simulation rather than assumed.

- **Risk numbers have been asked to do work they were never designed for.** For three decades the industry has used the risk matrix to "prove" that controls reduce risk. The matrix is an excellent prioritisation tool. Used as a measure of safety, it can create a comfortable illusion. Real assurance comes from measured control effectiveness.
- **The people who manage risk deserve the science behind it.** Neuroscience, behavioural science, forensic investigation, reliability engineering and data analytics have matured enormously. Bringing them into everyday risk management is the most powerful lever the industry has not yet fully pulled.

What You Will Discover

This book is organised into five parts.

- **Part I — Foundations** traces how risk management evolved, why it keeps reinventing itself, and the ten shifts that define the modern discipline.
- **Part II — The Mining Arena** examines the lessons the industry has learned from tragedy, the roles of the boardroom and the law, risk in South Africa's "Rainbow Nation", and the fundamentals of industrial, construction and project risk.
- **Part III — The Human Factor** explores human conduct, error, habits and the psychology of "casino safety", and makes the case for forensic-quality investigation.
- **Part IV — Mastering Controls** is the technical heart of the book: control effectiveness, key risk indicators, bow tie analysis done properly, critical control management and response planning.
- **Part V — Culture, Framework and the Future** connects risk appetite, trust and culture to a robust risk management framework, and describes the risk professional the industry will need in the decades ahead.

Each chapter closes with key takeaways, and the Conclusion gathers them into a practical agenda for action.

About the Author's Experience

The author has spent 44 years in risk management. In March 1997 he founded SheRisk, a risk management training and consultancy practice serving the mining, construction, industrial and petrochemical sectors from offices in Bronkhorstspruit (Gauteng) and Riversdale (Western Cape), with a new office in Kathu in the Northern Cape. He has lectured in risk management at the University of Pretoria's Faculty of Mining Engineering and helped to develop and present the c-MiRM course through Enterprises University of Pretoria. Many of the examples in this book are

real events in which he was directly involved as an investigator, facilitator or adviser. Their context has been preserved because their lessons are real.

What This Book Will Do for You

If you are an executive, it will help you ask the questions that reveal your organisation's true vulnerability, rather than the comfortable ones that confirm what you already believe. If you are a mine manager or engineer, it will give you practical methods to measure whether your controls work. If you are a risk or safety practitioner, it will equip you to move from compliance to influence. And if you are a student of the discipline, it will show you where the profession is heading.

Some of what follows will challenge long-held practices. That is intentional. Every idea in this book is offered in the same spirit in which the industry achieved its greatest safety gains: a refusal to accept that the next event is inevitable.

Let us begin where risk management itself began — a very long time ago.

PART I: Foundations

Before we can change the way risk is managed, we need to understand where our habits came from. Part I follows risk management from its ancient roots to the modern enterprise, shows how the discipline keeps reinventing itself, and sets out the ten shifts that separate yesterday's practice from tomorrow's.

Chapter 1: The Long Road of Risk

> *"Risk and time are opposite sides of the same coin, for if there were no tomorrow there would be no risk." — Peter L. Bernstein, Against the Gods: The Remarkable Story of Risk (1996)*

Anyone who sets out to find the "true origin" of risk management soon discovers a rich and contested landscape. Renowned specialists have documented the subject thoroughly, but each places the starting point somewhere different. Some trace the discipline to the birth of probability theory, others to the commercial use of the word *risk*, others again to the rise of insurance. Each view holds part of the truth. The most useful conclusion for a practitioner is a simple one: risk is not only about probability, and it is not only about opportunity. It is about people making decisions under uncertainty — and people have been doing that since the beginning of time.

From Eden to the Ark

Consider the oldest story many of us know. In the Garden of Eden there was an identified tree, a clear instruction, a persuasive voice and a decision. The outcome was uncertain to those who made it, and the consequences were profound. Was that decision not risk-based? Was uncertainty about the outcome not present? In a very real sense, the discipline we now call risk management is as old as human choice.

The same can be said of Noah's Ark: a hazard foreseen, a structure designed to withstand it, a long lead time before the event, and a response plan that included every living thing that needed protecting. It remains one of the most memorable illustrations of preparing *before* the trigger event rather than after it.

Defence in Depth Is an Ancient Idea

As far back as history is recorded, communities protected their cities with walls, moats, watchtowers and gates — layers of defence that would be familiar to any modern barrier analyst. Political and military rivalry was always opportunity-based: what can we gain, and what is at stake? Remarkable foresight and deployment of control mechanisms were demonstrated in ancient strategy. And the anxiety of predicting tomorrow pushed some people to become explorers. Nomads formed new clans, languages and futures for their followers. Today's risk disciplines have their own trailblazers and catalysts in exactly the same way.

The construction of the Egyptian pyramids shows risk thinking maturing. Whatever theory one prefers about how they were built, the questions the builders had to answer were risk questions. Who made up the labour force, and what was its culture? What did the stones weigh? How would weather affect the work? Would the structure remain stable for millennia? Each of these belonged in the design and execution of the task — and the pyramids are still standing.

Military history adds another lesson. Through the later Roman Empire and the early medieval period, every encounter revealed a new threat and prompted a new response: heavily armoured cavalry, the stronger Hunnic composite bow replacing the older Scythian design, longer swords, and the progressive replacement of scale armour by mail and lamellar armour. When the stirrup finally arrived, it transformed what cavalry could achieve. Each improvement was a control upgrade prompted by a lesson learned.

During the High Middle Ages, after AD 1000, technological and agricultural innovation allowed trade to flourish, the Medieval Warm Period increased crop yields, and the population of Europe grew significantly. Society organised itself into structured obligations — peasants owing rent and labour to nobles, knights owing military service to their overlords — and the founding of universities brought a philosophy that joined faith to reason. Centralised states reduced crime and violence. In other words, people built systems to manage uncertainty at the scale of whole societies.

> **Groundbreaking insight:** The control-improvement cycle is as old as civilisation. Every era has believed it was close to a defence that "fits all". Every era discovered a new threat. The mature organisation does not chase the perfect control; it builds the capability to keep learning faster than its hazards evolve.

This is where the philosophy of *realism* becomes useful. Realism holds that many of the mental shortcuts we call biases are not errors at all, but practical ways of dealing with a real world in which people speak, make mistakes, mislead, change their minds — and in which more time always brings more change. A risk manager who accepts that reality designs systems for people as they are, rather than as procedures wish them to be.

Words Matter: Danger, Threat and Risk

We often use *danger*, *threat* and *risk* as if they were interchangeable. They are not, and precise language is the first control in any risk assessment.

- **Danger** comes from the Anglo-French *daunger*, which carried the sense of "power to harm".
- **Threat** is recorded in Old English before the year 900, from *þreat*, meaning pressure or oppression.
- **Risk** entered English in the 1660s from the French *risque*, itself from the Italian *risco*. The commercial, insurance sense of the word appears by 1719, and the anglicised spelling *risk* is recorded from 1728. From the start, then, risk carried the meaning of a chance of loss in a venture undertaken for gain.

Games of chance, interestingly, are depicted in Egyptian tomb paintings from around 3500 BCE, but it was only after the development of a numbering system that people could calculate odds and simulate opportunity (Harper, *Online Etymology Dictionary*).

The Modern Milestones

The best-defined timeline of modern risk management is the work of H. Felix Kloman, a Fellow of the Institute of Risk Management (London). Reflecting on Peter Bernstein's *Against the Gods* as a fitting end to his list of milestones, Kloman stressed the importance of communication and warned against the "perils of excess quantification". His conclusion is worth carrying into every boardroom: whatever title we attach to the thinking process — risk management, enterprise risk management or strategic risk management — "it will continue to be a part of the human experience" (Kloman, 2009).

Georges Dionne, Professor of Finance and holder of the Canada Research Chair in Risk Management, shows that the formal *study* of risk management began only after the Second World War (Dionne, 2013). It grew hand in hand with post-war industrialisation, when engineers developed techniques to improve the effectiveness and reliability of process and manufacturing plants. Many of the techniques listed in ISO/IEC 31010 were born in that period:

- **Failure Mode and Effects Analysis (FMEA)** emerged from a United States military procedure in 1949.
- **Fault Tree Analysis** was developed at Bell Laboratories in 1962 for the Minuteman missile programme.
- **Hazard and Operability Studies (HAZOP)** were pioneered by Imperial Chemical Industries in the 1960s.

Understanding where these concepts came from — and what problems they were designed to solve — is essential for the profession. A technique applied outside its original purpose can still be valuable, but only if the practitioner knows what it was built to do.

History Informs; It Does Not Govern

The study of history has great academic value, and it teaches humility. But history belongs to the past. Risk lives in the future. The most important shift any risk manager can make is to use historical information to *understand mechanisms* — how energy is released, how controls degrade, how people adapt — rather than to set the direction of the future. Risk managers of tomorrow must prepare for sophisticated analysis in every aspect of the discipline, and risk management must take its place as a recognised, skilled profession.

Key Takeaways

- Risk management is as old as human decision-making; formal study of it began after the Second World War.
- Defence in depth, lessons learned and control upgrades are ancient ideas — the modern challenge is learning faster than hazards evolve.
- Precise language (danger, threat, risk) is the first control in any assessment.

- Know the original purpose of each technique you use.
- Use history to understand mechanisms, not to predict the future.

In the next chapter, we see how risk management keeps reinventing itself — and how some of its own inventions, from blame to compliance, can be transformed into strengths.

Chapter 2: Metamorphosis — How Risk Reinvents Itself

> *Imagine two mines with identical risk registers, identical procedures and identical legal appointments. One has gone five years without a serious event; the other has just recovered from a disaster. Which is safer? The honest answer is that the registers cannot tell you. What tells you is how each organisation thinks — and that thinking is the subject of this chapter.*

Risk management has mutated over many centuries, passing its traits from one generation of practice to the next. In antiquity, people faced danger under humanitarian and social constraints shaped directly by the uncertainty of life. Over time, that instinctive response matured into a sophisticated defence mechanism we now call enterprise risk management. Along the way, we also developed some habits we did not plan.

The Avatars We Created — and How to Transform Them

The aftermath investigations of man-made disasters reveal recurring patterns of organisational behaviour. Think of them as *avatars*: personas that the system adopts, often without anyone choosing them.

- **The avatar of blame** looks for the person who broke the rule. Its positive twin is **the avatar of learning**, which looks for the conditions that made breaking the rule likely.
- **The avatar of compliance** asks "did we do what was required?" Its positive twin is **the avatar of ownership**, which asks "did we do what was needed?"
- **The avatar of manipulation** — the newest arrival in the boardroom — adjusts resources and numbers to influence how outcomes look. Its positive twin is **the avatar of honest measurement**, which lets the data speak even when the message is uncomfortable.

Praise follows success just as blame follows failure, and both are driven by the same instinctive mechanisms. Recognising these avatars is the first step to transforming them. An organisation that names its own patterns gains the power to choose different ones.

> **Groundbreaking insight:** Blame, compliance and manipulation are not character flaws in individuals. They are system behaviours. Change the system's questions, and the avatars change with them.

The Brain at the Heart of Risk

The human brain sits at the centre of this metamorphosis. Industrial, forensic and criminal psychologists have studied how people think and act on matters that influence their cognition and emotions. Neuroscience, behavioural science, neuroeconomics and decision theory are closely connected disciplines, and together they offer an exciting window into the interplay between risk culture in the boardroom and the scientific blueprint of the mind.

The practical lesson is powerful. Risk management can be grounded in science rather than in perception, opinion or argument. The more we understand how people actually process uncertainty, the better we can design controls that work with the brain rather than against it.

Risk Homeostasis: The Inner Thermostat

Gerald Wilde's theory of risk homeostasis (Wilde, 1998) suggests that each person carries an inner "target level" of risk that they are willing to accept. When the environment changes, people adapt their behaviour to keep their perceived risk close to that target. Make a road feel safer, and some drivers will simply drive faster. In a boardroom, a workshop or a stope, people base decisions on their perception of risk measured against this inborn target, and they adjust to remove any inconsistency between the two.

The theory remains debated among academics, and more research is needed to establish reliable effectiveness criteria. But its value to the practitioner is clear. Risk homeostasis offers an alternative lens on what we usually call "behaviour" or "attitude", and its offsets can be examined as part of root cause analysis. If a new control makes people *feel* safer without making them safer, it may quietly shift risk somewhere else.

Legislation: The Floor, Not the Ceiling

Legislation is a modern marvel for setting guidelines and monitoring regimes. It establishes a minimum standard, protects the vulnerable and gives employees important rights, including the right to refuse dangerous work. Employers and employees both have a pivotal role in meeting these requirements.

The opportunity lies in remembering what legislation is designed to do. It sets the floor; it cannot build the house. Regulations, codes of practice and standard procedures are written to protect people, and they do. But when the volume of rules grows faster than understanding, the worker can be placed on a pedestal where personal initiative has little room to operate, and investigations tend to end with the finding that "someone broke a rule". The strongest organisations treat compliance as the starting line. They invest at least as much time in risk assessment as in investigation, and they invite statutory stakeholders into their risk assessments, not only their incident inquiries.

What the Pandemic Taught Us

COVID-19 offered a vivid lesson. When the pandemic reached countries outside China, legislation, lockdowns and movement restrictions were enforced quickly and firmly. In many communities, however, people lived up to ten in a home, with houses one or two metres apart, and the control strategy was overwhelmed by the realities on the ground. What ultimately changed behaviour was awareness — the numbers and images communicated through every media platform — and the homeostatic adjustment that followed.

The lesson for risk managers is this: good response strategies are pre-developed sets of actions, based on sound statistical analysis, with defined trigger points that enable effective management of significant events. *Pre-developed* means prepared after the possibility of the event has been identified, but before the event begins. Risk management anticipates events; it does not merely catalogue the ones we already know.

Take a closer look at your own risk assessments:

- Do they portray history, or do they describe future events that have not yet happened?
- Do they indicate *when* to expect an event, using measures of risk velocity?
- Do they show preventive and mitigating measures, together with their effectiveness against approved criteria?

The Fence Paradox

Pasquale Cirillo's essay *Of risk, fences, and unavoidable falls* (2018) describes a paradox every miner will recognise. Place people in a situation where failure is possible, and some will be challenged by their own minds to take the risk. Failing and learning is part of how humans develop. In the workplace, however, failure can mean injury. Legislators, stakeholders and authorities respond by building fences — regulations, directives and work stoppages. The fences are necessary. But they can also modify risk perception, so that people underestimate the danger on the other side. When a fence fails, it can fail for many people at once.

The way through the paradox is not fewer fences. It is better understanding: making the hazard visible, measuring the fence's condition continuously, and building a culture in which people respect the risk because they understand it, not only because they are told to.

Leading in Extraordinary Times

It is tempting to believe that we live in uniquely extraordinary times: an uncertain economy, conflict on the doorstep, a demanding industry. That framing creates pressure on owners and can justify a "state of exception" in which executives, like legislators, suspend, amend and approve new constraints at speed. The fence line looks correct right up to the moment of disaster.

When a tailings dam fails, the shock is profound. The failure of Vale's Córrego do Feijão dam at Brumadinho, Brazil, in January 2019 claimed 270 lives and led directly to the Global Industry

Standard on Tailings Management (2020). Closer to home, the collapse of the Jagersfontein tailings facility in the Free State in September 2022 claimed lives and swept away homes in the community below. In the aftermath of events like these, the questions are always the same. Where were the people who designed and monitored the preventive controls? Where were the owners, executives and stakeholders?

The most constructive answer is to build the moral obligation to protect fellow workers and neighbours into the culture *before* anything happens. When authorities and communities raise a concern about a structure, the mature response is to appoint the right expertise and provide transparent risk guidance — not to wait for the investigation. Risk does not fall because someone has been blamed. It falls when controls are understood, measured and maintained.

Designing for the Failure We Already Know

Some of the most useful lessons come from equipment we trust every day.

Spring-applied, hydraulic-release brakes are required on vehicles operating in many mining environments. They are a valuable control. Yet, like any engineered system, they can and will fail under certain circumstances by design. Over time, operators can become so confident in the brakes that they stop anticipating a runaway. Manufacturers, engineers and regulators all accept the residual risk. When an event occurs, the investigation can too easily conclude that "human error" was the root cause.

Now add a fire suppression system to the same vehicle: an automatic device triggered by rising temperature, releasing extinguishing medium through nozzles. Who decides where the nozzles go? Typically they are aimed at the engine compartment and electrical systems. But heat generated in the brakes can travel into the wheel assembly and ignite hydraulic fluid there. After a fire, the investigation reveals that the nozzles were in the wrong place. The risk was accepted by design; it was simply waiting for the right circumstances.

Conveyances tell the same story. After the Vaal Reefs disaster of May 1995, when a runaway locomotive fell down a shaft onto a double-deck cage and 104 mineworkers lost their lives, the industry resolved that nothing similar would happen again. Coming as the Leon Commission of Inquiry into mine safety completed its work, the tragedy added urgency to the reforms that produced the Mine Health and Safety Act of 1996. Yet conveyances still fall down shafts in the industry, often without injury, and those near-misses are priceless information. Each one tells us something about a system that we can improve before the next event finds a cage full of people.

> **Groundbreaking insight:** The most dangerous controls are the ones we trust completely. Every engineered control has known failure modes. The mature organisation lists them, monitors them and designs the next layer of defence around them — before an investigation does it for them.

Regulation in the Age of Rapid Change

Technology now moves faster than legislation can be drafted. A regulation aimed at a rapidly changing industry may be outdated by the time it is promulgated. This is not a criticism of regulators; it is a structural reality. Practically, risk managers identify the risk, evaluate its impact and set up controls to monitor it. Regulations, by their nature, tend to follow events: an event, an investigation, a new rule.

The opportunity is for the profession to lead. When risk specialists assume that risk is real and present — rather than accepting that regulation alone governs it — they prepare for the next event before it happens. Mathematical calculation and simulation help us understand the nature of exposure, but the subject always remains human intervention and human judgement. Model rigorously, and then apply common sense.

The Positive Impact of the Bad

Human society transforms constantly, and with it the social structures in which we manage risk. Transformation carries two embedded risk factors: uncertainty and incompletely understood outcomes. Ulrich Beck, the German sociologist and one of the most cited social scientists in the world, argued that metamorphosis is not about the side effects of the good, but about *the positive impact of the bad* (Beck, 2016). Catastrophic events, once thought unthinkable, force the world to see itself differently.

That is precisely the role of the modern risk professional: to analyse possible outcomes, recognise the potential good that can be drawn from them, and use that insight to change the future before the future arrives.

Key Takeaways

- Blame, compliance and manipulation are system "avatars"; each has a positive twin — learning, ownership and honest measurement.
- Risk homeostasis reminds us that controls which make people feel safer may simply move the risk.
- Legislation sets the floor; risk owners build the house.
- Good response plans are prepared before the event, with defined trigger points.
- Every trusted control has known failure modes — design for them in advance.
- Near-misses are the industry's most valuable free lessons.

In the next chapter, we examine the ten shifts that are changing the very nature of risk — and what each demands of a modern mining organisation.

Chapter 3: The Changing Nature of Risk — Ten Shifts for the Modern Mine

> *Did you know that, according to research published by the Institute of Risk Management in its Enterprise Risk magazine (Winter 2021), only about half of businesses report actively fostering a risk culture — while close to eighty per cent of investors say they are prepared to back companies that improve their Environmental, Social and Governance (ESG) performance? The gap between those two numbers is one of the greatest opportunities in modern mining.*

The commercial and operational environment of every mining company has changed, and it will keep changing as new technology arrives. The Fourth Industrial Revolution has created an exciting and demanding arena for risk managers. Those who have assessed the velocity of emerging risk are already upskilling their systems and preparing for the new. Governments and stakeholders worldwide are looking to reduce costs and secure sustainable revenue, and investors increasingly focus on ESG performance and political stability. A strong risk culture has become a source of strategic relevance.

New practices always begin as embryos. They have yet to prove themselves, and established habits are comfortable. The organisations that thrive are those that embrace credible sources of advice — including decades of research by bodies such as the Institute of Risk Management — and build a culture that is curious about the new rather than suspicious of it.

If risk is changing, we must examine the factors that compose it. Ten shifts define the journey.

Shift 1: From Second-Hand Experience to First-Hand Insight

Experience is valuable, but risk has moved on from yesterday — in its controls and in the response it demands. Copying and pasting information from one risk assessment to another, or from one data source to another, produces documents that look complete but describe a world that no longer exists.

The shift is to reveal the fundamental nature of each risk and its current exposure. Stakeholders should understand the consequences of relying on time-past information, and should expect redundancy to be built into the system, so that they can make timely decisions, reduce consequences and avoid future surprises.

Shift 2: From Fragmented Appraisals to Integrated Foresight

A narrow focus on one emerging risk, unconnected to the rest of the system, closes windows of opportunity. The holistic approach set out in ISO 31000 and in the implementation philosophy of the Institute of Risk Management offers a better way. After a major disaster, it is natural for all

attention to swing to the preceding event. The integrated organisation keeps one eye on the next event as well.

Four practical areas show what integration means:

- **Honest information for resource allocation.** When priorities become personal rather than organisational, resources drift away from the actual risk. Transparent, shared risk information keeps the focus where the exposure is.
- **Valuing the risk function through every business cycle.** Every executive will experience periods of downscaling and upscaling. Historically, "non-core" services, including risk management, have been among the first to be reduced, and engineers or production managers have sometimes been moved laterally into risk roles to retain their knowledge. The mature approach recognises risk management as a designated proficiency, plans restructuring with a formal risk assessment, and writes job specifications that protect the company's risk objectives.
- **Integrated accountability with service providers.** Outsourced experts bring valuable capability and usually carry professional indemnity insurance. But liability cannot simply be transferred by contract, and a court — or a community — will look at the whole system. Evaluate each provider's capability, efficiency and efficacy, and how they integrate with your objectives, culture and appetite.
- **Joined-up control of structures.** A water use licence cannot govern the risk of a dam wall failure, whether the dam stores tailings, dirty water or run-off. Structural integrity is a separate discipline with its own trigger mechanisms and its own experts. When the licence and the structural risk assessment are aggregated, the organisation understands its actual risk rather than assuming it is under control.

Resilience does not mean managing one fragment of risk. It means being operationally equipped for anything. Disasters can come from unexpected success — a long, steady decline in lost-time injuries that breeds complacency — just as much as from unanticipated failure: a fatality, a dam wall failure, a total communications outage or large-scale fraud. Resilient systems are dynamic and rebound rapidly.

Shift 3: From Number Games to Meaningful Measures

Since the mid-1990s, the South African mining industry has used semi-quantitative risk evaluation as a central part of its operational risk assessment. The intention was sound: these were enterprise risk management principles, and many believed they would work on the shop floor. Over time, however, a habit took hold of using the matrix to show that, once a few controls had been recorded, risk had become "lower" — without any defined measure of those controls' effectiveness.

The shift is to recognise that the objective is to manage the risk, not to show a lower number. Operational risk assessment remains a management function that demonstrates due care. Its value lies in what it reveals, not in the score it produces.

The same principle applies to financial risk. Hedging against fluctuations in commodity prices is an established practice, and within those same factors lie rich opportunities. Base maturity claims on approved, appropriate criteria. Avoid the subjectivity that creates a perception that all is well, and drive towards mature, distinct platforms for risk anticipation and response.

Where risk and control owners have, in the past, altered risk levels or control effectiveness ratings against professional advice, the event data eventually reveals it. Decision-making that is shaped to fit a desired outcome can drift into "groupthink" — a psychological phenomenon too often followed by disaster. The antidote is a decision-making process in which professional risk advice is recorded, respected and visible.

Shift 4: From Imitated Surveillance to Authentic Monitoring

Every individual brings their own perception of risk to an organisation. Psychological research identifies two broad mindsets that shape a person's predisposition. The first is the extent to which people are spontaneous and challenging, or organised, systematic and compliant. The second is the extent to which people are cautious and anxious, or optimistic, resilient and fearless — sometimes to the point of recklessness. Several validated psychometric tools can measure these predispositions.

Every person also carries a personal balance of moral values that influences daily decisions. Some individuals are predisposed — susceptible — to doing things in a particular way, shaped by biological, psychological and environmental factors. They may imitate a control in their own manner, sincerely believing it is correct, and stand by that decision as events unfold around them.

At senior management level, the balance of risk types in a team strongly influences its collective perception of risk, its willingness to take risk, how information is shared and how decisions are made. Supervisors mirror that culture, and the workforce mirrors the supervisors.

The shift is to understand these predispositions deliberately: to build teams with a balanced risk profile, to recognise when a control is being imitated rather than performed, and to monitor controls through objective data rather than personal impression.

Shift 5: From the Illusion of Control to the Reality of Control

Psychology shows that people tend to overestimate the influence of their own behaviour on events. We believe we control outcomes that are, in fact, uncontrollable, and our brains favour the rosier interpretation.

Consider a practical example. To work safely on a conveyor belt, an artisan must isolate the power at a substation 800 metres up a decline tunnel. Faced with that walk, the mind looks for an easier

way. Why not pull two pull-wire lanyards along the belt instead? They must be reset before the conveyor can restart, after all. The control heuristic suggests that the easier option will often prevail — even though it is not an isolation.

The shift is for risk professionals to design control structures that anticipate these illusions, and to give managers systems that show them the impact of emotional decisions. Put the isolation point where the work is, and the illusion disappears.

Shift 6: From Beehive Clustering to Dynamic Analysis

A two-dimensional risk matrix is a polygon: it places information into cells. It is useful, but it has limits, and it can hide large deviations. Grouping risk into a box reduces the flexibility to change, alter and adapt to risk indicators.

The trigger mechanisms of a risk event require an understanding of the threats we anticipate and the hazards we use to do business. The earlier we can control the timeline of an event, the more effectively we can reduce its impact by calculating risk velocity.

Software matters here. Generic terminology and generic outcomes are convenient, but they can force a risk into a cell of the beehive while the operational environment around it changes. Hazard drift, risk velocity and predisposition — the real stimuli of events — disappear from view. A home-made bow tie built in a spreadsheet may look similar to the output of internationally recognised bow tie software, but it cannot offer the same linking, auditing and verification capability. Quality tools are an investment in the integrity of decisions.

Professional risk managers continuously improve their systems. Part of their proficiency is to resist stagnation and to keep observing new criteria, applying practical analysis through approved methodologies so that systems evolve at the speed of the risks they manage.

Shift 7: From Imitated Responses to Opportunity-Led Identification

The mining industry has developed several default methodologies for risk assessment. Historically, these centre on events known to occur — they populate the industry's history. Very few embrace opportunity identification or resilient event identification.

Event identification may well be the area of the risk management framework with the greatest room for growth. Notice how much insight appears in hindsight: the research papers and conference presentations that follow a significant event. The aim is to bring that quality of thinking into the risk assessment *before* the event.

This begins with a human capital assessment. Before going further, determine whether the people available have the capability, capacity, knowledge and access to information to understand the risk in all its features.

> **Case study — Sinkholes in the coal sector.** An assessment of sinkhole risk at a coal operation needed a professional rock engineer and a geotechnical manager with direct experience

> of the hazard. Initially, a safety manager and an emergency coordinator led the session, with controls pre-assigned from past events — effectively a reproduction of a master's degree research paper. Only when the more experienced professionals joined did the critical identification process evolve into a genuine operational risk assessment. That assessment became the breeding ground for further studies of how sinkholes interact with old mine workings dating from 1922 and with hot coal. The lesson: the risk assessment does not start with past events and controls. It starts with identification — stakeholders, communication, scope and criteria — and with the right people in the room.

Shift 8: From Argument to Synergy

The ability to argue a point about risk does not demonstrate proficiency in it. Proficiency shows in the ability to bring an approach together.

This shift concerns not the people involved in incidents, but the people responsible for implementing the risk management structure. Some practitioners have developed great skill in interrogating every change, debating definitions and resisting alternative methods. Advocacy has its place. But when every discussion becomes winner-takes-all, the system loses the synergy it needs.

There is very good news here. Over the years, managers' ability to perceive risk and set control principles has evolved substantially — in some areas beyond that of the specialist staff who advise them. Safety practitioners who carry legal appointments were always intended to have an advice-giving obligation more than a policing one. The question for the industry is whether we have equipped them with the capacity to give that advice, or whether we still expect them to perform as production resources.

Shift 9: From Combined Functions to Specialised Professions

In many organisations, legal safety, occupational health, occupational hygiene, environment and risk were combined into a single department under one functional head. The intention was efficiency. In practice, each of these professions has different objectives, methods and bodies of knowledge. A medical doctor is not a structural engineer; an environmental specialist will naturally steer focus and budget towards the environment; and a medical practitioner reporting to a safety manager or occupational hygienist is a structure that serves neither profession well.

The shift is to recognise each discipline for what it is and to integrate them through governance rather than through a single reporting line. A risk *specialist* should hold a professional qualification and belong to an approved professional body. Separating the legal obligation function from the risk advisory function allows risk professionals to give advice from a professional platform and to subscribe to the work ethics of their profession, just as engineers do.

Shift 10: From Coincidence to Causal Understanding

Sometimes a risk assessment appears to achieve its objectives purely by coincidence: the circumstances simulated in the assessment happen to match reality, without anyone understanding the causal connection to vulnerability. This creates a perception of success that the system's integrity cannot support.

Examining the causal factors of risk means scrutinising the traits and conditions that drive the actions leading up to an event — and giving that work a place in the time we spend on risk assessment. Events are combinations of changing factors, physical and intangible, that shape the magnitude of possible outcomes. When a risk assessment is designed to justify a decision that has already been taken, it can conceal a disaster in the design.

The function of risk management is to cultivate a culture of understanding risk: to give time to the uncertainty surrounding release mechanisms, to the way magnitude changes as risk velocity builds, and to the true effectiveness of our control framework. A mature control framework is designed to articulate, anticipate and avoid these antecedents.

> **Groundbreaking insight:** The ten shifts share one direction of travel — from describing what happened to understanding what *will* happen. Every one of them can begin tomorrow morning, in the next risk assessment you facilitate.

Questions Every Mining Leader Should Ask

- Are our risk management and resilience functions living disciplines, or have they become philosophies and slogans?
- Is resilience a by-product of good leadership in our organisation, or is good leadership a by-product of resilience?
- Have we moved beyond behaviour as the default explanation for every event?
- Do our historical data help us when things go wrong, or only when things go well?

Historical data and analysis are excellent at describing success. The opportunity for the mining industry is to build business continuity on forward-looking understanding as well. Renaming an approach, redistributing blame or redoing a risk assessment changes little if the same principles remain underneath. Changing those principles will influence education, budgets and resources — and it will be worth it. Mining companies that develop a culture of agility, with roadmaps that combine defensive and offensive strategies, will lead the next era.

Key Takeaways

- Replace copied, second-hand information with first-hand understanding of current exposure.
- Integrate risk across functions, service providers, licences and structures.

- Measure control effectiveness; do not let the matrix score stand in for it.
- Design controls for the way people actually think — including the illusion of control.
- Put the right people in the room before identifying risk.
- Recognise each profession for its distinct contribution, and integrate through governance.

In Part II, we move from principles to the mining arena itself — beginning with the hard-won lessons the industry has learned from tragedy.

PART II: The Mining Arena

Mining has done more than almost any other industry to learn from its losses. Part II honours those lessons and asks how to learn them without the loss. It examines tragedy as a teacher, the role of the boardroom and the law, risk in South Africa's unique context, and the fundamentals of industrial, construction and project risk.

Chapter 4: Lessons Forged in Tragedy

> *Did you know that when the controls in a typical mining risk assessment are analysed, more than ninety-three per cent turn out to be administrative — controls that depend on a person doing, acting and responding exactly as a procedure, standard or rule describes? Every one of those controls is a request made of a human being at the moment of greatest pressure.*

Since risk management practices first entered the mining industry, we have developed differing views about *what* we are managing: context, vulnerability, threat or outcome. Today, photographs, video, incident reconstruction, simulation and advanced scanning allow us to document events in remarkable detail and to analyse risk almost in real time. We have been doing risk assessments for decades. The next step is to use that wealth of information to identify root causes with the same precision.

Experience shows a clear line running from a person's planning, instruction and observation to the point of failure. With more than ninety-three per cent of controls being administrative, we have placed an extraordinary burden on the worker. Shifting even a fraction of those controls towards engineered and automated solutions is one of the most powerful moves available to the industry.

Managers have recognised that the challenge goes deeper than identifying hazards. It means understanding *how* hazards materialise, the release or trigger mechanisms of the initiating event, and how hazards drift out of context or control — changing their characteristics in ways that are difficult for the eye to recognise.

Replacing an Old Narrative

"The mining industry is dangerous, poorly managed, has a weak safety culture, and its managers are careless." This statement is heard after almost every incident, usually from people who have never seen the inside of a mine. It does not reflect the industry the author knows. Mining companies do not set out to injure anyone. Managers respect their employees, employees respect their managers, and when it comes to human life, nobody plays with numbers: humanity raises the "no-go" flag.

The more useful exercise is to understand the dynamics that allow that old narrative to persist — and to replace it with a better one.

- **Professional recognition.** Working in risk management does not, on its own, make a person a specialist. A specialist is defined by designation, experience and registration with a professional body. The more the industry embraces professionally registered risk practitioners, the stronger its voice becomes.

- **Slogans versus mindsets.** Popular slogans catch the eye, and everyone follows. Zero Harm is a good example: few people know that it was intended as a *mindset*, not as a measurable outcome. In the same way, critical controls can become "silver bullets", and a good accident record can be read as proof of a sound management system. One success story should not become the vehicle for the future, and change should not wait for disaster.
- **Culture and maturity.** International research by the Institute of Risk Management has suggested that fewer than ten per cent of companies have high risk maturity and roughly sixty per cent have low maturity. Seen through the public's eyes, such figures appear to explain failure. Seen through a practitioner's eyes, they show how much potential is still to be unlocked.

Imperfect Perfection

Risk management has evolved along an ascending pathway towards maturity, and each disruptive event has fuelled the next step. Like a cyclone, it is driven by pressure and change: legal obligations, company risk appetite and industry objectives. After grievous events, the need to alter course becomes overwhelming. Each catastrophe in the mining industry has revealed opportunities to improve the system and raise risk maturity.

There is a caution, too. Systems can weaken, just as a cyclone weakens over land — not only because of resistance to change, but because of pressure to adopt interim solutions. The goal is to keep the energy of improvement without waiting for the storm.

> **Case study — The ladle explosion.** At a smelter, an experienced crane driver operated an overhead crane that moved ladles holding about eight tonnes of molten metal. He lived with a known medical condition, bipolar disorder. Like every disaster, this one was a sequence of smaller events that accelerated towards the centre.
>
> The crane driver for the next shift called in late, so the operator was asked to extend his shift — for the second time that week, taking him to 16 hours against a normal 12-hour cycle. While the crane waited, a full ladle stood idle and formed a solid crust on top. The procedure was to move the ladle to a designated area where a rock drill operator would follow a risk-based procedure to break the crust before the metal was poured into ingots.
>
> It was shift change. The incoming crew performed housekeeping and used water to wash away debris spilt during the night shift, leaving water on the concrete floor outside the furnace building. When the rock drill opened a hole in the crust, the whole crust fell away and molten metal overflowed the edge of the ladle, which the crane driver had tilted into position. Molten metal meeting water produces a rapid, violent steam explosion. Thirteen employees were seriously injured, an adjacent building caught fire, and a veld fire started more than 200 metres away. None of the injured were directly involved in the task; they were simply passing by, unaware of the exposure in that area. The crane driver was not injured, but became a silent witness to the tragedy.

> The investigation focused on the rock drill operator, and the remedial measures matched those of six previous events in the same activity: redo the risk assessment, revise the procedure, discipline the operator, retrain and enforce the rules. The deeper questions were different. Why was a high-risk activity designed and positioned in an area that was always busy with people moving between buildings? Why was it not enclosed? Why did the shift-extension decision not take account of fatigue and a known medical condition?

The lesson of the ladle explosion is that the investigation system was designed to find behaviour as a cause. Over time, the industry has begun to evaluate the effectiveness of its risk management systems alongside the systems that validate people's hazard competency. The remaining challenge is to put a manager's understanding of the hazard into the worker's hands. Managers see hazards differently. The worker must recognise the moment a hazard moves from a stable to an unstable condition, drifting from normal, and must then execute the task by following a procedure that may be written in complex language. Clear, plain-language procedures are a control in their own right. And the many employees who live with medical conditions deserve to be included, respectfully and confidentially, in our risk management strategies.

When a Risk Assessment Confirms a Decision Already Made

Some risk-based decisions carry hidden consequences. The 12-hour shift cycle is a good example. At many mines, employees travel long distances and form lift clubs to save fuel. An hour on the road, twelve hours at work and an hour back adds up to fourteen hours, before time in the change house is counted. Add a rotation between day and night shift every month or every three months, and the disruption to the body's circadian rhythm is severe — with obvious implications for fatigue and for family life. A well-being programme may then be designed to support the outcome of the original decision.

When a risk assessment is used to endorse a decision that has already been taken, it stops being a risk assessment. The mature approach is to assess *before* deciding, and to include the people who will live with the outcome.

Breaking the Disaster Curve

History shows that accident statistics have declined after each significant mining disaster. Plot the data by decade and the pattern is unmistakable: a spike, a rapid fall to a much lower level, and then stabilisation at a new plateau of acceptance. South Africa's own history is marked by such turning points. The Coalbrook disaster of 1960, in which 435 men died when pillars collapsed, gave rise to the research that produced the Salamon–Munro pillar design approach. The Kinross fire of 1986, which claimed 177 lives, transformed thinking about flammable materials underground and self-rescue. Vaal Reefs in 1995 reshaped shaft safety and legislation.

Each step was a genuine advance. The opportunity now is to evolve our risk management framework through proven strategies and active pursuit of improvement, so that the next step down the curve does not require a disaster to trigger it.

Residual Risk Is a Description, Not a Number

Although we want our risk management systems to be proactive, we have, almost without noticing, turned the risk matrix into a reactive tool. Less mature systems adopt a principle of "proofing": controls, mitigation and treatment are each given a numerical value to show that the risk level has fallen.

The definition of risk was never a number. In the context of risk identification, ISO Guide 73 describes risk in terms of risk sources, events, their causes and their potential consequences. Residual risk should be described the same way. What remains after treatment? *Describe it*. Stakeholders want to know what risk is left, not how much a score has decreased. When an event happens and we ask ourselves, "Where did that come from?", the answer is often that the description of residual risk was replaced by a number.

> **Groundbreaking insight:** The risk matrix is a prioritisation tool for treatment options — nothing more and nothing less. Unless a recognised quantitative study, with simulation, shows that risk has decreased, no score should change on the strength of a person's or a group's perception. Almost every catastrophic event in mining history was recorded on a risk assessment somewhere. You can alter the number, but a conveyance can still leave its guides for as long as we use conveyances.

The use of the matrix as a measure of safety since the mid-1990s is perhaps the industry's most valuable lesson in this area. Tragedy has shaped a clearer view: the matrix tells us where to look first. Control effectiveness tells us whether we are safe.

There is a legal dimension as well. In the author's experience, no court has asked for a residual risk level to establish whether stakeholders took measures that were reasonably practicable. Residual risk ratings are not used to show that a delict did not occur. When an event has happened, the evidence required concerns negligence and causation, not whether a residual risk score was met.

Key Takeaways

- More than ninety-three per cent of typical mining controls are administrative — shifting even some towards engineering is a major opportunity.
- Replace the "dangerous industry" narrative with evidence of professional, measured risk management.
- Investigate the design and positioning of tasks, not only the behaviour of the last person involved.
- Assess before deciding; never use a risk assessment to endorse a decision already made.

- Residual risk is described, not scored; the matrix prioritises, control effectiveness assures.

In the next chapter, we move from the shop floor to the boardroom and the statute book, and then to the wider South African context in which mining risk is managed.

Chapter 5: Governance, Law and the Rainbow Nation

> *Imagine a board meeting in which the chief financial officer and the head of material risk present a single, integrated picture: what could stop the company meeting its objectives, what could take a life, how the two are connected, and which controls protect both. In most mining companies those conversations still happen in different rooms. Bringing them into one room is the boardroom's greatest risk opportunity.*

Risk in the Boardroom

Resilience has traditionally been defined as the ability to return to the status quo after a disturbing event. Many events have shown that the harm to a company's reputation can exceed the direct cost of recovery, although history also shows that the impact on share price is often short-lived. The real test of resilience is the ability to recover during and after intervention.

Resilience belongs inside the risk management process and inside operational demand. Residual risk measures should show how affirmative steps reduce risk velocity before and during an event; when those factors are documented, they become powerful decision tools. The objective is to address root causes during the risk assessment itself. Very few methodologies build this in proactively. Bow tie analysis is one that does: it identifies threats, establishes control regimes and shows the board the system's vulnerability, so that the board can provide resources to strengthen capability.

Boardrooms today rightly focus on systems perceived as critical to preventing events. The opportunity is to balance critical control management with **resilience system analysis**: an equally rigorous look at vulnerability, capability and the resources needed to recover.

As shock absorbers, boardroom risk decisions should target vulnerability, capabilities and resources, integrating risk and resilience into strategy and programme planning. This adds value through an articulated, multidimensional, cross-sectoral and vertically integrated process. That process requires strategies and policies approved by the risk committee, with two primary objectives: to prevent events and to reduce consequences.

Beyond the Lagging Indicator

Prevention strategies adopted through risk assessment are well entrenched in mining. Although companies share a common commitment to Zero Harm, their methodologies differ significantly, as anyone who has served on several boards will recognise. Reducing injury rates is still the dominant measure of success in the boardroom: to show progress, fatality rates must fall.

The industry deserves great credit for its effort to reduce accident rates. Many programmes aim to improve injury rates by reducing the magnitude of events — in other words, consequence management, turning potential fatalities into injuries. Rollover protection, pedestrian detection and warning systems, airbags, cab design and automatic fire suppression all belong here, and all save lives.

There is a subtler lesson in how some controls interact with people. Narrow roads between structures with a clearance notice, speed humps, multi-road "spider" crossings and crash barriers can all lead operators to assume that a safe situation exists. If an event then occurs through a slip, a lapse, absent-mindedness or a deliberate violation, the design of the control can increase the outcome. People adapt to their environment, and they develop a subconscious capacity to challenge controls. Consequence management is valuable, but it works best as the second line, behind controls that keep the energy contained in the first place.

Strategically, risk management is about much more than lagging indicators. Risk in the boardroom concerns all the risks to which the company is exposed — including the financial decisions that determine its viability and liquidity. Historically, safety risk owners have operated in one silo and financial risk owners in another, with little movement towards marrying the two. Integrating them is one of the greatest leaps the industry can make.

The Month's Flavour

Even well-founded programmes can be redirected by the latest incident. A trackless mobile machine runs over a pedestrian, and immediately every resource turns to mobile machines, their brakes and their interaction with people. The response is understandable. But the most effective programmes are built to absorb such events without losing direction. Risk culture should reflect the tone at the top, and priorities should be set by the magnitude of the risk source and its velocity — not by the most recent accident.

Managers and boardroom risk owners can manage risk proactively. They should be empowered to challenge risk, and to deal with events if they occur, within the parameters of good corporate governance.

Risk and the Law

The law has a profound influence on risk management. This section examines where prescription can constrain good risk decisions, and how a risk-based approach can strengthen compliance rather than compete with it.

In physics, *matter* is anything that takes up space and can be weighed. Legal obligations also take up space in risk decisions: the time spent on them can be measured, quantified and weighed. Companies exist to generate income, create wealth and build sustainability; the law sets the rules within which they do so. Much of our health and safety legislation was written in response to blood spilt in the workplace, and it is driven primarily by consequence. Its great strength is that it

protects people. Its natural limitation is that it was not designed around the economics of a business.

A revealing exercise is to count how often the phrase "the employer must" appears in mining health and safety regulations. The author stopped counting at more than one hundred and forty — still rising with each amendment — compared with only six duties placed on the employee. These obligations are essential, but they also shape behaviour. After an event, they become the yardstick against which managers are judged by regulators, stakeholders, unions, employee representatives and, increasingly, public opinion. Under that pressure it is natural for a *compliance* culture to grow alongside — and sometimes in place of — a *risk* culture.

> **Case study — The "approved" stamp.** In the 1990s, when the first mandatory codes of practice were issued, mining houses complied and submitted their documents to the relevant department for approval. The department stamped the front page "approved". After the first fatality that ended up in court, the manager pointed out that he had complied with the approved practice. The reasoning was sound: the guideline was prescriptive about content, and it had been scrutinised and approved by the regulator. The court turned its attention to the department, and from that day forward the department no longer approved codes of practice, exemptions or assessments.

The lesson is constructive. Ownership of risk cannot be transferred to the regulator — and the regulator wisely stepped back from trying to own it. Mandatory codes of practice are most valuable when they guide a risk-based approach rather than prescribe its outcome. A risk-based approach allows the manager to determine what is essential for the specific operation, rather than relying on what an external drafting team considered important.

In the same period, management made a mammoth effort to find alternative approaches that would redirect the industry towards safer operation. Many of those approaches are now captured in best practice guidelines and journals. The law and the risk-based approach are at their strongest when they work together: the law setting clear duties, the owner deciding how best to meet them.

Risk in the Rainbow Nation

Historical philosophies have shaped the modern strategic role of risk professionals in South Africa, producing a risk control framework with clear principles in the mining industry. Over many years, some systems have evolved and others have regressed. This section explores the forces that shape risk control in South African mining — and the remarkable resilience the industry has shown.

South Africa is one of the most progressive economies in Africa, and its mining industry has been regulated for safety since 1911, when the first mining safety regulations were promulgated under the Mines and Works Act. The broader governance environment has, at times, been challenging. The country's placement on the Financial Action Task Force (FATF) "grey list" in February 2023 was a clear signal of governance weaknesses. Equally clear was the determined national effort that followed: South Africa was removed from the grey list in October 2025. That turnaround

demonstrates a principle every risk manager knows — when risk is named, owned and systematically addressed, it can be brought under control quickly.

Engineers and Financiers: Two Kinds of Excellence

The financial sector has registered professionals who uphold risk governance and compliance. The mining sector produces professional engineers and operational managers who often have no formal risk management education — yet those engineers and production managers have built physical risk control systems that are, in many respects, more rigorous than any financial risk control strategy in the industry.

Part of the reason is visibility. An injury cannot be hidden; it requires medical intervention and must be reported to the authorities. Misuse of funds or petty theft, by contrast, can happen in the background and may only be reported when someone notices or a system flags it. Silent witnesses can become partners in non-conformance until "it's the way we do business". Bringing the same visibility and discipline to financial and governance risk that the industry already applies to physical risk is a significant opportunity.

Measuring What the Annual Report Reveals

Risk culture is the values, beliefs, knowledge and understanding of risk shared by a group with a common purpose. South Africa would benefit from a formal approach to measuring governance and risk compliance culture across sectors. As the author argued in *RM Professional* (Winter 2015), there is a tendency to appoint a financially qualified person at board level to be responsible for risk. That approach can limit the board's knowledge of material risk control. Saving lives cannot be achieved by protecting monetary value alone.

Look at a mining company's annual report and ask what measurable, functional governance and risk information it contains. In the author's analysis, on average about eight per cent of critical risk and performance indicators relate to material risk, and fewer than six per cent of principal risks are material risks. The evidence suggests that *protection* currently outweighs *prevention*. The distinction matters:

- **Protection** manages risk to secure the company's financial stability. A project may be terminated, mothballed or downgraded to protect capital expenditure.
- **Prevention** deploys strategies to explore risk opportunities and avoid loss — for example, investing in collision avoidance systems, with the maintenance policies to sustain them, because they save lives.

A balanced annual report would show both.

From Protective Recommendations to Industry-Wide Learning

Much of the industry's direction has been shaped by recommendations arising from serious accidents. Those recommendations tend first to protect the risk owner. Implementation often happens only at the site where the event occurred; the lesson can then disappear from the wider industry until, years later, it resurfaces as legislation. Accident investigations can also drift away from the basic principles of risk management.

The opportunity is enormous: share learning across the industry immediately, through structured, open learning platforms; adopt adaptive engineering; and give risk *opportunities* the same attention as risk *threats*. Such a shift would be a substantial change of mindset for the mining and metals industry, requiring well-designed policies and an expansion of existing systems — beyond identifying risk sources, assessing them and assigning a matrix value — to include tracking, shadowing and simulating risk opportunities and acknowledging their outcomes.

Integrating the Tiers

Systems that support risk response activities can be simple, and many have been adopted with varying success. A tiered approach is common: lower-tier risks are placed in silos and handled by subject matter specialists, while the upper tier operates independently. Over time, this detachment can produce a protective culture and a disjointed process. The remedy is integration — clear information flows between tiers, shared language and a single view of risk from the shop floor to the board. (The related question of combining distinct professional functions into a single department was discussed under Shift 9 in Chapter 3.)

Building the Framework at Board Level

Strategies and risk architecture are developed at board level, where risk appetite is matured and promoted. At a minimum, the upper-level risk management framework should include:

- **Risk structures** that define roles and responsibilities for internal and external control.
- **A risk management strategy** that gives direction for philosophies and policies.
- **A risk management system** with detailed information, audit protocols, rules and procedures, approved by the board.

In parallel, a lower-level strategic management process is articulated on the basis of international best practice and guiding principles, with the business risk owner as an active participant. When the strategy rolls down into operational activities, it can be seen to have arrived on the shop floor when the material risk owner recognises it in the verification of controls. Chapter 15 develops this framework in detail.

A Story of Innovation

Over several decades, the South African mining industry has seen the full range of practice. The author recalls a risk manager asking why it is necessary to identify hazards or establish the magnitude of a risk source, rather than simply listing events. He also recalls critical control "best practice" designed for a highly sophisticated plant, complete with QR codes, being applied in settings for which it was never intended. These moments are part of every maturing discipline.

One conclusion stands above them all. Future research will show that the South African mining industry was innovative and made genuinely ground-breaking survival decisions during some of its most challenging times — through political, fiscal and legislative change. That is the foundation on which the next generation of risk management will be built.

Key Takeaways

- Balance critical control management with resilience system analysis at board level.
- Consequence controls save lives, but work best behind controls that contain energy at source.
- Integrate financial and material risk into one boardroom conversation.
- Ownership of risk stays with the owner; the law sets the duty, the owner designs the response.
- South Africa's FATF grey-list exit shows what named, owned and systematic risk response can achieve.
- Report prevention as clearly as protection.

In the next chapter, we look at the ground beneath our feet — literally — and at the fundamentals of industrial, construction and project risk management.

Chapter 6: Building on Solid Ground — Industrial, Construction and Project Risk

> *Imagine sinking a shaft a few hundred metres away from where the orebody model said it should be. Nothing collapses. Nobody is hurt. But every tonne of ore for the next forty years travels further than it needed to, and the cost of that one early decision is paid every single day of the mine's life. Some of the most consequential risks never appear in an incident report.*

Geotechnical Science: Risk Management's Closest Ally

Risk management and geotechnical science are natural partners. Both deal with uncertainty, and both seek to minimise deviation from a designed set of principles. Risk management tends to be more qualitative, so it benefits greatly from disciplines that can quantify. Civil engineering and geotechnical science are among the most valuable — and most underused — of these allies in the project and operational phases of industrial and construction work.

Compare two industries. Highways built over decades are periodically upgraded: widened for more traffic, resurfaced for repair or greater load-bearing capacity. Slope failures beside these upgrades, roads that flood and water running dangerously across the surface during rain are familiar sights. Resurfacing brings two decisions that deserve a risk lens. A new surface laid on top of an already disintegrating one inherits its weaknesses. And each new layer raises the road surface, reducing the clearance under bridges. A tanker carrying liquefied petroleum gas that strikes an overhead bridge can cause a massive explosion with multiple fatalities. The road was safe on the day it was designed; the cumulative change made it unsafe.

Now consider a mining company tunnelling underground for a rail haulage. Here the company is acutely sensitive to how ground, soil and water react to disturbance — whether from geotechnical conditions, hydrology or other external loads and stresses — and it is legally obliged to employ geoscience and geotechnical analysis.

What Construction Can Learn from Mining

In parts of the building and construction industry, geotechnical investigation receives far less attention. Technical data may be thin, and the financial pressure in planning new housing developments, multi-storey buildings, drainage and stormwater systems can be intense. Civil engineering and geoscience often meet in different boardrooms. Ask for the evidence and you may find one or two borehole samples representing a large project. Two samples cannot support a sound risk decision.

The consequences are visible in smaller towns. Buildings have been constructed on top of old waste sites, only to show severe structural damage within a few years, leading to court cases and

demolition. In coastal towns, houses have been built on unstable material; combined with poor stormwater management, the result is ground movement and structural damage. In such cases, investigations sometimes reveal that a professional decision was taken without any geotechnical information at all.

The mining industry, by contrast, has a firm grip on subsurface conditions. Sampling, mapping, analysis and interpretation are embedded in mine design and planning. Geotechnical modelling and rock engineering define the expected behaviour of rock and rock masses under stress, and numerical modelling is used to assess the risk that invalid data could lead to misunderstood design problems.

> **Groundbreaking insight:** Mining's geotechnical discipline is one of its great exportable strengths. The construction and municipal sectors could avoid many structural failures by adopting the same standard of subsurface evidence that every competent mine planner takes for granted.

The construction and building industry is ready to embrace risk management principles more fully. Its legislative framework still leans towards a historical safety approach of inspections and observations — a lagging, reactive model. Larger construction companies have adopted risk-based approaches and appointed experienced risk professionals to drive change and support their risk appetite. Smaller companies, though fully law-abiding, often lack the expertise and capacity to implement the methodologies. For them, "risk" can still mean "safety". Industry bodies, universities and larger contractors can make a significant difference by sharing knowledge, language and tools with their smaller partners.

Project Risk Management

Project risk management is a subject in its own right, but its fundamentals deserve attention here. The first principle is integration: project risk management must connect to the company-wide risk structure. Separate systems drive different decision models and drift away from company objectives.

Project managers lead the process and appoint experts who understand project requirements. One of those experts should be a risk manager who oversees risk for the project's entire life cycle — a competent facilitator who knows the relevant assessment and decision techniques. Working with the project manager, the risk manager assesses project risk at each phase and stage and controls access through each "gate".

Setting the project's risk appetite is not the project manager's role. Based on the risk assessment at each phase, the risk manager develops a risk appetite statement that reflects risk-based decisions and informs the capital funding structure. Practically, that means knowing **when to pursue risk and when to pull the plug**: when to mothball, terminate or discard a project, or part of it, within defined time frames.

Lessons from the Information Flow

Many project lessons concern the flow of information, and its effect on project efficiency and key risk indicators can be enormous. The same issue lies behind a significant number of industry disasters. Real examples include:

- A tailings dam whose construction ran late, so that tailings were pumped into a river for several days.
- Process plants mothballed because capital was reallocated mid-project.
- An application for a water use licence that delayed the start of a project by six months.
- The discovery of eight-hundred-year-old graves mid-project, which caused a four-month delay while they were respectfully relocated.

Each was avoidable with earlier, better-connected information.

Three Kinds of Project Risk

All project risks can be classified as **preventable**, **strategic** (calculated risks taken in pursuit of reward) or **external** to the project (Kaplan and Mikes, 2012). Changes in project scope make preventable and strategic risks harder to manage. Seven challenges deserve particular attention.

Uncertainty. Uncertainty differs between greenfield, brownfield and bluefield projects. Greenfield projects have little interaction with existing work; the others interact directly with existing operations, with greater potential impact. Uncertainty in a dynamic system distorts risk velocity calculations, deviation analysis and sensitivity analysis. At the same time, burying the analysis in complexity can delay, derail or halt a project. Unused buildings, half-built bridges on new national roads and steel structures rising into thin air are monuments to unresolved uncertainty. In mining, the consequences may be less visible — like the shaft sunk a few hundred metres from its optimal position — but they last for the life of the mine.

Scope and context. Early involvement improves risk identification. When stakeholder assessments that clarify the work and its social and environmental implications are skipped, consultation suffers, transparency is questioned and a mature risk culture can be lost. A formal stakeholder engagement approach from day one prevents all three.

Optimism bias. Optimism motivates individuals, but in a group it can blind a project team to the inherent risk of the project. People in authority may soften the reality of the circumstances being assessed. Resistance to change, the opposite tendency, has the same effect and is often rooted in limited experience of the project domain. Risk managers and facilitators should be able to recognise and record these influences, because they will shape the response plans.

Reliability of information. Adequate, reliable data are the most important risk decision factors for a project and for the operation that follows it. Exploration that produces too little information for sound modelling has contributed to closed-off decline shafts in the platinum sector, difficulties in public rail tunnelling, and open pits that had to become underground mines.

The ideal mining project generates its own capital by mining an outcrop while the process plant is being built. If exploration reveals major geological disturbances requiring a decline to more stable ground, that advantage is lost — and the loss is greater if an optimistic project or risk manager underestimated it.

Communication. Risk assurance depends on communicating the right information. Transparent, timely communication with stakeholders has an equal effect on risk response. Strong communication channels build understanding of the risk, its controls, the cost of mitigation and the sequential impact. Mixed signals erode trust.

Budgeting. From a risk perspective, project budgeting is not merely an accounting function. It is interlinked with every project activity and time frame, all of which carry a cost, and it is the single most important risk calculation factor in any project. Extending the delivery date can damage the project and stakeholder trust; compressing the schedule by ignoring known risk delays has the same effect. This balancing act is a tangible risk. Project managers should enable the risk manager to apply appropriate techniques to identify it correctly, and the team should put mitigation and response plans in place together. Including the risk manager in these decisions also reveals secondary risks — smaller downstream impacts that combine into something significant. And remember that "harm" includes harm to the company's goals, stakeholder relationships, union trust, employee well-being and the social environment.

Complex dependencies. Multidisciplinary, multilateral, multidimensional projects create complexity. The more interdependent the elements, the greater the risk of subjectivity, and misalignment can propagate outward and upward through the project. The dynamic environment — market trends, technological change, regulatory change and unexpected events such as COVID-19 or the discovery of ancient graves — adds the most uncertainty of all.

What Works

Companies that value risk management, invest in their risk practitioners and strengthen the risk function throughout the project life cycle overcome these challenges. Proven approaches include interactive risk workshops, risk tolerance analysis, scenario analysis, consequence research, strategic intelligence, and expert collaboration using tools such as forensic scanning and real-time monitoring.

Project risk managers need proficiency in three key areas:

- **Context:** an understanding of the risk framework for the project domain (industry, technology, geography), the life cycle from cradle to grave, costing (from exchange rates to end-user demand) and resource allocation (internal and external).
- **Analysis:** diagnostic skill in qualitative and quantitative risk analysis, including software tools, simulation and modelling to improve the accuracy of evaluations.

- **Communication:** the ability to explain complex concepts, change (adaptive engineering) and consequences (problem-solving) to cross-functional teams and stakeholders, creating an environment in which synergistic solutions emerge.

The Human Side of Projects

Projects are run by people, and well-documented decision biases appear in every project team. Identifying them early and planning for them is far more effective than discovering them in an incident investigation — especially in project environments where punitive cultures can drive high labour turnover.

- **Overconfidence:** managers who trust their abilities so strongly that they underestimate significant project risks.
- **Confirmation bias:** seeking information that confirms existing opinions, and discounting or softening risk information that does not.
- **Deference to dominance:** teams that become accustomed to decisions made by dominant stakeholders, even when those decisions do not reflect sound risk criteria. When arrogance, silence and fear of raising the topic combine, a project becomes prone to groupthink.
- **Loss aversion:** managers who are more sensitive to potential losses than to the benefits of opportunities, and who therefore avoid calculated risks.
- **Anchoring:** decisions fixed on the first piece of information received, distorting professional judgement.
- **Sunk-cost commitment:** large capital investments that make it hard to "abandon ship while it is still afloat", even as losses grow.
- **Short-termism:** prioritising short-term objectives over long-term risk, creating the conditions for legacy risks.

The remedy is a structured approach that harmonises attitudes, built on three segments: a **project risk management structure** (governance, stakeholder engagement, assurance and reporting), a **project risk management strategy** (methodologies aligned with the company's risk appetite, with policies translated into objectives and result indicators), and **project risk management systems** (the techniques, protocols, measurements and verification through which implementation happens). The same three-part architecture underpins the enterprise risk management framework described in Chapter 15.

Key Takeaways

- Geotechnical evidence is a risk control; mining's subsurface discipline is a model for construction.
- Cumulative small changes (like road resurfacing) can erode a design's safety margin.

- Integrate project risk with the enterprise risk structure, and give the risk manager a role at every gate.
- Know in advance when to pursue risk and when to pull the plug.
- Classify risks as preventable, strategic or external, and manage each differently.
- Plan for decision biases at the start of the project, not after the incident.

In Part III, we turn to the most complex element in every risk system — the human being — beginning with the enigma of human conduct.

PART III: The Human Factor

Every risk system is ultimately operated, maintained, bypassed or rescued by people. Part III explores how people perceive danger, why errors happen, what habits do to our judgement, and why investigations must evolve from finding fault to understanding cause.

Chapter 7: The Enigma of Human Conduct

> *"We cannot change the human condition, but we can change the conditions under which humans work." — James Reason, "Human error: models and management", BMJ (2000)*

Honestly and without prejudice: can we fairly hold a person responsible for an accident without first understanding the person? Humans have shaped and changed their environment since the beginning, and they have never stopped. By that very involvement, we accept all the risks and opportunities created when people interact with the world around them. The same engagement has given us the ability to perceive, understand and judge human conduct — and it has shown us how easily our own perceptions can mislead us.

The aim of this chapter is not to repeat the work of the great philosophers and psychologists, but to build practical links between what is known about human conduct and how risk is managed in mining.

Mind and Matter

Start at the beginning. What is the difference between mind and matter, and how does one influence the other? Alfred Binet, in *The Mind and the Brain*, drew the distinction clearly: mind is the domain of psychology, while matter is the realm of physics. The interaction between the brain (the knowable mind) and the surrounding environment (foreseeable matter) determines how a person acts or reacts — within the limits of what that person knows.

In practical terms, every workplace hazard has two features: what the mind must do once it recognises the danger, and what the surrounding matter does to the mind through our senses. Together, they determine our ability to manage the hazard.

Beyond Fight or Flight

Many approaches to managing behaviour and human error have focused on the "fight or flight" response, in which the amygdala, a processor of emotions, triggers hormones that prepare the body to confront or escape danger. These ideas, increasingly expressed in the language of neuroscience, have value. But they tell only half the story.

The other half is the state in which a person over-reacts, or responds with full capability, to a threat that will not actually cause harm. This, too, reflects the ability to develop a rational response to danger. In risk assessments it can appear as a participant who is strongly risk-averse, who tends to underestimate the effectiveness of controls or overestimate consequences, and for whom emotion plays a central role in debate. Such participants are valuable: they test assumptions. The facilitator's task is to channel that energy into evidence.

How We Think Before We Act

Instinct and emotion motivate behaviour. Human conduct is the mind's instinctive reaction as it interacts with matter — the perceived danger around it. To understand where error arises, we must take the thinking process apart.

- **Actions** are the process of doing something to achieve an objective of the mind. They arise from inborn capacity and from responses to sensory stimuli, which we call instinct. Emotion allows us to reason from what we know and to distinguish right from wrong during the action. Confidence in certain senses and a deep survival instinct are central features.
- **Conduct** is how a person behaves within the restrictions of surroundings, situations and operating conditions.
- **Habits** are the repetitive ways in which we do things in the same manner. They develop from a subconscious effort to balance good practice with instinctive ambition, and the dominant disposition encourages a particular way of performing an activity. The work environment, operating systems and procedures are a breeding ground for habits. Consider a haul truck operator driving the same road between the loading point and the tip for eight hours. This is not about fatigue, which is managed through breaks and monitoring systems. It is about the operator's objective of maintaining a certain standard at a certain pace — the dominant trait that yields a habit.

A problematic event with obscure causes often becomes clear the moment we ask about habitual contributing factors. When a witness describes a task as "the way we always do it", the habit is visible. The usual response is that the task relies on skill and experience, the product of excellent training. But there is no clear boundary between skill and habit. Skill is possible only where habit exists, and habit exists wherever skill has been achieved. Training moulds skill, skill patterns habit, and habit makes actions easy. Work cannot flow without habitual activity — and this automation of brain function deserves a place in every accident investigation.

- **Fault** implies that a disposition has fallen short of a desired standard. A person's mental state — perception of risk, experience, beliefs and memory of past events — strongly influences whether a fault is a breach of a rule or the product of a habit. Fault can also arise from the brain's autopilot mode, from subconscious engagement with another activity, from psychological "noise", or from reckless omission.

During accident investigations, we tend to focus on *legal* fault — negligence, intent or consent — rather than the *mental* state described above. Investigation processes in the industry are well aligned with establishing blame, and they tend to pursue fault in one direction. Very few track the habitual or restraining forces involved. Adding those dimensions is one of the clearest opportunities in this book.

Understanding Error

How does all this relate to human error, the ancestor of the "unsafe act"? Before any law, what determined whether something was a mistake? If I am cold, I seek warmth; if I am hungry, I seek food. In the same way, I seek a safe working day and a scratch-free return home. The reasoning behind those actions varies by every possible degree.

Human error remains one of the most thought-provoking characteristics of humanity. Several causative features help to bring it into focus. Consider them through a single scenario: an underground tunnel, an approaching vehicle and pedestrians.

- **Cognition shapes the decision.** When encountering a dangerous condition, decisions are influenced by emotion, instinct, habit and perception. One pedestrian stops and looks for a safe place to stand; another keeps walking, judging that there is enough space between the vehicle and the sidewall.
- **Situational awareness and individual differences** make responses unpredictable, and hard to reconstruct afterwards. They vary with how a person manages stress, fatigue and operating conditions. The pedestrian who kept walking was under production pressure, with several activities to finish before leaving the mine. At the time, the decision made good sense.
- **Past consequences shape risk perception.** The pedestrian knows that people have been run over or crushed against the sidewall, yet that knowledge does not change the decision. He perceives that he is in control and that the outcome can be predicted.
- **Errors are often the product of several decisions by several people.** The vehicle operator loaded long material that protruded into the space between the vehicle and the sidewall. Another person chose to walk behind the vehicle because it was travelling in the same direction. The oncoming pedestrian, dazzled by the headlights, looked down at the footwall to avoid tripping in front of the vehicle. The vehicle was fitted with a proximity detection system designed to slow and stop it near a person. Over time, drivers and pedestrians came to rely on that system, and were reluctant to believe that it could be in a failure mode. The safer, modern technology had quietly replaced the older decision to stop and wait for the vehicle to pass.
- **Investigations must look into the decision itself.** After an event, the investigator follows a structured questioning process that is easy to predict in hindsight. Investigators trained through short, one-dimensional programmes rarely have the tools to explore the psychological processes behind a decision. Incorporating the understanding of how decisions are constructed into causal investigation methodologies is essential. A questioning line driven by yes-or-no answers cannot reach it.
- **Lessons learned should capture good decisions too.** Continuous improvement should ask which features of the human mind contributed to decisions that ended in unwanted events — and, equally, which features produced the *wanted* outcomes and good business. The next maturity step is to learn from success as rigorously as from failure.

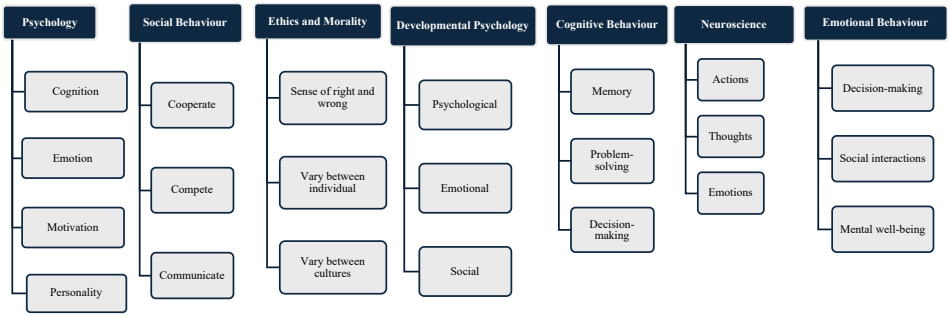

Figure 1: The complexity of human behaviour

Figure 1 illustrates the breadth of disciplines that shape human behaviour — and why simple programmes rarely change it. Decision-making is a skill that improves with practice. Understanding the human mind's capabilities and limitations, and how it is influenced, offers a pathway to understanding behaviour. With proper education, a mature risk decision-making model creates a mature awareness of where errors are likely.

The industry can move from task risk assessment towards **potential error analysis** — identifying, in advance, the points in a task where decisions are most likely to go wrong. This gives line supervisors the opportunity to close gaps in decision-making before an event. Where this cannot be achieved, automation and technology lead the way towards Zero Harm.

A further opportunity lies beyond near-miss reporting. A mature risk management system encourages the reporting of human error itself, and analyses the underlying causes. This culture change, at supervisor and employee level, requires a detailed programme built on good communication and trust.

What Shapes Risk Perception

Every person perceives risk differently: how hazards are understood, how energy flows, how conditions influence a risk source. Psychological, cognitive and environmental factors combine to shape that perception, and it changes with age and time in the workplace. Some individuals need more support to develop a strong sense of risk. Several factors, often missed in investigation and causal analysis, explain why.

- **Settled habits** form when actions are performed unconsciously and often compulsively, reinforced by repetition, long periods on the same task or physiological factors.
- **Persistent behaviour patterns** — rebelliousness, insubordination, disregard for the rights of others, bullying and harassment — are disruptive habits that undermine the safety of everyone around them and deserve firm, early attention.
- **Health and neurodevelopmental conditions** can affect how a person identifies intangible risk, communicates about risk and manages it in practice. Conditions such as

post-traumatic stress, depression and bipolar disorder can also influence habits. Handled with respect, confidentiality and appropriate medical support, these factors can be accommodated in safe work design, as the ladle explosion in Chapter 4 showed.

- **Historical exposure** — beliefs, norms, cultural background, peer influence and social pressure — cultivates different risk tolerances. Lack of exposure to a particular risk source limits a person's ability to anticipate it.
- **Cognitive biases** lead people to over- or underestimate risk, or to stay safely in the middle and avoid any decision that might affect their ego. Listen for phrases such as "not on my shift", "impossible" and "it has never happened". People who believe their risk judgement is always correct, because outcomes have so far been positive, become overconfident — and overconfidence is where disaster is born.
- **Intellectual and technical understanding** of cause and effect varies. People who find technical information difficult may develop coping habits: downplaying risk, using generic statements that suggest understanding, or opening an early debate that diverts the assessment to another department. Good facilitation makes technical content accessible to everyone in the room.

Risk perception and awareness can grow or fade, and they are subject to all the complexity of human behaviour. When people realise they have made an error, their response depends on disposition, severity and circumstance: denial, increased stress, guilt after a fatality, apology, seeking help or shifting blame. Each can be observed in the aftermath — and each deserves a compassionate, structured response.

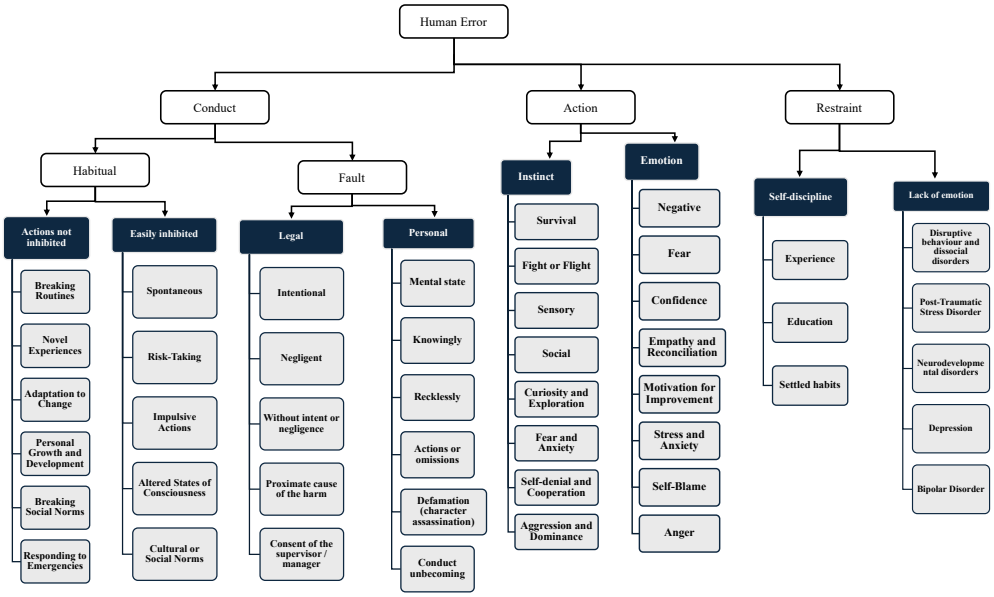

Figure 2: Human error in relation to human behaviour

Figure 2 summarises how human error relates to the behavioural factors shown in Figure 1.

The Psychology of "Casino Safety"

Psychology is often avoided in risk management, sometimes simply because the subject seems daunting. Yet uncertainty has always exposed psychological patterns in the workplace, and none is more instructive than what the author calls **casino safety**.

The gambler believes his judgement is correct and keeps placing bets. Wisdom and folly are as different as day and night, yet, like day and night, they meet at dusk and dawn — brief moments in which intelligent people make reckless decisions. Carelessness, the root of recklessness, connects directly to likelihood. Adjusting the figures in a risk matrix on the strength of perception is no different from gambling. Repeated, it becomes a habit, and in a boardroom it becomes groupthink.

When a subject becomes too technical during a risk assessment, people fall back on their existing knowledge and habits, and may adopt a generic attitude. The event then becomes subject to the roll of the dice.

Consider the pre-conditions for the failure of a hoist rope. The engineer reasons that the rope has a safety factor and remains conditionally safe until that limit is reached. The risk assessor starts at zero deterioration and treats any deviation as an increase in the likelihood of failure. A participant without knowledge of hoist ropes may argue that the rope will fail regardless of what is done. Like the gambler, that participant is interested in the outcome, not the likelihood. A good facilitator brings all three perspectives into one evidence-based view.

Horizon scanning challenges every assessor, especially in finding the path between recklessness and sound risk-taking. The gambling mindset can stay quiet during hazard identification and then debate risk levels at length, returning to the starting point so that the decision feels like its own. Objectivity is what the wise person brings to the table; it is the facilitator's most important gift to the team.

Wisdom and the Average Person

A legal obligation exists to ensure that people are not exposed to thermal stress that could cause heat- or cold-related illness. The wise approach is to include as many factors as possible in assessing each person's risk profile. The less mature approach is to generalise: to adopt a temperature limit that would harm the average, "reasonable" person — the lowest common denominator — as a benchmark. Anyone more vulnerable than that benchmark is then at higher risk.

A worker collapses while re-entering an underground workplace. The investigation focuses on the obligation: how was the person exposed? Whether the risk had been rated low or high no longer matters. If the investigation is designed only to establish who was in breach, that is what it will find. The industry, with so many people exposed and such serious efforts to prevent major events, deserves investigations that ask what made the exposure possible — and individual risk profiles that make it less likely in the first place.

Imitation, Passion and Emotion

Other psychological traits influencing behaviour and habit include imitation, passion, affection and emotion. Imitation is especially common where leadership style or risk culture allows it to spread. A front-line supervisor may share the habits of the section manager, who in turn mirrors the general manager. When leaders model excellent risk behaviour, imitation becomes a powerful force for good — which is why visible, authentic leadership matters so much.

In criminal psychology, passion and affection are features of the individual mind that affect physical emotion. Desire and care both shape emotion. Some individuals show a healthy fear of risk; others do not, and tend to challenge risk with the same mindset. They are not criminals because they have little fear. And experience shows that punishment alone does not change behaviour: the successful prosecution of white-collar crime, even with sentences of more than thirty years, has not stopped such crime from increasing. In the same way, risk-taking in the workplace has not disappeared because of investigations, blame or attempts to change behaviour; it has simply taken new forms.

Think of the road. How many crashes are caused by emotion that never becomes road rage? The actions and inactions of other drivers strongly influence our desire to reach our destination safely. In the workplace, colleagues and surroundings provoke emotions throughout the shift, and prolonged exposure turns those emotions into habitual responses.

When the Senses Mislead

The mind uses the senses to detect its surroundings. Hearing, smell, touch, sight and taste detect physical risk; for intangible risk we rely on cognition, perception and what we call gut feeling. These connect to our instinctive responses of flight, fear or aggression. Everyone has different sensitivities, which affect their judgement and behaviour. In the workplace, the most common sensory errors involve **distance**, **illusion** and **reflection**.

- **Distance.** It is not the professional golfer who is best at judging distance; it is the caddie. In the workplace, distance from a risk source can be the difference between life and death. A community may perceive the outflow zone of a tailings dam as smaller than the operator does, and build homes closer to it.
- **Illusion.** Underground, with only a cap lamp for illumination, the beam can create tunnel vision and optical illusions. What the person perceived remains in memory and is recalled during investigations, potentially misleading our understanding of pre-conditions, sequence and cause. The beam also limits the eye's ability to judge size, which affects the accuracy of pre-task risk assessments and witness statements. When the scene is reconstructed in virtual simulation, we are often astonished by how different reality was from perception.
- **Reflection.** Glare and reflected light, as in the tunnel scenario above, can redirect attention at the critical moment.

Risk assessors and accident investigators should be aware of the magnitude of these effects. Psychologically, we tend to accept illusions of distance and size without question. Making them visible is a simple, powerful improvement.

What Works

How, then, do we work with the enigma of behaviour? There is no quick fix: no single book, short classroom programme or behaviour-based system will solve it. The industry has invested significantly in behavioural programmes over the past two decades and has learned a great deal from them. The greatest returns now lie in three areas:

- **Cultivate powerful habits.** Design tasks, training and environments so that the safe way becomes the automatic way.
- **Provide clear, industry-specific standards and rules** written in language the workforce understands.
- **Strengthen risk assessment capability** to identify the inciting causes that lie within the human capital and intellectual resources of the operation and the industry.

Health and safety practitioners, risk managers and supervisors do not need to be psychologists. But they must understand the risks that arise from human actions and habits when exercising their duty of care. Where there are people, there will be risk events. Conference podiums, journals, political platforms and court findings will not on their own bring events to an end. What will make the difference is a risk management system designed for human nature as it really is — one that respects the extraordinary power of the human brain and builds controls that work with it.

Key Takeaways

- Skill and habit are inseparable; investigate habits as contributing factors.
- Look for mental fault, not only legal fault.
- Errors usually involve several decisions by several people — and trusted technology can quietly change those decisions.
- Move from task risk assessment towards potential error analysis.
- Recognise the “casino safety” mindset; objectivity is the facilitator’s greatest gift.
- Senses mislead: distance, illusion and reflection belong in every assessment and investigation.
- Build powerful habits, clear rules and stronger human-factor risk assessment.

In the next chapter, we apply these insights to investigation — and make the case for a forensic standard that moves the industry from blame to understanding.

Chapter 8: Forensic Investigation — From Blame to Understanding

> *Did you know that a typical serious-incident investigation in mining may take five days or more, while the average time a risk assessment team spends on a single hazard is fifteen to twenty minutes — including the recommendations and the capturing of information? The risk sources, the trigger events, the causes and the outcomes are the same in both. Imagine what would change if we reversed the ratio.*

Understanding the purpose of investigation in mining has been as challenging as understanding the risk assessment process. Confusion and the instinct to fix blame have become deeply entrenched, and over time even experienced professionals have come to accept investigation outcomes that a forensic scientist or a court would question.

Where Our Methods Came From

Most of the investigation methods used in mining were developed by engineers and safety specialists, not by people trained in law or forensic science. Incident causal analysis methods built on James Reason's organisational error model, such as ICAM, were developed in Australia from aviation experience and later adapted to mining. They brought valuable structure to investigation and gave the industry a common language for organisational factors. They also have natural limits. Investigation is a discipline as old as civilisation, stretching from the Twelve Tables through Roman law, developed by professional jurists, to modern forensic science. Few of today's industrial methods draw on that body of knowledge.

Fixed questioning sequences, such as repeatedly asking "why?", are easy to apply but can produce non-factual outcomes. Witnesses are seldom challenged, and they may — entirely sincerely — supply details that are incorrect. Analysis built on such investigations then produces comparisons that create an illusion of correctness. Over time, that illusion can become the accepted norm.

> **Groundbreaking insight:** Investigation methods built on error models taught the industry to look beyond the individual. The next leap is to add the disciplines that courts, scientists and accident boards have used for generations: factual causation, legal causation, evidence validation and independent review.

Raising the Standard of Causation

Applying the "but-for" test of factual causation — would the outcome have occurred *but for* this act or condition? — is a simple, rigorous way to strengthen remediation. At present, the quality of investigations is seldom verified by any person or institution, yet causal analysis is used for predictive monitoring and its outcomes can find their way into legislation. Without accurate

investigation data, and without understanding how the data are modelled, conclusions risk misleading the industry.

Traditional methods draw raw data, which is little different from information. What is required is *intelligence*: work done on the data to add value and significance in the context of the event. The author offers this view knowing it is an opinion — but it is an opinion formed over many investigations, and the opportunity it points to is real.

A successful investigation is not measured by liability or prosecution. It meets the principle of legal causation with ethical quality. Investigators and regulators can strengthen their work by building on sound legal principles: they do not need to establish the position of the injured relative to the person in charge, nor search for a provision that was breached. Their task is to establish cause.

The old school is credited, rightly, with reducing the *frequency* of major events. The new school must reduce their *consequences* as well. Conveyances still fall down shafts, tailings outflows still affect communities, and labour disruption still reaches the mine gate. Forensic-quality investigation is how the next generation of improvements will be found.

The Forensic Standard

Forensic investigation is widely used in other fields, and its literature is rich and well-founded. Its principles are powerful in reducing both events and consequences. Aviation offers a proven model: under Annex 13 of the Chicago Convention, the sole objective of an accident investigation is the prevention of future accidents, not the apportioning of blame or liability, and investigations are carried out by independent bodies. Major industrial inquiries, such as Lord Cullen's public inquiry into the Piper Alpha disaster of 1988, have shown what independent, multidisciplinary investigation can achieve — in that case a complete transformation of offshore safety regulation in the United Kingdom.

Better systems may take more time, because they need accurate and reliable data, a solid foundation and access to comprehensive research. Forensic investigation has several essential elements, and the integrity of the whole system demonstrates due care, diligence and consistency.

- **Trained, experienced professionals.** Gathering correct data requires an understanding of risk sources, release mechanisms and causes. The investigator should understand research methodology, draw sound conclusions, follow lines of defence, avoid assumptions and draft constructive recommendations. Stakeholder trust is best maintained by investigators with exposure to a judicial function.
- **A wide range of disciplines.** Experts and scientists are involved from the early stages. In some countries, investigations of major events are led by universities with knowledge and research capacity in every relevant discipline, which also secures impartiality and standardises the methodology.

- **Only proven facts.** Forensic engineering, forensic psychology and digital forensics are seldom used in current accident investigation. A quality investigation validates information and evidence before classifying it as factual, and may include back-analysis by professionals to ensure reliable results.
- **Transparent documentation.** An investigation that cannot withstand scrutiny in court, or cannot be understood by the public and the witnesses, will not achieve its purpose. Credibility depends on transparency throughout the in-loco inspection, classification and evaluation of evidence, interpretation and the fairness of the outcome.

Adopting forensic investigation will require a revision of regulatory and operational approaches. Independent investigators would need the authority to take photographs and biological samples, take custody of documents and data, and interrupt operations where necessary. It suits organisations that are already mature and are ready to move to the highest level of maturity. The opportunity has been at the industry's door for a decade. The organisation that opens it will lead the industry in safe operation.

Unintended — but Never Causeless

Many risk assessments describe events as *inadvertent*: "inadvertent movement of a dump truck", for example. Technically the description is correct, but labelling an event as unintentional can imply that it had no cause. Bow tie analysis, widely used in the industry, starts from a defined event and identifies its causes and consequences. If facilitators assume that almost every event is an unintentional or unplanned release of energy — reasoning that nobody plans an accident — they may stop looking for the mechanisms behind it.

Then, when an event occurs, investigation techniques often look for a deliberate act or omission by the person responsible. A system that focuses on *unplanned* events when assessing risk, but searches for *intent* after the event, is internally inconsistent. Consider two examples.

> **Two everyday events.** Two dump truck drivers reverse out of a congested parking area after a thorough pre-use inspection, a safety talk and an awareness discussion, and their trucks collide with minor damage. The collision avoidance system was installed but deactivated for the congested parking area. Neither driver planned the collision.
>
> A laboratory assistant accidentally bumps the outlet valve of a chemical container while setting up a cleaning process, causing a small spill. Because the chemical is hydrofluoric acid and several people are exposed, the event is classified as negligence on the part of the assistant.

These are two of hundreds of everyday events whose meaning changes depending on the investigation method used. Consistency between how we assess risk and how we investigate events is a simple and powerful improvement. The question in both cases is not "who intended this?" but "what design, condition or decision made this outcome possible — and how do we prevent it everywhere?"

Investigate Before the Event

Historical human error analysis offered a broad framework for human factors in task execution. Factual investigation goes further: it does not assume that people will, at some point, make a mistake. An integrated approach examines a broader spectrum, beyond artificial simulation. Moving from reactive to proactive means **investigating risk source events during the risk assessment process** — without witnesses, with external influences removed, with reality checks on control effectiveness, and with resilient systems showing opportunities.

In simple terms, criminal law looks at past events: if causation is proven, together with certain other facts, the guilty are charged. The mining industry already applies a similar logic proactively when it drafts critical control standards. Take one step back: the risk sources, trigger events, outcomes and causes are the same before and after an event. So why spend so much time on investigation and so little on risk assessment?

Silent witnesses cannot speak, but those close to them can tell us how little was done before the event. If we give hazards the same depth of analysis in the risk assessment that we give events in the investigation, investigations will confirm what we already knew rather than reveal what we missed. Regulators require investigation; they participate, recommend actions and may issue fines. Inviting them into risk assessments as well would be a significant step forward.

> **Groundbreaking insight:** Accidents should not merely be investigated; they should be *analysed in the context of the risk assessment that preceded them*. When the two processes share one evidence base, trust grows — and the organisation of the future is built on that trust.

The Causation Path of a System Failure

Event causation aims to unpack the correct sequence of occurrences leading to an event — the causation path. System failures are rarely analysed with the same rigour as equipment failures; they are often recorded simply as causes, and accident methods tend to stop once underlying causes have been identified. Understanding the causation path of a system failure is crucial for both response and prevention. Reliable systems share several characteristics.

- **Defined context and scope, with system monitoring indicators.** A system without key monitoring indicators may fail under demanding circumstances. At a minimum, monitor how often the system has been challenged and how often it has partially or totally failed. A fitness-to-work examination at placement is effective only on the day it is done: a haul truck driver who develops a heart condition, and does not report it for fear of losing income, illustrates a system failure. Similarly, one code of practice may set the minimum age for operating mobile equipment on a mine at twenty-one, while another permits drivers with a Code 08 licence to operate — a disconnect between the owners of the two codes that only a system view will reveal.

- **Consistent trigger events.** The trigger events identified in the risk assessment should be the same trigger events used in the investigation. They should not be confused with immediate or underlying causes. Trigger events have pre-conditions, and pre-conditions have causes within a system. Examining them reveals the operational environment in which the system was developed and how it adapts to change. A person's state of mind reveals how the system can be breached and the reasoning behind active failure pathways.
- **Factual contributing factors.** Contributing factors are factual influences and should not be mixed with uncertainty or assumption. Any assumption, even if documented, should be set aside. In many investigations, human factors are classified as unknown or left to the assumptions of an investigator or witness with little knowledge of psychology. Reviews of investigation outcomes show that "human error" is frequently treated as a fact. This is perhaps the clearest opportunity for improvement in current causal analysis and "why" questioning practice.
- **System reliability analysis.** Reliability engineering monitors and evaluates equipment in detail. The same principles can be applied to systems. The methods used to determine at what condition a bearing or motor shaft will fail can determine a system's threshold before failure — its capacity to withstand interference from people or the environment. The transformation is to superimpose quantitative reliability engineering onto system reliability, which is currently judged qualitatively, using semi-quantitative techniques that validate decisions before a change is recommended.
- **Organisational factors from the top.** Traditionally, organisational factors include context, policies, procedures, training and culture. Contemporary event analysis focuses on the tone at the top, management involvement and the design of the risk management system. Simulating events and their impact on reputation, the social environment and governance is an important part of continuous improvement.
- **Genuine system improvement.** A desktop review of existing investigations shows a familiar pattern: find fault, identify root causes, define remediation and record lessons. Following a predefined questioning pathway from the start limits what can be found. Drafting a procedure, reviewing the risk assessment, enhancing training, taking disciplinary action and redesigning a conveyor take-up pulley all share one feature: on their own, they are not system improvements. Real improvement measures a system's actual effectiveness, efficacy and efficiency, and changes it.

Established methods have served the industry well and have many dedicated practitioners. Their limitations are now well understood, and the profession is ready for the next step: investigation led by people who specialise in forensic investigation.

Key Takeaways

- Give hazards in risk assessments the depth of analysis we give events in investigations.

- Strengthen investigations with factual and legal causation, evidence validation and independent review.
- Learn from aviation and major public inquiries: investigate to prevent, not to blame.
- Be consistent: the same trigger events and the same logic before and after an event.
- Treat “human error” as a starting point for analysis, never as a finding of fact.
- A procedure, retraining or discipline is not, by itself, a system improvement.

In Part IV, we enter the technical heart of risk management — controls — and begin by replacing ambiguity with clarity.

PART IV: Mastering Controls

Controls are where risk management meets reality. Part IV is the technical heart of the book. It replaces ambiguity with clarity, shows how to measure whether controls actually work, and sets out a mature approach to key risk indicators, bow tie analysis, critical control management and response planning.

Chapter 9: Clarity of Control

> *"For a successful technology, reality must take precedence over public relations, for nature cannot be fooled." — Richard P. Feynman, Appendix F to the Report of the Presidential Commission on the Space Shuttle Challenger Accident (1986)*

ISO 31000 defines a control as a measure that maintains or modifies risk, and notes that controls may not always exert the intended or assumed modifying effect. The definition is deliberately broad. Its breadth is one reason why almost anything can end up on a risk register as a "control", whether or not anyone understands its objective. *Anomalies*, by contrast, are the unexpected: deviations and irregularities. Variations in perception create uncertainty, and uncertainty is where events begin. Ambiguous controls — controls whose purpose is unclear or difficult to understand — hold employees and systems to ransom. They carry an expectation of malfunction, and they undermine continuous improvement.

Why Ambiguity Creeps In

Controls for material risks are often designed primarily to meet legal requirements. That gives them a strong foundation, but it can also hide opportunities. People naturally do what is expected; anything more might appear to expose a manager to additional liability, so a generic approach can feel safer. With intangible risks — a cyber-attack on payroll, or security intelligence in the diamond industry — controls may be deliberately obscure, for good security reasons. In complex systems, the dynamics themselves make it hard to identify system behaviour and to monitor effectiveness.

Mature risk management structures can map exactly where ambiguous controls appear in a system. The process relies on detailed analysis of intended system behaviour, anticipated deviations and response plans based on risk velocity and critical control points, expressed quantitatively. It typically includes periodic system diagnostics followed by detailed forensic examination to expose vulnerabilities.

> **Case study — The automated winder directive.** Following a conveyance accident in a shaft, a regulatory directive required that whenever a conveyance transports people, a winding engine driver must be present at the controls. Several winders in the industry had been operating automatically for many years, successfully and without exposure to human error in their operation — which was exactly what automation was designed to achieve. Under certain conditions, these automated winders were operated manually by qualified winder electricians and technicians, not by conventional winding engine drivers, who did not have the specific knowledge or experience. The directive, issued with the best of intentions, placed human error back into a system that had been engineered to remove it. The lesson is constructive: regulatory responses to major events are strongest when professional engineers familiar with the specific technology help to shape them.

The Concept of Risk Control

To appreciate risk control and risk response, it helps to revisit the classic categorisation models and compare them with mature contemporary strategies. The hierarchy of controls, James Reason's Swiss cheese model, trigger action response plans, barrier analysis, control matrices, layers of protection, control failure probability, control reliability and risk responsiveness each have a specific purpose. Each also contributes, when misapplied, to an illusion of effectiveness. Adaptations of these models have been made over time, some helpful and some not. The discipline is to use each for what it was designed to do.

Two habits deserve particular attention.

Turning a control failure into a hazard. Roof bolts control the convergence of an underground excavation. If the hazard is recorded as "inadequate support" rather than as the gravitational energy of the rock mass, the risk assessment shifts its focus from the energy to the control — and the true risk source, which has not changed, disappears from view. Always name the energy; then list the controls that manage it.

Diluting consequence with likelihood. Consider a risk assessment of a collision between a bus carrying more than fifty employees and a third-party vehicle. A team can spend hours debating the likelihood that all passengers will be fatally injured. The sound approach is simpler: first assess the likelihood of the event, and only then the maximum reasonable consequence, taking into account the number of employees and the occupants of the other vehicles involved. When ranking becomes more important than understanding, risk management objectives are diluted — and requests for a second ranking simply to "show the improvement impact" of controls take us further from the truth.

Black Swans in the Mine

Nassim Nicholas Taleb's black swan metaphor describes rare, high-impact events that we underestimate beforehand and explain away afterwards. At a mining conference in 2011, the author presented a paper on the black swan metaphor in the context of mining. The research and discussion made one thing clear: the industry was busy with activities believed to reduce the number of fatal accidents, while the rare, catastrophic event remained largely untouched by the risk matrix.

Consider a room full of risk experts who remain unconvinced that a consequence score drawn from a subjective impact table will change focus, reallocate resources or measure control effectiveness. Should a lower number at least show better control? Not necessarily. The "number game" is most readily accepted when the point of prevention is placed at the unwanted event, when it should be placed at the trigger or initiating event. These scenarios occur daily when a risk assessment departs from its essential practices.

Control philosophy rests on two features that give rise to risk: the calculated prediction that a specific risk event exists, and the consequence should it materialise. Each subdivides into risk

velocity, risk magnitude, hazard drift and risk deflection, all of which belong in the control framework. A low likelihood does not mean an event will not happen soon; it may happen a second after the assessment is signed.

The Challenger disaster of 1986 is the classic illustration. On the eve of the launch, engineers raised concerns about the performance of the solid rocket booster O-rings in cold weather. The Rogers Commission found that the burden of proof had, in effect, been reversed: rather than being asked to prove that it was *safe* to launch, the engineers were asked to prove that it was *unsafe*, and without conclusive data the launch went ahead. The purpose of recalling it here is not to second-guess complex decisions, but to ask whether we can define control effectiveness and sustain a quantitative monitoring system strong enough that the burden of proof always rests where it belongs.

Rethinking the Hierarchy of Controls

The hierarchy of controls remains a valuable guide, but it contains a myth or two. Mining uses risk sources to do business: an electrically driven hoist motor cannot be eliminated. In practice, elimination is rare — unless one counts replacing coal-fired power with solar power, or removing an electrical cord from the floor to prevent a trip. The mature view is that we identify risk sources and adjust them to a state in which they serve as a resource and in which we have the competence to manage any outcome or change.

Administrative controls depend on a person's actions, the behaviour of the interaction and the practicality of implementation. People act not on the risk they are exposed to, but on the risk they *perceive*. They assess risk through emotion, which is where the control heuristic meets the decision. Because administrative controls are easy to adapt to personal disposition in order to "get the job done", they need the strongest verification.

The traditional view that personal protective equipment is less effective than administrative controls deserves a second look. An electrician carrying out live fault-finding in a full arc-flash suit relies on administrative procedures to maintain the suit's effectiveness; and a sign instructing people to wear goggles is itself administrative. PPE and administrative controls are interdependent, and both rely on behaviour.

Controls Have Objectives

The reality is that every control has an objective, and the aim is to capture all the risk *modifiers*. Modifiers fall into two groups: those that release the risk source, and those that manage the event once it has begun.

- The first and best option is to design controls that prevent or mitigate the modifiers, not the event.
- Control effectiveness is measured in the same two groups.

- The primary interest is the first objective: managing prevention and the likelihood of the event.
- The second objective, mitigation, is consequence management. It is essential, but on its own it moves us away from the Zero Harm philosophy.

The adaptation from older to newer approaches is slow compared with the rapid growth in control demands created by new technology. Closing that gap is essential if our understanding of effectiveness is to keep pace with the complex, sophisticated systems now entering the workplace.

Control Effectiveness

Effectiveness is the degree to which a control can be confirmed to achieve its intended objective. In practice, it tells us how often the associated hazard release or trigger event was encountered, and how the control actually performed against its objectives.

Effectiveness reflects not only a control's theoretical ability to manage a risk, but its actual performance across five attributes. This book uses the same five throughout:

- **Reliability** — does it perform its function every time it is demanded?
- **Availability** — is it in place and ready when needed?
- **Survivability** — does it continue to function under the conditions of the event?
- **Dependability** — does it operate in time, consistently, on demand?
- **Compatibility** — does it work with the other controls, people and systems around it?

Effectiveness can be expressed quantitatively or qualitatively, in absolute or relative terms, and is usually reported together with a description of residual risk.

Measuring the Response

Control response can be measured as the ratio between the number of times a control was challenged and the number of times it achieved its intent. If the response is observable, specifiable and quantifiable, it becomes an operational key risk indicator. Many such measures have been developed over the years, with credit to the pioneers who showed what an effectiveness measure should look like. In essence, they compare challenges with failures over time. A pedestrian detection system on a mobile machine illustrates the options, in order of preference.

- **Real-time system data.** The system is designed to detect a pedestrian at a specified distance, warn both the operator and the pedestrian, and slow and stop the machine automatically if necessary. The best measure is the system's own data: the number of interactions compared with the responses taken.
- **Test data.** During planned maintenance, the system's operability is tested. Record how many times it was tested, how many times it worked as designed and how many times it failed.

- **Before-and-after event data.** Compare the number of unwanted events over a defined period before and after the control was introduced.
- **Bypass data.** The system can be overridden in specified situations — a hard park bay, a workshop, or a congested loading area beside a shovel. Every bypass creates an opportunity for an event and should be measured: how many times the system was bypassed, and how many proximity detections were recorded while it was bypassed.

Every effectiveness measure needs an unambiguous description of the control and its intent, the hazards and trigger events, the design qualification and the operational performance levels. It is also essential to define what constitutes an *unwanted*, *intermittent* and *wanted* failure.

- **Unwanted failures** are flows of energy in an unwanted direction, releases of energy towards a target, or losses of structural integrity. They can be identified in the early stages of failure, and the point at which failure becomes imminent can be determined quantitatively. If the risk assessment has identified the correct point of failure, back-analysis can establish the critical point at which the situation can still return to normal. That is where the control's objective should focus, and where its response can be measured. Engineering has led the way with equipment and structural condition monitoring. The opportunity now is to protect those systems from cost-driven drift from preventive towards breakdown maintenance.
- **Intermittent failures** occur periodically, without notice, and are hard to define. They reveal cracks in existing controls. A control that fails only occasionally may be judged "acceptable", and qualitative reliability systems that use percentage acceptance levels can miss intermittent failure entirely. A control may be rated effective because it sits at the engineering level of the hierarchy and shows implementation above ninety per cent, yet its intermittent failure rate has never been measured. Measuring it is the mark of a mature system.
- **Wanted failures** are designed into a system and can be defined early in design, but are often overlooked in risk assessment. The humble hook on a shaft conveyance is a perfect example. For years it provides a stable connection between the hoist rope and the conveyance. But in an overwind, if every other safety device has failed to stop the conveyance, the humble hook must *fail by design*: as it passes through the spectacle plate, it opens and releases the rope before the conveyance is pulled into the sheave wheels, allowing the conveyance to drop safely onto the jack catches.

Layers of Protection

Layer of Protection Analysis (LOPA) offers a semi-quantitative way to assess control effectiveness. On its own it may not be the best performance measure, but combined with bow tie analysis it can discriminate powerfully between controls.

- Controls of similar design and intent are aggregated into control families.

- Controls requiring different levels of assurance are systematically selected.
- Defined criteria isolate the most critical controls, on which other controls depend.
- Dependent controls are subject to the effectiveness of a critical control: if it fails, they fail with it.

The defence lines designed through LOPA provide excellent quantifiable response points along the event pathway. They influence and measure the reduction of risk velocity, and they are fundamental to a trigger action response plan.

Consequence Controls

Consequence controls have similar effectiveness criteria but need to be measured during and after events — which means they must be tested through simulation, without waiting for a real event. ISO 31000's risk treatment options refer both to control effectiveness and to changing the magnitude of consequence. If an organisation's safety objective is Zero Harm, consequence controls deserve as much attention as likelihood controls. Traditional control thinking struggles here for several reasons:

- Organisations perform best with tangible risks, whose consequences are obvious; intangible risks require more detailed analysis.
- Many risk assessments are a record of the team's collective knowledge and opinion, so controls and their effectiveness are subjective, especially where human factors are involved.
- Many organisations use a control effectiveness technique alongside bow tie analysis, but do not use it as a criterion for selecting critical controls.
- Escalation factors are often omitted from effectiveness criteria, so critical control regimes become administrative and a supervisory burden.

For many years, organisations have reviewed risk assessments and added new controls. Far less emphasis has been placed on evolving effective control strategies that target the release mechanism and place controls as close as possible to the risk source. The natural tendency is to reduce the likelihood of the event — the brain takes the easier path. Consequence control effectiveness should be quantified in relation to the magnitude of the risk source, through a comprehensive process. Templates are useful starting points, and the best advisers and software suppliers now build control effectiveness analysis into them. Informed risk managers will choose the partners who simplify controls, not those who complicate them.

Five Dimensions of Effective Control

Likelihood and consequence control effectiveness share common ground but should be considered in distinct categories. Using all the available dimensions of control moves us beyond traditional effectiveness ratings. Target effectiveness in five dimensions:

- **Shift left:** move the control framework along the timeline towards the threat, as close as possible to the release mechanism of the trigger event, and ultimately into the pre-condition phase.
- **Shift right with purpose:** make a similar shift towards controls for subsequent events, mitigating downstream effects.
- **Think in three dimensions:** move control analysis from a single characteristic to all consequence categories — health, safety, financial, production and more — revealing the total consequence domain. This requires bow tie software that can integrate control frameworks; spreadsheets are not designed for it.
- **Cascade through the organisation:** evolve effectiveness measures through the hierarchy of supervision and management, down to the worker and task level.
- **Control the escalation factors:** escalation factors that defeat control objectives need their own controls and effectiveness measures.

A Worked Example: The Moving Conveyor Belt

The criteria for measuring control response reside in the design of the control and its objectives. A control's ability to function under normal and abnormal conditions shows in how well it manages the magnitude, frequency and velocity of the encounter. Historically, controls were developed as risks arose and as fires flared up. A far more practical approach is a methodology founded on the timeline dimensions of control. The risk of a moving conveyor belt illustrates how it works.

Segmentation — populate controls in functional units.

- Guarding on the conveyor belt is a single preventive unit.
- A position indicator that reports whether the guard is closed or open to a control room is a single monitoring unit.
- An alarm that sounds if the guard is opened while the conveyor is running is a single first-response unit.
- An interlock between the guard and the conveyor power supply is a single first-response control unit.
- A proximity device that trips the power supply when a person enters the danger zone during maintenance or training is a single first-response unit.
- An automated mechanical interlock that prevents runback of a tail or drive pulley during gravity imbalance is a single first-response unit.
- An automated release of the belt take-up tension if a person is drawn into the belt is a single emergency-response unit, acting on the failure of the first-response units.

Structural dimensions — define each unit's objective and failure mode against the release mechanisms it must impede, reduce or terminate.

- Guard: may not prevent every revolving-part release mechanism.
- Position indicator: fails with faulty detection instrumentation.
- Alarm: fails mechanically; test it manually and with noise instruments at defined levels and distances.
- Proximity device: fails when dust accumulates on the sensors.
- Mechanical interlock: fails when structural integrity is exceeded under abnormal overload.
- Belt tensioner: releases tension only over a limited distance from the take-up, and inclined systems remain under gravity tension.

Simulation — use two- or three-dimensional simulation to test control response before implementation.

- Position control devices — rip detectors, belt alignment switches and pull-wires — at their design locations along the entire length and installation height of the conveyor.
- Activate technological systems for each scenario, such as interlocks operated from the control room or locally.
- Apply human activities under normal and abnormal conditions: pedestrians, maintenance, functional testing in motion, inspections.
- Quantify the magnitude of the risk release before and after each control unit activates.

Quantification — measure the five effectiveness attributes (reliability, availability, survivability, dependability and compatibility).

- Define control reliability confidence intervals that reflect actual test results at a chosen significance level.
- Critical controls can then be defined on quantifiable evidence: automated controls with a high effectiveness rating and a simulation confidence of more than 95 per cent (a significance level of 0.05).
- Keep simulated information for the entire life cycle of the control.

> **Groundbreaking insight:** We can no longer base our major risk control effort on the opinion of a single subject matter expert. With today's technology, every critical control can be simulated, tested and measured before a worker ever depends on it.

From Opinion to Evidence

Applying these principles requires a resilient process and education in the dimensions and functional capacity of controls. The challenge for organisations is to design control frameworks integrated with real work activities and explained in a way that every worker understands, in the language they speak. Good progress has been made in recent years, and several organisations have raised their appetite for demonstrating due care.

Rigorous control structures are essential on the journey towards resilience. Organisations will need to invest in risk research and, where appropriate, partner with educational institutions. Qualification and quantification of the risk source, from cradle to grave, remains the most important element in specifying suitable and sufficient life-cycle controls. Capture those controls in real-time event scenario simulations, verify their effectiveness with factual research data, and design controls that reduce impact at the point of release, at every change in magnitude and throughout the life cycle.

Key Takeaways

- Name the energy, not the control failure, as the hazard.
- Assess likelihood first, then the maximum reasonable consequence; do not rank twice to "show improvement".
- Place prevention at the trigger, not at the unwanted event.
- Measure effectiveness through five attributes: reliability, availability, survivability, dependability and compatibility.
- Define unwanted, intermittent and wanted failures for every critical control.
- Use segmentation, structural dimensions, simulation and quantification to design controls.

In the next chapter, we turn effectiveness into early warning: key risk indicators that measure what matters before the event.

Chapter 10: Key Risk Indicators — Measuring What Matters

> *Imagine a primary crusher that knows when it is unsafe — that closes its own access, warns everyone nearby and isolates its own energy before a person can reach the danger zone, and that cannot be overridden by anything except its own control system. In that plant, the machine decides when it is safe, not the person. That is where key risk indicators are taking us.*

Success is built on measuring performance against a standard. The standard is as important as the performance criteria, and understanding the hazard and risk source tells us what to measure, how to measure it and how to compare it with established best practice.

What a Good Indicator Measures

Information for measuring risk performance should share four features:

- **The pre-event stage:** the period during which the energy and risk source are stable and serving the business.
- **The release pathway:** how the risk source manifests in the workplace, and which mechanism will discharge it in a particular direction.
- **Speed and magnitude:** how quickly, and with what force, energy travels from the point of no return to the target of harm, loss or impact on objectives.
- **Transitions:** between the stable phase and the point of loss, a risk source passes through transfer stages and changes of state, each producing signals of failure. The energy that must be released, or the forces that must be applied, to produce a change of state per unit of time can be measured and validated.

Take a primary crusher. Rather than measuring the reasons for breakdowns or accidents, seek information that measures the crusher's performance. Analysing performance reveals each energy source with its function, and each has a stable, a controlled and an uncontrolled state. The uncontrolled state is not one in which controls have "failed"; it is the state in which the energy's magnitude and speed of release exceed the capacity of the control parameters. These phases define key performance areas: the crusher's electrical current and any sudden increase or decrease; the amount of interaction with people outside the crusher; and the number of times a person attempts to enter the crusher, and whether that happens in a stable or unstable state. Broken down into actual data comparisons, these areas show the forces involved and the capacity of the control — which is, in principle, the critical risk performance indicator.

The Mature Measurement System

Measuring risk performance well requires a mature risk management system whose vocabulary has already moved from numbers to big data, and from engineering controls to artificial intelligence. Monitoring critical controls only through manual checklists is the first step on a long road. Asking questions during a visible felt leadership visit is valuable for engagement, but it is not a substitute for monitoring critical risk performance.

Automatic monitoring can track energy, forces, movement and change, producing a far more accurate basis for managing out-of-specification indicators. Artificial intelligence can take over high-risk assignments, using big data to support self-regulating control within a controlled environment. Machinery such as a primary crusher should not require a person to operate it or approach it closely; if access is needed, the system should close access, warn and stop all energy in time, and resist override by any system other than its own.

Best practice is a useful benchmark, but it is not necessarily the best practice for *your* risk. Legislation, likewise, sets the minimum requirement. Defining the pre-event state establishes the standard for a controlled state within design parameters, and with it the area to monitor and the criteria to measure against.

Transformers: Measuring at the Right Point

Transformers are used in every mining operation and carry a significant risk of power loss to a plant and of explosion. Considerable time and money go into preventive measures, often following insurance and legal audits, and those recommendations are generally based on a past event or on perceived best practice.

Key risk indicators should focus on the early stage of the event — the pre-condition state. For a transformer, that means the insulation and cooling design criteria, not whether a fire barrier was installed between transformers. The forces affecting dielectric strength can be measured as a rate (how quickly breakdown voltage deteriorates) and quantified (for example, in kilovolts per centimetre) to detect change. At the other end of the event, many systems monitor moisture in the oil only through an indicator, without actual data on moisture content — and miss the rapid increase in tank pressure that moisture can cause, which is what produces the explosion. The question is therefore: what is the best practice — the design state, or the moisture indicator? The answer is to measure the design state.

> **Groundbreaking insight:** The closer an indicator sits to the point where the risk source is released — where it begins to drift or change velocity — the more effective prevention becomes, and the fewer resources are needed for mitigation.

Confidence: The Missing Ingredient

Managers need assurance that monitoring systems produce enough information to support their decisions with confidence. Very few monitoring systems in mining consider confidence levels when defining achievement. Operators should use confidence levels, and the quantity of data needed to verify a control's efficacy. Risk managers of the future will design key risk indicators at 95 per cent confidence, and management will be able to analyse the deviations.

A sound indicator structure uses samples of a defined accuracy, a specified degree of acceptable variability and a size suited to the organisation. Control effectiveness is affected by measurement variability and sampling error. Suppose a control owner samples compliance with isolation and lockout procedures — through physical inspection or from control response data — and finds 90 per cent compliance. With an adequate sample and a margin of error of five per cent, the owner can conclude with 95 per cent confidence that between 85 and 95 per cent of all employees comply (Israel, 1992). Without that discipline, the 90 per cent is simply a number.

The familiar 80/20 Pareto principle is a helpful rule of thumb, but it is not a mathematical law and it carries no sampling precision. Occupational hygienists have used statistically designed sampling for years when measuring airborne pollutants and noise. Translating the same principles to risk sources with more variables is challenging, but achievable. Monitoring regimes should be designed so that random, once-a-week inspections are not the primary evidence. Critical control verification through occasional inspections and observations by managers will rarely meet the statistical requirement; continuous data will.

Choosing the Right Analysis

The data analysis method matters as much as the data. For a descriptive analysis of compliance, frequencies and deviations from a mean will suffice, and any number of observations will do. To determine multiple causes of deterioration, covariance between causes that share a trend, or relationships between causes, more advanced methods such as log-linear analysis are needed.

What does this mean in practice? Quantified risk criteria produce a measure of deviation from a practice that defines the likelihood of an event and, with a defined accuracy, outline the consequence of a breach. As confidence in control effectiveness is quantified, risk assessment maturity increases. Key risk indicators then provide accurate data for early warning systems that proactively stop an event, or deflect it to a wanted state.

> **Case study — The cage door and the penthouse.** A double-deck shaft conveyance carried thirty persons per deck, with numbers controlled by a counting turnstile at the start and end of each shift. Tradition had it that at the end of the shift, everyone wanted to reach the surface as quickly as possible — especially when there were several shafts to travel through.
>
> One day the turnstile was defective, and loading reverted to manual counting. Employees crowded in front of the conveyance, and the onsetter could not control the push from the back. The last person to board did not know the count, and squeezed himself in so that the door could

> be closed. The onsetter managed to close the cage door, but it bowed outwards with the pressure of the people inside, with the door interlock intact and working. As the conveyance left the station, the bowed door caught on the penthouse above the station, was torn from its runners, and four persons fell down the shaft before the conveyance could be stopped in mid-shaft. It was a tragedy without precedent at that operation.
>
> A cage door makes people inside feel safe: they cannot see the danger outside. The deeper questions concern the system, not only the individuals. Why did a single defective turnstile remove the only quantitative control on load? Why was there no detection of door deformation before the conveyance moved? Key risk indicators with detective, predictive and surveillance monitoring rules — load sensing, door position and profile monitoring, and an automatic hold on departure when either was out of specification — could have prevented it.

What Falling Statistics Do Not Tell Us

If declining industry accident statistics create a feeling of safety, it is worth pausing. In 2025, South African mining recorded 41 fatalities, its lowest annual total ever — and by mid-September 2026 the year-to-date total had already exceeded it. A falling trend reflects history, not the future. If we do not understand the human nature of risk, history will repeat itself.

Standards, procedures and engineering specifications are all essential. But they are only as good as the mathematical modelling we use to understand the variables of human nature that must deal with them in the seconds before an event. Key risk indicators are how we bring those seconds into view in advance.

Key Takeaways

- Measure the pre-event state, the release pathway, speed and magnitude, and transitions.
- Monitor performance of the risk source, not only the causes of past breakdowns.
- Place indicators as close as possible to the point of release.
- Design indicators with defined confidence levels (typically 95 per cent) and adequate samples.
- Continuous data beats occasional inspection for critical control verification.
- A falling accident rate describes the past; leading indicators describe the future.

In the next chapter, we explore the technique that brings threats, controls, consequences and indicators together on a single page — bow tie analysis, done right.

Chapter 11: Bow Tie Analysis Done Right

> *Did you know that in a mature bow tie for an open pit slope, the top event is not "slope failure" at all, but "driving forces exceed resisting forces"? Move the top event a few steps closer to the hazard, and preventive controls you thought you had become mitigating controls — and controls you never knew you needed suddenly appear.*

Risk management is about coordinating activities and directing controls for identified risks. It rests on decisions informed by sound risk evaluation, analysis and good corporate governance, which together produce the company's risk appetite. Every phase of the risk management process — assurance, scope and context, assessment, control response, monitoring and review — has its own designated methodologies. For many years, the mining industry had reasonably suitable systems for every phase except control and monitoring. In seeking a solution, it turned to a technique rooted in the fault tree and event tree analysis of the 1960s and 1970s. The bow tie diagram is often traced to hazard analysis course material at Imperial Chemical Industries in 1979, and it gained global momentum when Shell adopted it in the early 1990s in the wake of the Piper Alpha disaster. It has since evolved into a well-defined technique for evaluating threats and controls and analysing barriers.

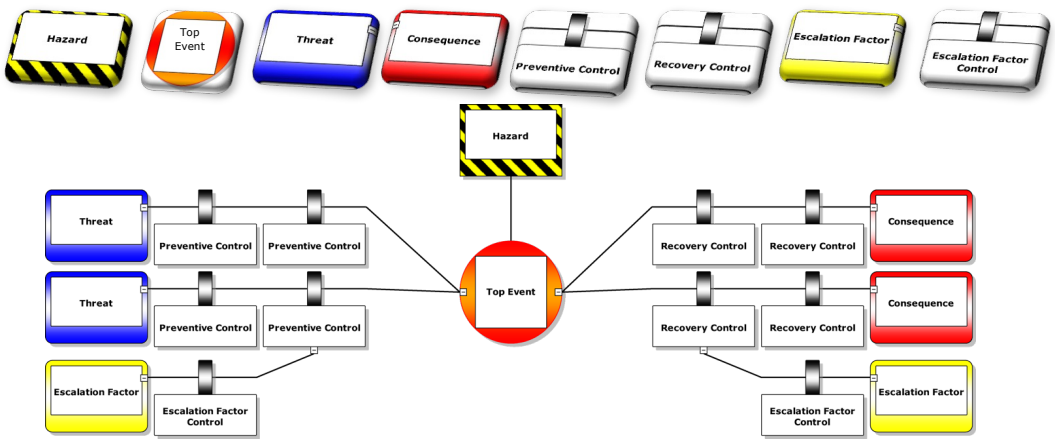

Figure 3: Bow tie analysis terminology

How Mining Adopted the Bow Tie

True to its character, the mining industry developed a common approach, drawing on the long experience of the petrochemical and chemical industries, which were already showing results. After extensive stakeholder consultation, the International Council on Mining and Metals (ICMM) published its health and safety critical control management guidance, comprising a good practice guide and an implementation guide, which it continues to update. Readers should always refer to the latest version. The aim here is not to reproduce that guidance, but to examine how its principles have been implemented in practice.

Implementation illustrates a phenomenon first described by the German psychologist Wilhelm Wundt in 1886: the *heterogony of ends*, in which the pursuit of a goal produces side effects that gradually change the goal itself. The intention was to follow the established objectives of bow tie analysis, with critical control management as a natural extension. In practice, in many organisations the objectives shifted towards the anticipated rapid reduction in fatality rates. Bow tie analysis became, for some, merely a way to identify critical controls, rather than a full process of barrier analysis and advanced control management. Layer of protection analysis, integral to the original thinking, was often left unexplored.

The industry had a well-defined programme and was on a sound track when the guidance arrived. Mature companies set the pace. Others, newer to the technique, sometimes simplified it into basic verification questions on spreadsheets. The opportunity now is to return to the full purpose of the technique — and the reward will be visible the next time a significant event is *not* followed by helicopters landing and communities gathering at the gate.

Six Lessons from Taleb

Nassim N. Taleb, Daniel G. Goldstein and Mark W. Spitznagel (2009) described six mistakes executives make in risk management. Each has a direct lesson for bow tie practice.

- **We think we can manage risk by predicting extreme events.** A baseline risk assessment followed by a bow tie does not prove that we can manage the event. The bow tie's value lies in revealing vulnerability, not in predicting the future.
- **We are convinced that studying the past will help us manage risk.** Past accident analyses, statistics and trends are useful, but they cannot on their own establish control effectiveness or likelihood.
- **We don't listen to advice about what we shouldn't do.** "What not to do" — avoiding fragile designs, removing single points of failure — is often more valuable than adding another control.
- **We assume that risk can be measured by standard deviation.** Statistical averages can hide the extremes that matter most.
- **We don't appreciate that what is mathematically equivalent isn't psychologically equivalent.** How a risk is framed changes how people respond to it — which is why human factors belong inside the bow tie, as escalation factors.
- **We are taught that efficiency and shareholder value don't tolerate redundancy.** Profit forecasts that leave no room for risk variation are fragile; markets react sharply to major events.

The positive conclusion is that the industry is very well positioned to identify significant exposure to low-probability, high-impact events, and is prepared to invest in managing critical controls. Combining engineering controls with a genuine understanding of human psychology will close the remaining gap.

What the Bow Tie Is — and Is Not

Risk management is not accident investigation, and accident investigation does not, by itself, prevent accidents. Managing risk means anticipating, deterring, preventing and mitigating unacceptable risk — taking risks on calculated, pre-set criteria aligned with the company's risk appetite.

Every mining company has a formal risk management policy that deals with risk, qualitatively or quantitatively, at different levels: from corporate enterprise risk management through middle management to worker involvement. Some use three tiers, others four. The number of tiers matters less than the interaction between them and the flow of information and risk knowledge.

Baseline and issue-based risk assessments tend to depict history. The so-called high risks have remained much the same for decades, shifting occasionally with the latest event. The industry relies heavily on recorded consequences to decide what counts as high risk, and every new event gives the impression that existing controls cannot cope. Bow tie analysis, by contrast, lends itself to scenario setting. It has revealed numerous events never recorded before, exposing absent or misplaced controls. Understanding the framework in which the bow tie is used is therefore critical to the quality of its outcome.

To use the bow tie well, understand its language and philosophy:

- It is **not** a risk assessment, an investigation or a problem-solving technique.
- It **cannot** weigh one risk against another, force a decision that one control is more effective than another, or eradicate the risk.
- It **is** a way to manage a suite of controls against predetermined criteria that make the risk acceptable for our work — and it reveals the vulnerability of risk owners.

The risk assessment produces a prioritised list of risks. The bow tie then provides the specific suite of controls that manage a particular threat or outcome, and identifies those that are critical to managing systems, physical activities and their monitoring.

Start with the Life Cycle

Information gathering for a bow tie is as important as for any other risk activity. An issue-based risk assessment in the form of a **life-cycle risk assessment** reveals the hazard in its complete form, from cradle to grave. A life-cycle map shows how the hazard changes in magnitude, how it interfaces with other hazards and external forces, how it drifts, how release velocity increases, and finally the mechanism that releases the energy at a point of no return. That information is plotted on multiple event timelines — one hazard, several outcomes — and the timeline identifies the point at which control is lost: the top event.



Figure 4: The top event

The Top Event Belongs Next to the Hazard

Mature bow ties place the top event close to the hazard; less mature ones place it at an unwanted event or even at a consequence. In an open pit, the hazard is the gravitational force acting on the pit wall. A mature company uses “driving forces exceed resisting forces” as the top event. A less mature one uses “slope failure”. The least mature uses “inundation of people by a slope failure”.

The position of the top event determines how the control framework is plotted. Place it too far to the right, and preventive controls become mitigating controls, and vice versa. A simple example: earth leakage protection is sometimes plotted as a *preventive* control for electric shock. Anyone familiar with electrical systems will point out that earth leakage protection operates *after* current has begun to flow between the source and earth — it is a mitigating control. The positioning of the top event is therefore the most important single decision in a bow tie.

Hazard	Resilient	Pro-active	Compliant	Re-active	Basic
Hanging wall	Driving forces exceed resisting forces	Changes in hydraulic forces / increase in stress fields	Loss of cohesion	Uncontrolled fall of ground	Fall of ground, fatalities, loss of production
Flammable gas	Flammable gas ahead of cutting	Explosive atmosphere	Ignition source in contact with flammable gas	Uncontrolled release of flammable gas	Gas explosion, production loss, damage to equipment
Tailings dam wall	Change to the undrained shear strength of the encrusted tailings	Change in the water table position between the zone of saturation and the zone of aeration	Increase in phreatic surface (tailings) – pore water pressure under atmospheric conditions (pressure head is zero)	Overflow of tailings	Dam wall failure, environmental pollution
Trackless mobile equipment	Inability to operate equipment under normal neurological conditions	Deteriorating integrity of brake system	Inadvertent movement	Loss of control over equipment	Collisions, overturning, production loss
Structures	Increase in the rate of deformation	Loss of tensile strength	Loss of structural integrity	Uncontrolled failure of structure	People trapped under structure, loss of equipment

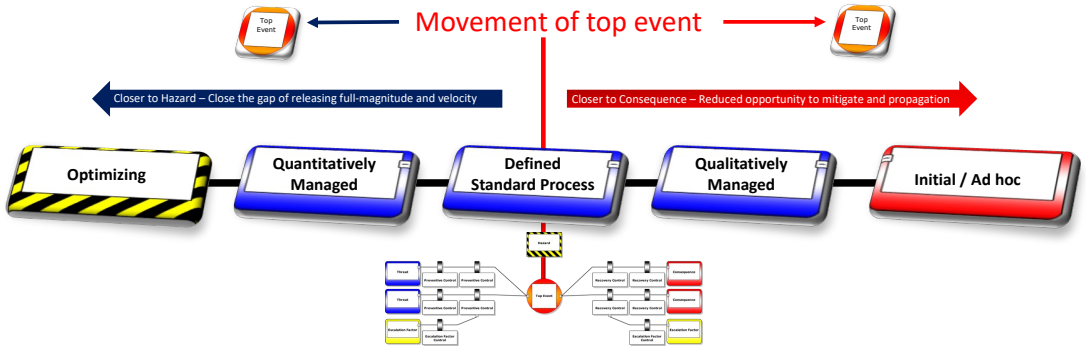

Figure 5: Shifting the top event as maturity grows

Bow Tie Families

Even with the correct top event, the life-cycle assessment may show that the control framework needs strengthening for the event's pre-conditions or downstream outcomes. It may then be necessary to build several bow ties for the same hazard, and a bow tie may extend into another dimension of risk — from a material hazard to a financial risk source. The analysis starts with a **parent** bow tie and is broken down into smaller, mature **child** bow ties. These move left and right along the event timeline, or across risk categories, but they can all be linked into a **bow tie family**. Good software secures top event families, control families and activity families, and links them to investigation patterns and layers of protection. Together they form a "cube" containing everything relevant to one hazard, from cradle to grave, highlighting control effectiveness and the risk owner's vulnerability.

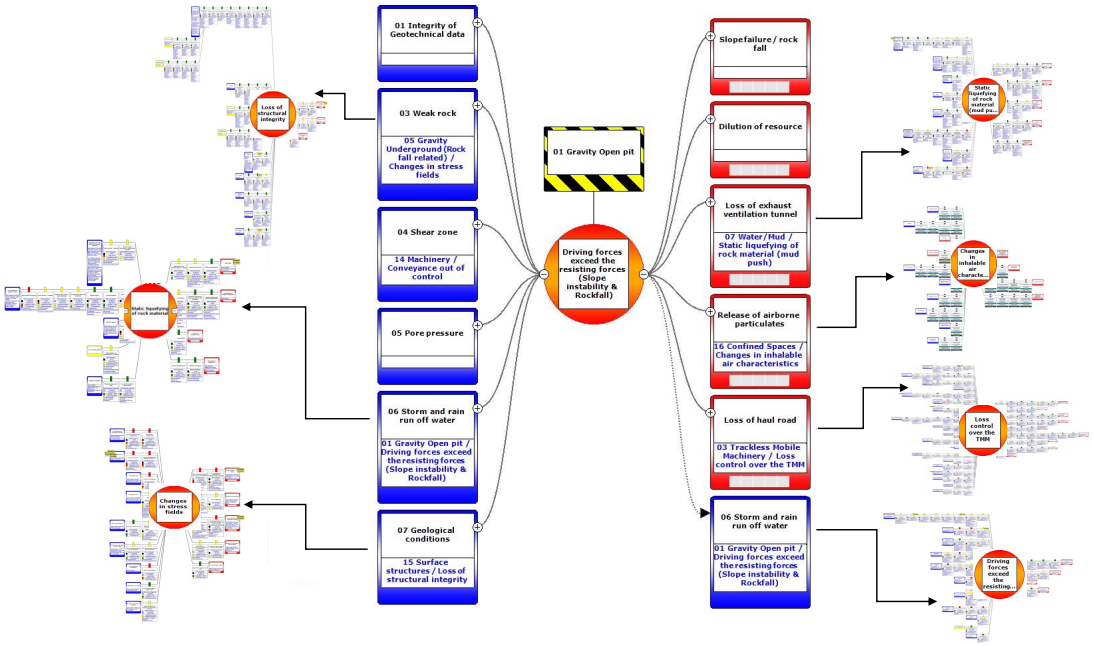

Figure 6: Bow tie family diagram

A single bow tie analyses one section of a longer causal chain — the crucial phase leading to an event. Usually that captures the critical factors, but some issues are so complex that we must look further back. Software can chain bow ties together to extend their reach.

Case study — The transformer explosion bow ties. A team of professional engineers developed a timeline of a transformer explosion. Building multiple bow ties from it revealed both the power of the method and how easily the top event can be misplaced.

- Transformer oil has three essential functions: it maintains electric fields, cools heated electrical components and impregnates the paper insulation.
- Transformer oil is a dielectric — a medium that transmits electric force without conduction. The point at which the oil is compromised, creating conditions for an internal arc, became the critical point on the timeline. Every control in a bow tie must be measurable; here the measure is dielectric strength, expressed through breakdown voltage — the oil's ability to withstand electrical stress. Standards such as IEC 60296 specify minimum breakdown voltages for new and treated oil.
- Oil-filled transformers contain two further dielectric materials: paper (on walls and windings) and, in some designs, varnish on the windings. The bow tie used all three during layer of protection analysis.
- Understanding dielectric strength directed the team to the causes of breakdown: operating temperature, increased internal pressure, moisture, contamination by particulate matter, and the formation of gases.
- Under conducive conditions, energy is released extremely rapidly as an electrical arc. The arc vaporises and ionises the oil; the resulting gas bubbles rapidly increase static pressure inside the tank until the tank structure fails violently. Oil mist is expelled under pressure and, if ignited, produces a massive explosion.
- High operating temperature contributes to the breakdown of dielectric material and is often associated with the explosion.
- Contaminants — moisture, impurities introduced during filling, bubbling during top filling, and products of arcing — were major contributors in several industry events.
- The gases released include combustible methane, hydrogen, ethane, ethylene and acetylene, all identified as part of the release mechanism in a previous transformer explosion.
- The arc, generating explosive gas mixtures released when the tank ruptures, effectively describes the initiation of the explosion.
- Finally, repeated elastic and plastic deformation of the tank's steel structure under pressure cycles, even without rupture, can degrade the tank's pressure rating over time.

From this description, the early stage of the event begins where *breakdown* occurs. The arc carries high thermal and radiant energy and affects the cooling properties, but it is better treated as an escalation factor — like the gas volume and oxygen space above the oil — than as the leading

cause. The formation of water, acids and sludge in the tank and windings occurs after the top event and is noted as an escalation of the impairment of cooling and insulation controls.

Despite guidance on the definition of a top event, the team initially chose “increase in temperature” as the top event. The analysis that followed raised questions about missing controls critical to preventing an explosion. The team reconvened, followed the guiding principles and tracked the event more carefully. The result was a sequence of top events — thermal energy, insulation breakdown and then oil — and a coherent control framework.



Figure 7: Transformer explosion timeline

The potential magnitude of an oil-filled transformer fire or explosion is beyond dispute, especially in a shaft or main intake airway, where it can mean multiple fatalities and the release of toxic gases. In the author’s view, no company should have an appetite for oil-filled transformers underground, whatever scientific justification is offered, because no control can adequately mitigate that consequence and the cost of replacement cannot outweigh a human life. Likelihood and consequence must not be blended: statements that the likelihood of multiple fatalities has been “reduced” by monitoring critical controls on the transformer confuse the two sides of the bow tie. Likelihood is controlled on the left of the top event, by preventive controls; consequence is controlled on the right, by mitigating and recovery controls. In the five bow ties that followed, the only effective mitigation identified was to reduce the magnitude of the hazard by removing the oil altogether — not containment, not fire barriers, not extinguishing media.

Set the Decision Criteria First

Before facilitating any bow tie, define the decision-making criteria. They ensure uniformity, prevent bias and limit subjectivity. Validate them, free of uncertainty, so that they reflect the company's risk appetite.

- **Preventive control acceptance criteria** define what counts as a control: an act of a person, a physical object or a technological system that prevents or monitors the release of the threat or the occurrence of the top event.
- **Mitigating control acceptance criteria** explain at what stage after the top event each control applies, the sequence of controls and their trigger points. They should acknowledge that the consequence cannot be eliminated.

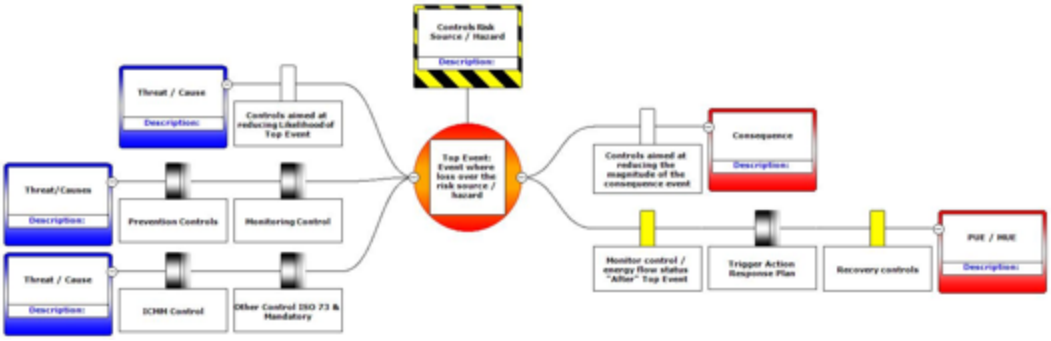

Figure 8: Positioning of controls in a bow tie

- **Control effectiveness criteria** set qualitative and quantitative specifications for reliability, availability, survivability, dependability and compatibility, with questions that define objective outcomes for failure rates, compliance and impediments.
- **Layer of protection criteria** add to effectiveness by defining the stability of each control within the whole framework. These quantitative criteria require substantial data supported by monitoring, research and statistical analysis.
- **Escalation factor criteria** cover four categories, each explaining why a control can never be fully effective: natural perils (for example rain or mist), critical services (for example diesel supply), inherent mechanical features (for example elastic deformation) and human physical or psychological constraints (for example circadian rhythm or impaired vision).
- **Control framework criteria** describe the hierarchy and layers of controls — supportive, dependent, verification and mandatory. They are linked to qualitative effectiveness criteria and should be used only until monitoring has produced enough data for the criteria to become quantitative.
- **Competence criteria** define the minimum expertise needed to secure effectiveness, extending beyond technical competence to risk competence and maturity.

- **Monitoring criteria** specify the activities within each control (for example calibration or bump testing), their frequency and the nature of verification. Note that calibration is not itself a control; it is a monitoring activity for a control such as a gas detector.
- **Auditing criteria** set designated questions for field verification and compliance surveys, using a tiered approach.

Threats: One Step Back on the Timeline

With the top event and criteria secured, identify the threats. Threats are the causes of the top event — one step back on the timeline. At what point did a pre-condition become a risk that could produce the top event? The word *cause* invites people to think of accident causes, including control failures and human error, which do not belong among the threats. Human error belongs among the escalation factors.

Threats lead directly to the top event, singly rather than in combination, and do not cause the consequence directly. The angle of a waste dump directly causes the slope to become unstable; the unstable slope loses cohesion and causes the consequences, in that sequence. If the top event requires two or more threats to occur together, the top event has been defined incorrectly. A fire, as a top event, needs at least three threats at once — fuel, heat and oxygen. Falls of ground and vehicle collisions likewise require several simultaneous threats, which disqualifies them as top events by definition and tends to produce "threats" such as "lack of support", "poor roadway conditions", "human error" and "failure to lock out".

Threats can be tangible or intangible. Select them carefully, with a clear explanation of their boundaries and trigger mechanisms, frequency of exposure and contribution to the top event. These details reveal commonality between threats and suggest where to combine or split them. Aggregation and segregation should be done by knowledgeable people who can see when a threat shifts back into a pre-condition or forward into an escalation factor. Rain, a natural peril, can be a pre-condition that becomes a threat — and it can also be an escalation factor on either side of the top event.

Consequences Create Scenarios

A top event propagates into multiple consequences. Production loss, injury, property damage and environmental impact are not consequences in themselves, but *outcomes* of consequences. A less mature bow tie has "unwanted fall of ground" as the top event and "fatalities" as the direct consequence. A mature bow tie has "driving forces exceed resisting forces" as the top event, with falls of ground as consequences that lead to discriminating outcomes.

Consequences create opportunity: they allow different scenarios to be constructed from the top event, even if the organisation has never experienced them, and so allow it to evaluate its readiness to respond. Simulation of scenario events is widely used to observe complex events virtually and identify additional or unverified consequences.

A good sign that a consequence needs more definition is when its controls become generic and repetitive. An emergency preparedness procedure that appears on multiple consequence pathways describes how to manage an accident scene in general, but may not explain how to recover a person trapped in a vehicle after a collision — including the specific step of deactivating the airbags before using hydraulic rescue tools to cut the vehicle.

The objective of the bow tie is to evaluate the effectiveness of controls for specific events selected after a risk assessment. Its purpose is to **prevent the top event and mitigate the consequences** — and preventing fatalities is the outcome that follows.

Barriers: Act, Object or Technological System

Risk management positions progressive barriers to manage a defined event. A barrier is any measure that acts against an adverse force or intention to maintain a desired state. Barrier effectiveness becomes the basis for decisions about opportunities, revealing the vulnerability of the risk owner.

First, classify barriers as those that manage likelihood (left of the top event) and those that mitigate consequences (right of the top event). This distinction separates prevention from response, and allows trigger points to be identified and event velocity to be monitored.

Each barrier has an objective related to the threat, top event or consequence. Objectives should be definite, specific, observable and measurable. By that definition, a procedure or an operator's competency is not a barrier. A barrier is an **act** a person performs, an **object**, or a **technological system**.

- **An act** is a hands-on action by a person using skills and knowledge, directed at the barrier's objective — for example, testing for flammable gas. Completing the checklist or calibrating the instrument is not the barrier; those support the effectiveness of the test. Acts include automatic, rule-based and taught performance. If every element related to one act (checklist, procedure, training, gas test) is listed as a separate barrier on the same pathway, it creates the illusion that the threat is well managed, when in fact there is only one act. Administrative functions are supportive and verification controls of the act, not acts themselves.
- **An object** is used during an act or operates independently, and may be linked to a monitoring system. Objects manipulate the flow of energy, can monitor change and can trigger alarms. Once installed, objects such as interlocks or earth leakage protection need no person to function.
- **A technological system** monitors barrier status, change or energy modification, either so that a person can act on the information or so that the system can act automatically at a set trigger. Big data and artificial intelligence can identify common error patterns and predict how effectively events will be managed.



Figure 9: Control families as defined by ICMM (green band) and ISO 73 and mandatory controls (grey band)

Supporting Controls and Effectiveness

Once barriers are classified, assign supporting controls. These comprise **dependency controls** — procedures, training, codes and design criteria that ensure the act is performed by a competent person, to specification, with the right tools in a controlled environment — and **verification controls** — checklists, maintenance reports, data analysis and competency assessments that verify the barrier's reliability, availability, survivability, dependability and compatibility.

Barrier effectiveness should be assessed formally against the criteria. Declaring a barrier "90 per cent effective" requires substantial data and must not be done lightly or subjectively. Bow tie software such as BowTieXP includes verification survey and audit modules that confirm effectiveness in the field, reflecting actual performance rather than opinion. The point at which the technique began to be used mainly to *select* critical controls, rather than to *evaluate* control effectiveness, is the point at which much of its value was set aside.

A mature bow tie shows the effectiveness of every barrier and its supporting controls against a norm that can be verified in the field against a performance specification. Field surveys provide real-time data that feed selection automatically. Barriers with similar objectives can be grouped into families with similar effectiveness. Validation covers the related activities and appraisals that ensure each act, object or system functions as designed.

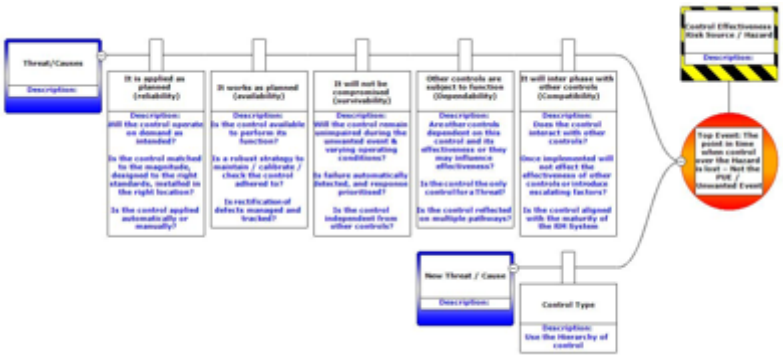

Figure 10: Control effectiveness criteria

A bump test of equipment to confirm that zero energy has been achieved is an activity connected to the test for zero potential; so is using an instrument to confirm zero potential in a hydraulic system. Appraisals define what effective performance looks like, and the most important verification questions concern systems, hardware, procedures and legal requirements, competency, maintenance and detection. List every part, item of equipment, substance and material involved in the barrier, because criticality may reside in a single component. In a spring-

applied brake system, for example, the critical element may not be the hydraulic pressure at all, but the setting that releases or applies the spring plate. Validation pinpoints that crucial part with reasonable accuracy.

Failure Rates and Escalation Factors

The other side of effectiveness is the failure rate, which monitoring data can substantiate. Failure rates show a barrier's limitations and sometimes reveal that it cannot reach its full potential because of escalation factors. Each barrier is assigned a level of protection — its independent capability to detect and prevent the top event or mitigate the consequence. Barriers that cannot maintain their control parameters are selected for escalation factor analysis.

Escalation factors fall into the four categories described above, and each needs additional controls. Avoid overlap between escalation factors and threats in the same bow tie. As a threat, rain reduces traction and causes a vehicle to skid on a gravel road. As an escalation factor, rain reduces the effectiveness of water management controls such as road camber and drainage. Escalation factors are, by their nature, opportunities to improve barriers by design, through action management or critical control management.

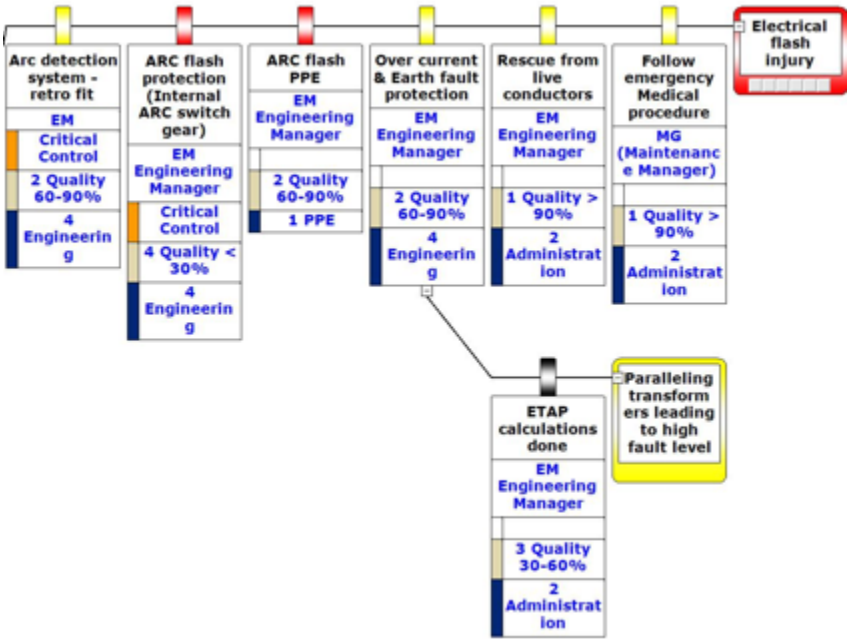

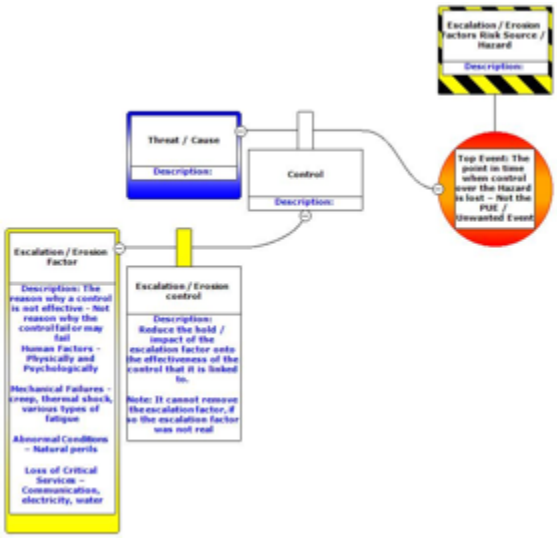

Figure 11: Escalation factor criteria and example

From Bow Tie to Residual Risk

With all barriers, supporting and verification controls in place, calculate the probability of failure and the maximum reasonable outcome. The calculation depends on layer of protection analysis and on mitigating controls functioning as intended. The result is a **vulnerability statement** and a description of residual risk requiring contingency plans. The risk assessment can then be updated with the residual risk description and ratings. This approach to residual risk avoids opinion-based statements and is specific to the sequence of events.

Some controls carry a distinctive characteristic: if they fail, other controls fail with them. Their breakdown radically changes the outcome, and the event propagates into a disaster at high velocity. A further benefit of the bow tie is that it highlights these controls. They are the **critical controls**: controls whose change carries a high, quantifiable magnitude, considerable velocity and immediate escalation.

Key Takeaways

- Return the bow tie to its full purpose: evaluating control effectiveness, not only selecting critical controls.
- Start with a life-cycle risk assessment and plot the hazard on an event timeline.
- Place the top event at the point where control of the energy is lost — close to the hazard.
- Threats act singly; if two must coincide, the top event is wrong.
- Barriers are acts, objects or technological systems; procedures and training support them.
- Set decision criteria before the session, and measure effectiveness in the field.

- Use parent and child bow ties to build a complete picture of each hazard.

In the next chapter, we take the critical controls identified by the bow tie and show how to select, specify and monitor them so that they deliver what they promise.

Chapter 12: Critical Control Management — Restoring the Window

> *Did you know that the method behind "critical control points" was developed in the 1960s, when the Pillsbury Company worked with NASA to make food safe enough to send into space? A single contaminated meal could end a mission, so the team stopped testing finished products and started controlling the few points in the process where hazards could enter. Mining borrowed that idea. The question is whether we borrowed all of it.*

The Romans were among the first to use glass in windows. Two thousand years later, glass still serves the same purpose: it controls the exchange of light, air and weather, while its colour and shape shape the character of a building. But a building's appearance can deteriorate quickly unless it is maintained. In 1982, James Q. Wilson and George Kelling described how a single unrepaired broken window signals that nobody is caring for a building — and invites more damage.

Critical controls arrived in mining in much the same way: first proven in other industries, then refined with guidelines and well-presented implementation strategies, offering a panoramic view through what seemed a small window. The opportunity they presented for drastically reducing fatalities was enormous. In some organisations, that window was soon covered with the graffiti of perception. This chapter is about restoring it.

What Critical Controls Were Meant to Be

The food industry learned, through events such as glass fragments discovered in baby food, that a single failure could close an internationally respected company within days. Hazard Analysis and Critical Control Points (HACCP) dramatically reduced such events, although they still occur.

The "graffiti" version of critical control management holds that a few controls, alone or in combination, will prevent accidents. That is only true if the concept is understood. The original objective was to verify that **all** the controls designed to manage an event's likelihood, or mitigate its impact, were actually in place — all of them, not a selected few. Next, the effectiveness of each control, and its interaction with others, was to be determined, and action taken to reduce vulnerability. Only after that holistic view would the critical controls be selected. At that point, a mature system can show quantitatively which controls have the greatest influence on the outcome: those that can prevent the event or, if that is not possible, reduce the consequence to an acceptable level.

Every link in a chain is critical. If the chain lifts an eight-tonne ladle of molten metal, some links may deserve more attention than others — perhaps the link closest to the heat, the one connected to the shackle, or the one in the middle. That decision must not rest on perception. The lessons

on human error and groupthink in Part III apply directly. The layers-of-protection data are the frame of the window; a window that does not fit its frame cannot do its job. One way or another, the industry will move towards data and artificial intelligence to make these decisions.

Monitor the Challenge, Not Just the Standard

Changing control definitions, designs and surveillance so that data points measure the performance of a complete event control system is challenging, but the previous chapters show what to monitor. Control performance standards and verification should specify how data and artificial intelligence will track each control. Rather than setting reliability simply to meet a legislative or other standard, focus on **how often the trigger event challenged the control, and how often the control functioned fully, partially or not at all**.

> **Case study — From padlocks to smart locks.** Energy isolation with lock-out and tag-out has been used in industry for more than three decades. Competent, rule-based artisans follow the procedure, and it is supported by detection practices and a well-documented permit system. Yet every year, serious incidents involving mechanical or electrical equipment still occur, and in largely rule-based systems, violations tend to be followed by severe injury. If the equipment, the locks and the tags stay the same, behaviour will stay the same, and future events will follow.
>
> To monitor even one aspect of isolation — the lock — organisations can adopt smart (intelligent) locks. These pre-programmed devices communicate with a control panel, through cellular or other networks, and inform the user and supervisor of any malpractice or deviation from the programmed norm. They can be cross-referenced with each machine's isolation plan, access requests, anomalies, energy type, lock mechanism and the validation of energy status. Research into digital lock-out tag-out systems with machine learning (Chen, 2022) shows how quickly the field is advancing.

Accepting electronic systems that gather structured data for artificial intelligence requires stable, secure monitoring infrastructure. Bluetooth, cellular networks, leaky feeders and Wi-Fi each carry their own risks, which must be managed — but together they offer a window of opportunity to move beyond paper-driven systems.

Selecting Critical Controls: Seven Steps

Which control is more critical? Selection should follow defined steps and not depend on any individual's perception, experience or qualifications.

- **Step 1 — Complete the risk assessment** using a technique that illustrates each control's associated risks within the life cycle and along the event timeline. Identify the critical point at which each control must act to be effective. Without these points, decisions become random and priorities misleading. Well-designed software can simulate events and controls.

- **Step 2 — Define each control's objective, function and dependencies** on other controls. Activities such as calibration that maintain control status enrich the decision.
- **Step 3 — Assign control families.** Combine interdependent controls into families, with the option of selecting individual controls where appropriate. Some may sit at escalation level and strongly affect the upper level of control.
- **Step 4 — Establish the reliability of each layer of protection.** This is the single most important element of selection. It determines what is needed to achieve an acceptable, quantitative risk level for each scenario, evaluating causes and outcomes and specifying protective measures. Use sufficient data to decide whether a control will function on demand and produce its intended outcome. Define the point of failure of each layer (each family of controls) and the trigger that activates the next layer. The criteria for each layer become the basis for monitoring control effectiveness.
- **Step 5 — Select controls that are engineered** in design and application, have an electronic monitoring system, or can support artificial intelligence control regimes.
- **Step 6 — Define the monitoring structure:** sensor placement and type to ensure accurate surveillance, wanted, discarded and pre-emptive engagements, and the response plan for malfunction. Monitoring is interlinked with the layers of protection, their failure and the trigger of the next layer. Mature systems automate these activities and prevent human interference.
- **Step 7 — Revisit the vulnerability analysis,** using back-analysis to confirm that current measurements correlate with forecast design data and parameters.

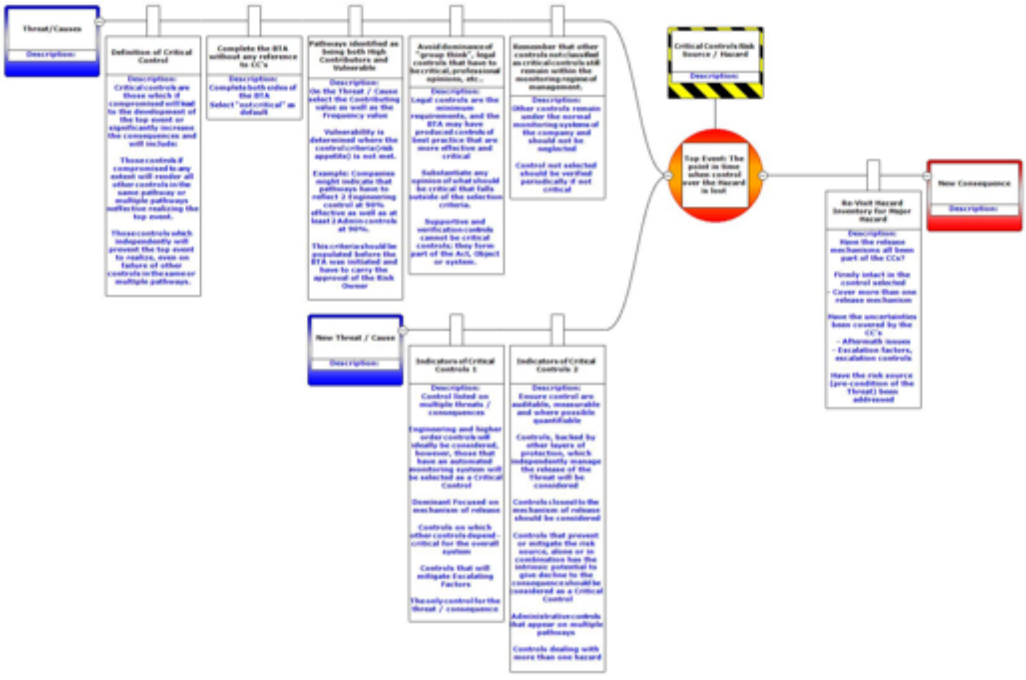

Figure 12: Critical control selection criteria

Guiding the Selection

Experience shows that selection can drift. The facilitator should guide the team in ten areas to avoid bias and discrimination between controls.

- **Groupthink.** Selection can be steered by a group towards the preferred outcome of one influential person. Beauty is in the eye of the beholder, and so are critical controls when there are no guiding policies. Consensus should rest on reasoning, evaluation and consequence analysis. Cohesive groups of specialists who have worked together for many years in the same environment are especially susceptible, which is why independent facilitation adds value.
- **A priori reasoning.** Facilitators must distinguish an a priori *assumption* from an a priori *truth*. Some controls, presumed by experts to be critical, have little effect on the outcome on closer scrutiny; spring-applied, hydraulic-release brakes on trackless mobile machines can be an example. Others are known, without the need for further evidence, to prevent an event on their own — physically disconnecting a power source by withdrawing the cable connection from the switchgear, for instance.
- **Human acts.** The act-object-system approach applies, but an act as a critical control has nothing to do with countering behaviour. It is a control *within* a task, not a task. Testing for zero potential on a capacitor is not a critical task; it is a critical control within the larger task of de-energising the capacitor.
- **Clustering.** There is a fine line between breaking a control into fragments and clustering fragments into larger units. Start with fragmentation, then cluster where objectives overlap, and show the fragmented options before selection — the critical control may be one small fragment rather than the whole family. Fire suppression on a trackless machine is a good example: the critical discussion should cover design, nozzle placement, extinguishing medium and quantity, feed-line routing, activation and the entire life cycle of the fire response.
- **Terminology.** Use language that the operator responsible for implementation understands. Technical terms such as “driving forces exceed resisting forces” are correct for the bow tie’s top event, but at the workplace “loss of cohesion” may communicate better and lead operators straight to the preventive critical control. The facilitator’s role is to translate without losing integrity.
- **Administrative controls.** An administrative control selected as critical is like a barking dog: noticeable, but it will not stop a determined intruder. A pre-start warning siren on a conveyor is valuable, but it is not a critical control.
- **Behavioural controls.** Administrative and PPE controls depend heavily on behaviour, and good practice is to exclude them as critical controls — usually after considerable debate. An arc-flash suit for electrical switching is a common example. The bow tie should show that the control is inadequate and that remote switching is the better solution.

Selecting a critical control to accommodate an unwillingness to improve is never good practice.

- **Habitual decisions.** Habits that become critical controls are a modern operational hazard. Safe declarations and permits to work can become documents that allow work to continue whether or not the tick marks are correct. Over the years, some operators earn a "PhD in tick-and-flick". Facilitators should identify these habitual controls and remove them from the decision criteria.
- **Mitigation-heavy frameworks.** The principal objective is to prevent the event, yet some bow ties carry many critical controls on the mitigating side and few on the preventive side. Self-contained self-rescuers and refuge bays underground are essential — but they are mitigation. Much capital goes to conveyor fire suppression, rescue equipment and facilities, all of which sit to the right of the top event. Balance the investment.
- **Earned, not bought.** Critical controls cannot be bought; they must be earned through evidence, design and verification.

> **Groundbreaking insight:** Critical controls cannot be bought; they must be earned. A control earns its place by evidence that it changes the outcome — not by its cost, its visibility or its popularity.

Specifying and Verifying

Managing critical control monitoring requires an electronic, collaborative system with iterative capability. Best practice combines risk assessment, monitoring, analysis and communication in one package. Professional bow tie software offers layers of protection, quantified risk levels, monitoring surveys, live automated updates and action management.

Selection is followed by control specifications, performance standards and verification. The information for this phase should have been developed during the bow tie analysis. Each control's functionality defines its owners' roles and responsibilities within the overall control framework. Specifications cover:

- **Competency** of the people who perform activities governed by the control, including instrument calibration, non-destructive testing, back-analysis, observations and settings.
- **Integrity principles** that keep the control unimpaired and operational on demand, including design principles, certification by inspection authorities and quality assurance.
- **Performance measurements** such as vibration limits, supplier-conducted primary and secondary brake response tests, and brake temperature limits with interlocks.
- **Response or recovery plans** that trigger the next layer of control — for example, automatic transfer from a failed primary brake to the secondary brake system and then to retardation and a controlled stop.

- **Key performance indicators** covering the period before the control is introduced, during its use and before it is retired, aligned with roles, responsibilities and risk appetite, and providing a tangible measure of the control framework's effectiveness.
- **Quantitative verification** based on those indicators: in-situ monitoring, deterioration, acceleration and deceleration of risk velocity, escalation rates, control distortion, sensitivity and deviation impact. A yes-or-no checklist generates paperwork and a database of little value, and obscures the true reliability of the critical control programme.

Before implementation, a test and validation phase follows. Risk modelling and simulation within defined parameters measure realistic control response, identify gaps, optimise the framework and close out lessons learned before any worker depends on the control.

Key Takeaways

- Verify all controls first; select critical controls only after the holistic view.
- Monitor how often the trigger challenged each control and how it performed.
- Follow a seven-step selection process anchored in layers of protection.
- Prefer engineered, electronically monitored controls capable of supporting data and AI.
- Guard against groupthink, habit, behavioural controls and mitigation-heavy frameworks.
- Replace yes-or-no checklists with quantitative verification.

In the next chapter, we look beyond the controls to the response itself — how to be ready before the event, and how to prove it.

Chapter 13: Risk Response Plans — Ready Before the Event

> *Imagine a dust cloud rising from an open pit and drifting across the road into town. Every person who sees it forms an opinion: a blast, a building collapsing into the pit, people trapped. Inside the mine, the control room is waiting for information. What happens in the next ten minutes was decided months ago — in the response plan, or in its absence.*

Many risk assessments stop once residual risk has been identified. Having ranked the risk once or twice — without and with controls — the team concludes that it has a number suitable for the risk profile. That approach quietly silences response and contingency planning.

Take one step back. Risk evaluation compares the results of risk analysis with the risk criteria to determine whether a risk is acceptable at a defined magnitude. Within its risk appetite, the business decides whether it is ready to bear the risk after treatment. Vulnerability, acceptance and tolerance can therefore only be determined once response and contingency strategies have been considered.

Response Planning Is Iterative

Response planning is a process of deduction and iterative measurement. The risk event is evaluated against the criteria, with its mitigating controls, to determine acceptance or tolerance. If the requirements are not met, the procedure is repeated until it succeeds or moves closer to the desired objective. This repeated effort to reduce magnitude has pushed the mining industry into research programmes and technologically advanced controls — one of its great strengths. Response plans cannot remain static or hibernate. In best practice, the monitoring and review function produces a response plan for each contributing factor of the event, supporting three levels of response.

- **Response to the trigger event:** terminating the release, or reducing the event's velocity and magnitude, so managing both likelihood and propagation. The plan evaluates the first trigger point, as described in the trigger action response plan (TARP), and measures the capacity of the control to manage the event. On an active open pit slope, for example, the slightest movement recorded at a survey prism should initiate a documented response with defined actions and responsibilities.

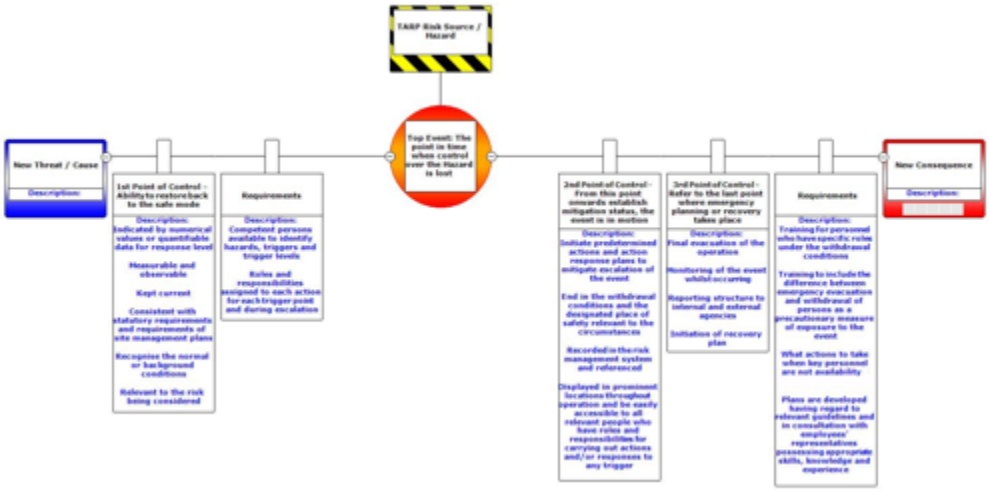

Figure 13: Response plan within the bow tie process

- **Response for each control,** in its sequence of application along the event timeline. Like a TARP, it defines the responsibility and authority for activating each control. In the slope example, the plan has multiple trigger points with defined actions before a slope failure — and also actions during and after the failure.
- **Response to the event if controls fail** under normal or abnormal conditions. Here the plan extends into emergency, recovery and contingency planning. A multiple-bench failure under abnormal conditions may include a mudslide, an inrush of water and loss of access to the ramp. Propagation scenarios test the effectiveness and resilience of the programme.

Contingency Plans

Contingency plans are backup plans for the surprises that arise during and after a risk event. They set out the sequence of actions that mitigate impact and restore the previous status quo. Best practice includes every scenario created during the significant hazard risk assessment and bow tie analysis, each with its own response and contingency plan. Every plan should share four features:

- **Continual measurement of effectiveness, efficacy and efficiency,** to streamline response actions, remove bottlenecks, avoid duplication, improve resource deployment and evaluate, through empirical evidence, the factors that influence outcomes — internal and external capacity, risk knowledge, changing context, implementation reliability, and resource availability and access.
- **Consistent training and education** for everyone with a role in the plan. TARPs are integral, and interim changes affect reliability and survivability.
- **Regular, proactive drills and exercises** that include internal and external parties. Engagement with local communities requires care and an understanding of consequence and impact; sensitive information can spark rumours before an event, and the plan must address communication accordingly.

- **Continual maintenance,** so that the plan remains aligned with operational demand, readiness for change and the liquidity of the business unit.

Resilience Proven, Not Assumed

An organisation's ability to bounce back to "business as yesterday" after a disturbing event reflects its resilience. Significant hazards in mining are associated with multiple fatalities, and such events profoundly affect the industry, local communities, markets and everyone who works in mining. Resilience has made the sector remarkably strong and adaptable over time; even with generic plans, it has repeatedly returned to normal after disruption. Specific, tested plans will make that resilience faster and more certain.

Assurance auditors verify the existence of response and contingency plans from evidence, but in the absence of actual events they may be confirming plans that have never been exercised. High-risk events should link directly to specific response plans, and a simulation programme should demonstrate the ability to respond. A risk assessment captured in a spreadsheet or a static library cannot, on its own, provide that evidence of due care. Mature operations have moved to real-time simulation, virtual interactive training and immersive environments that produce highly efficient response teams.

> **Case study — Lily Mine, February 2016.** At Lily gold mine in Mpumalanga, a crown pillar collapsed and a surface container used as a lamp room fell into the resulting sinkhole, with three workers inside. Despite a sustained rescue effort, they could not be recovered. Dozens of other workers underground were brought to surface safely. The event showed, in the most painful way, why response planning must extend beyond the first consequence — to ground stability during rescue, access to the incident site, communication with families and the community, and the long-term recovery of the operation. Many of those post-consequence controls can only be discovered by extending the bow tie past the consequence, before the event.

Finding the Trigger Points

Where do we find trigger points, and how do we define the actions? The opportunity lies in the company's bow tie programme. Linking parent and child bow ties allows movement through the three-dimensional arena of risk. A bow tie's top event produces consequences, each with a suite of mitigation controls, and part of that mitigation is the initiation of the response plan. What is usually missing is the rest of the response: recovery and contingency controls.

To discover them, shift the top event to the right along the release timeline, past the consequence, and develop a new bow tie with the consequence as its top event. This unlocks direct and latent impacts. In most organisations, populating these controls has been an eye-opener: they are often missing or generic.



Figure 14: Control and response points in a bow tie

The extended bow tie needs a particular format for post-consequence impacts. Return to the pit wall failure that sends a dust cloud towards the access road to town. Every person who travels through it or sees it forms a different opinion of what happened. Inside the mine, perceptions may also differ — did the crown pillar fail? Should we wait for the control room operator to respond? These scenarios become the centre of the impact assessment for the new bow tie, and the extended bow tie reveals the response controls that must follow the consequence.

Response controls for each post-event impact are directed at the magnitude and severity of the event. Escalation factors may appear, and this is where contingency plans must be well defined and tested within the bow tie control framework. Response and contingency plans differ in gravity and will not all activate at the start of the event. Roles and responsibilities for each critical point of each critical control must be defined and tested, and drills must be simulated to ensure chronological, timely and correct execution during a real event.

> **Groundbreaking insight:** A response plan that has never been simulated is a hypothesis. Extend every major bow tie past its consequence, populate the response controls, and test them — before the event tests them for you.

Key Takeaways

- Residual risk cannot be judged acceptable until response and contingency plans are considered.
- Build response at three levels: the trigger event, each control, and failure of controls.
- Every contingency plan needs measurement, training, drills and maintenance.
- Prove resilience through simulation; existence of a plan is not evidence of readiness.
- Extend bow ties past the consequence to discover recovery and contingency controls.

In Part V, we move from controls to culture — beginning with the relationship between trust, uncertainty and risk appetite.

PART V: Culture, Framework and the Future

Controls work only inside a culture that sustains them. Part V connects trust, uncertainty and risk appetite to a robust risk management framework, and closes with a portrait of the risk professional the industry will need in the decades ahead.

Chapter 14: Appetite for Risk Culture

> *Imagine refuelling a vehicle with an 80-litre tank. Some drivers wait for the warning light; others refuel at the quarter mark. Nobody wrote them a procedure. Each acts on an inner sense of what is acceptable, triggered by a measurable signal. That is risk appetite in its simplest form — and understanding it is the key to understanding risk culture.*

Trust grows as culture evolves; uncertainty and opacity grow as culture regresses. The gap between them is an opportunity that can be measured: an increase in trust, a decrease in uncertainty. To understand risk appetite, we must examine the traits that shape our perception and our trust when we decide whether to pursue or retreat from risk.

Three Kinds of Uncertainty

Uncertainty is shaped by information, and it comes in three forms.

- **Natural uncertainty** is the inherent variability of information. More measurements, smaller margins of error or more sophisticated equipment cannot reduce it. In practice, critical controls are measured to reflect the resilience of the risk management system and its position on a maturity scale. Counting answers to critical control questions produces a *compliance* status, which does not measure the quality of the control or the system behind it. Natural uncertainty calls for both quantitative and qualitative information. The qualitative element sets the level of uncertainty that affects a person's appetite: it alerts the mind to how trustworthy the information is, and so to hesitation before acting on risk.
- **Data-induced uncertainty** arises from data selection, cleanliness, variables and transformation. Modelling data to support decisions requires computational methods with estimated parameters and predictions, and modelling produces the information used for critical controls and TARPs. Inaccurate input gives inaccurate output. Accurate output costs money, takes time and becomes more complex as the margin of error shrinks. Algorithms may produce outcomes that do not match an expert's expectations, so artificial intelligence and the Internet of Things should be introduced with care — and with explanation, so that users understand the decisions being made on their behalf.
- **Randomness** arises when information is independent and its accuracy cannot be quantified, so the user accepts its imprecision. Managers who verify critical controls through walkabouts, inspections and observations, record the outcome and wait for the dashboard to show compliance are working with randomness. The uncertainty attached to small, random samples is unquantified.

In summary, all the information we use to determine risk levels, critical control effectiveness and maturity — everything that influences risk appetite — carries uncertainty. The more accurate the

information we share with the people who must respond, the more trust grows. Missing, unreliable, conflicting and confusing information is the framework of a regressing risk culture; accurate, timely and honest information is the framework of a growing one.

Trust as the Foundation of Resilience

In risk culture, resilience means addressing the sources of threats and triggers while strengthening a system's capabilities and resources to survive risk events and continue to function. Rising trust runs parallel to falling uncertainty, and together they form the foundation of a secure risk culture. Trust has its own uncertainty and depends on the same information variables.

Trust is part of the social makeup of the workplace, with a few key drivers: fear, consequence, embarrassment and retribution. Many efforts to improve safety by "increasing trust" have had limited effect, largely because our knowledge of the subject is limited and we have overestimated our understanding of the minds behind safety and risk culture.

Think of the trust a patient places in a surgeon without ever having experienced a medical procedure. Most people fear surgery, beyond any uncertainty about success. The more critical the operation, the greater the uncertainty and the more the patient must trust the surgeon. Only after waking from the anaesthetic does the patient know the procedure succeeded, and some uncertainty remains until full recovery.

Now reverse the roles. In the workplace, **you are the surgeon**, not the patient. Your colleagues trust you, and your decisions, made on uncertain information, affect them — under intense production pressure. Your risk appetite, your maturity in taking or avoiding risk, and your capability to act reflect the risk culture of the whole company, not only your own. This shift requires an understanding of how people react under different circumstances.

Building Trust That Lasts

Relationship building, leadership and accountability remain important, but in today's demanding mining environment they need new depth. The building blocks of trust are a person's actions. Those actions arise from instinct, emotion and habit, and they produce conduct that shapes how others perceive us. Criminologists and psychologists note that emotions, like the body, must be fed. Understanding employees' emotions offers a window into what they experience during the day. In a tightly controlled environment, any step to the left or right has consequences. Feeding emotions positively is a norm of social culture, and it shapes workplace risk culture.

Safety campaigns and behaviour-based programmes have raised awareness across the industry. To build trust that can be measured in the context of risk, the controls used to build it should meet the same standard as critical controls: observable, specific and measurable.

Risk culture will remain a lively topic as long as the industry believes that trust can be controlled. Trust is better understood as an equilibrium — one that is influenced by factors beyond the authority of any single leader, and that must be continually earned.

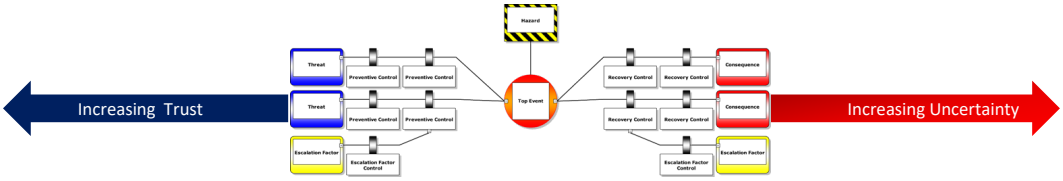

Figure 15: Bow tie analysis of trust and uncertainty

If trust is earned through consistent effort, it must be visibly reflected in a person or a group. Trust is dynamic: it is disturbed every time a person or group changes, and the change can be as ordinary as transferring a manager or supervisor to another section. A shift in team dynamics can increase risk-taking, introduce new habits and alter the group's risk appetite. Planning for trust during change is therefore a genuine risk control.

Appetite, Tolerance and Capacity

Understanding risk appetite gives insight into risk culture and into the industry's continuous improvement journey. Consider shaft winders and hoisting. Stringent rules, checks and balances are built into their monitoring and action management systems. Yet significant events still occur — in South Africa and worldwide, in almost every commodity and method, in vertical, inclined and declined shafts, in shallow and deep mines. Only a few result in fatalities, but each has repercussions throughout the industry and the public domain. Useful questions follow. Do engineers and production managers have different risk appetites? What is the current state of risk appetite? What should a best practice framework contain? What does the future prototype look like?

Risk appetite is shaped by two factors, **risk tolerance** and **risk capacity**, and the space between them is the **mitigation gap**.

- **Risk appetite** is the type, amount and level of risk that a company, a team of engineers or an individual is willing to take. It defines who you are and who you want to be. It can be expressed in clear statements that support the achievement of business objectives and are agreed by senior management. Ultimately, risk appetite is a position adopted by senior management in pursuit of its objectives, and it should be presented to the organisation as their position, not as a mandate from the risk manager. There is no one-size-fits-all expression of appetite. Frameworks help, but each organisation must chart its own path in line with its tolerance, deciding how formal, detailed and mature its statement should be.

Key considerations are how acceptable levels are set, the value placed on the related objectives and the trigger point for response. With tangible, material risks, the picture can become clouded.

If a conscious decision is taken to continue hoisting with a rope showing 10 per cent deterioration and an incident follows, a court may regard the decision as risky — even if it was within the law. Like the driver who refuels at the quarter mark, a person or team connects a risk level directly to a measurable item — 10 per cent, or 20 litres — and responds according to their inherent understanding of what is acceptable.

- **Risk capacity** is the threshold. Winding speeds and accelerations vary greatly between winders. One winder was designed to operate safely at 20 metres per second, yet raising and lowering personnel was limited to 15 metres per second. Many variables are considered in design to determine the actual speed for a particular shaft, supported by extensive dynamic analysis and tests and by the legal requirements at the time of design. That design threshold is the risk capacity.
- **The mitigation gap** lies between risk appetite (the trigger response level) and risk capacity (the threshold). Within this gap, upper and lower limits can be defined and used as key performance indicators for meeting objectives.

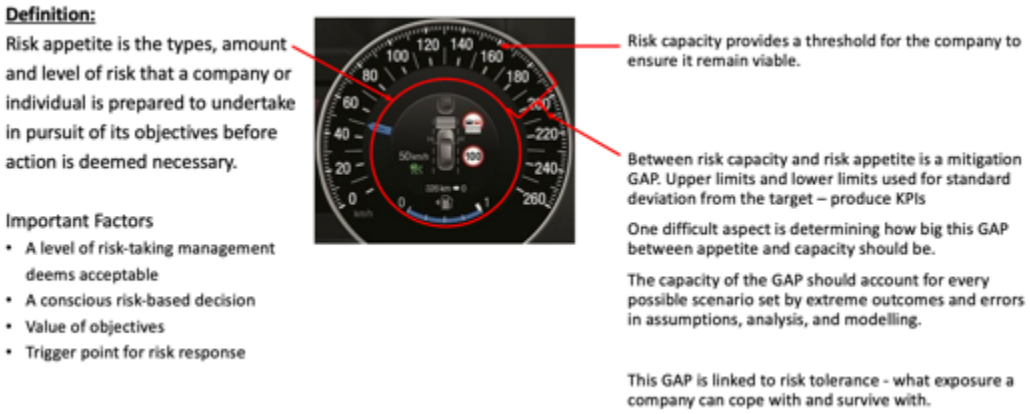

Figure 16: Summary of risk appetite

> **Groundbreaking insight:** Every material risk can be expressed as an appetite (the trigger), a capacity (the threshold) and a mitigation gap between them. Define the gap, set its upper and lower limits, and you have turned risk appetite from a boardroom statement into a shop-floor key performance indicator.

From Safety Management to Risk Management

Safety management and risk management have been integrated for many years, and the difference is worth clarifying. Safety management is, by design, a compliance practice: a law, standard or guideline sets the requirement, compliance is measured and non-compliance is corrected. Risk management manages risk within the parameters of the company's risk appetite, risk tolerance and risk capacity. Both are essential.

The next generation of risk managers and specialists should be able to develop risk tolerance statements, risk capacity philosophies and risk response plans within the boundaries of the

company's risk appetite. In practice, after a baseline risk assessment, each risk is connected to a defined tolerance (upper and lower limits), a trigger response that activates mitigation at the right time and, where necessary, initiation points for response and contingency plans. Current mining legislation, being largely prescriptive, does not yet allow full implementation of this approach, and managers are constrained in pursuing risk in several ways. Pedestrian detection and retardation systems, for example, are mandated as safety management rather than selected through risk management. Evolving legislation towards risk-based outcomes would unlock significant value.

These challenges reflect the complexity of risk management and the importance of well-designed risk appetite initiatives. Risk culture management measures internal performance and compares it externally. Appetite, tolerance and response are measured through an information network spanning operational levels. Although risk appetite is often seen as the driver of risk culture, the two are distinct. **Risk culture is a human matter**: it describes how people act or react in relation to risk. Measuring it means considering what influences people's behaviour towards risk — the tone from the top, how managers perceive and respond to risk, how employees act in turn, the resources available for risk, transparency and positive recognition.

Maturity and Culture Are Different Measures

Risk maturity and risk culture are often used interchangeably, but they measure different things. **Risk maturity** describes a person's or organisation's capability to perceive, manage and respond to risk, and its readiness to pursue risk for continuous improvement. Many maturity models use a staircase. A staircase has a bottom step, but it should have no top step: a model that declares "resilient" as the final level suggests arrival, which contradicts continuous improvement. Maturity models are excellent for benchmarking between departments and operations within a company, but less suited to comparing companies with different appetites and cultures.

Some models attempt to combine maturity and culture, often adapted from the financial sector. Any integrated model must measure the finer detail of appetite, maturity and culture. Superficial perception surveys, undirected psychometric tests, 360-degree ratings and similar questionnaires can drift from the objective. Each trait needs a proper breakdown: for appetite, the criteria used to set the appetite statement; for maturity, the capability and methodology to identify, understand and manage risk. That detail provides a framework of measures that reveals system opportunities and system regression, supporting continuous improvement.

Risk maturity and culture are specific to the nature of the operation, its management style and its environment, and risk professionals should develop the measures. Each operation deserves a model aligned with its own capabilities. Risk, in all its dimensions, does not lend itself to generic approaches — and the willingness to design a tailored model is itself a sign of maturity.

Key Takeaways

- Distinguish natural, data-induced and random uncertainty in the information behind decisions.

- Trust rises as uncertainty falls; share accurate information with those who must respond.
- In the workplace, you are the surgeon — your appetite shapes others' safety.
- Express material risks as appetite (trigger), capacity (threshold) and the mitigation gap.
- Safety management ensures compliance; risk management operates within appetite, tolerance and capacity.
- Measure maturity and culture separately, with models tailored to each operation.

In the next chapter, we bring appetite, culture and controls together into a robust risk management framework.

Chapter 15: A Robust Risk Management Framework

> *"A bad system will beat a good person every time." — attributed to W. Edwards Deming*

A good risk register informs corporate decisions, and a good risk register begins with a risk management framework. ISO 31000 sets out eight principles that describe effective and efficient risk management and provide the foundation for managing risk. Expressed as what risk management should deliver, they are:

- **Integrated:** risk management is part of all organisational activities, not an add-on.
- **Structured and comprehensive:** a consistent methodology for each stage of the process produces comparable results.
- **Customised:** the framework and processes are proportionate to the scale, scope and context of the organisation.
- **Inclusive:** stakeholders participate appropriately and in good time.
- **Dynamic:** the system anticipates, detects, acknowledges and responds to change.
- **Best available information:** decisions use the best information available and openly flag its limitations.
- **Human and cultural factors:** the human and cultural influences on every risk decision, control and response are recognised.
- **Continual improvement:** well-defined performance measures and lessons learned drive improvement.

An effective framework becomes a culture. It defines capabilities and practices, and integrates strategic decisions to create, preserve and realise value. Its scope is broader than internal control: it includes strategy setting, governance, stakeholder engagement and performance measurement.

Three Pillars

At executive level, the framework rests on three pillars.

- **Risk management structure** — the basis of the management approach. It establishes the structures that manage the framework, inform and engage stakeholders at every level, and provide assurance and effective reporting.
- **Risk management strategy** — the methodologies, aligned with the company's risk appetite. Philosophies, arrangements for embedding risk management and policies are translated into objectives, with result indicators for benchmarking and for reporting to the board.

- **Risk management systems** — where implementation happens: the designated techniques, protocols, measurements and verification that precede reporting.

Executives and senior operational managers play different roles. Executives approve strategic decisions, set boundaries and oversee execution; they define risk parameters, categories and limits. Senior operational managers align strategy, resources, processes, people, reporting and technology to achieve objectives in line with established values. The Institute of Internal Auditors' Three Lines Model complements this picture by clarifying the roles of operational management, risk and compliance functions, and independent assurance (Douglas and Anderson, 2015).

The framework instils risk discipline in three ways: through risk management philosophy, risk appetite and the control environment.

Risk Management Philosophy

Risk management philosophy is the collective set of beliefs and principles of executives and stakeholders that shapes risk culture. It guides how the company considers risk — from identification to mitigation, from strategy to daily activity. Many companies treat it as a mission statement, but it needs more substance than a declaration: it must carry value through "the way we do business".

A holistic approach includes the board and executives, who apply the framework just as rigorously as anyone else. Any change to normal operations, whether an opportunity or a risk, should follow the same principles as those applied by an employee cleaning spillage under a conveyor with a shovel. Significant corporate changes such as restructuring deserve a formal, systematic risk analysis using the same principles. Such high-level risk assessments are still too often carried out only after the impact has been felt — in rising stress-related illness, post-traumatic stress and the human cost that follows poorly planned change.

The shift-cycle decision described in Chapter 4 is a telling example. When senior management decides first and the risk assessment is then guided towards the intended outcome, stakeholder interests are set aside, and with them the company's values: the travelling time to and from work at night, limited time with family, the risk of hijacking, stray animals on national roads and much more. In one case, when the objectives were not met, the company restructured again a year later, and only then experienced the full impact of the original change.

A philosophy cannot simply be assigned to an operation. It must be embedded until it becomes culture — from the top down, from the boardroom into the hands and minds of the workforce.

Risk Appetite

Risk appetite reflects the risk management philosophy and in turn shapes risk culture and operating style. It is usually set out in a risk appetite statement that frames the risks the company should accept, the risks it should avoid, and the strategic, financial and operating parameters within which it should operate (Rittenberg and Martens, 2012). Appetite is a delicate balance:

flexible enough to respond to changes in the business environment, but not so flexible that it is constantly shifting. Frequent changes erode value.

The Control Environment

The COSO Internal Control framework identifies five principles of the control environment that support strong risk governance:

- a commitment to integrity and ethical values;
- board independence and oversight;
- management's establishment of structures, reporting lines and authorities;
- a commitment to competence; and
- accountability.

A chain is as strong as its weakest link, and a business is as good as its people. These five principles make that truth operational.

Integrating Financial and Material Risk

Applying a robust framework strengthens the impact of governance on risk culture and, ultimately, on the achievement of business objectives. The mining industry has many well-structured policies and procedures with excellent intent. The next step is to integrate the framework across all levels, departments and functions. Today, risk philosophies, appetite and internal control often differ between financial and material risk decisions; the matrices differ, and financial decisions frequently take precedence.

This misalignment is understandable. Material risk specialists have often had limited exposure to financial risk principles, and the two systems have developed separately. Stakeholders may be more comfortable appointing a financial manager as risk manager than appointing a risk specialist as financial manager. The robust framework needs both disciplines, with distinct roles and responsibilities, working within one integrated process.

> **Groundbreaking insight:** The most powerful risk framework in mining will be the first one that gives financial risk and material risk one philosophy, one appetite statement and one control environment — so that protecting the business and protecting people are recognised as the same objective.

Key Takeaways

- Build the framework on ISO 31000's eight principles.
- Structure it in three pillars: structure, strategy and systems.
- Apply the framework to boardroom decisions and restructuring with the same rigour as to shop-floor tasks.

- Keep risk appetite stable but responsive.
- Use the five COSO control environment principles to anchor governance.
- Integrate financial and material risk into one process.

In the final chapter, we look ahead to the risk professional the industry will need — and the skills that will define the role.

Chapter 16: The Future Risk Professional

> *Imagine a risk manager whose dashboard does not report last month's injuries, but shows, in real time, which controls were challenged today, how each performed, which are drifting towards their limits and which decisions an algorithm has already taken to keep energy contained. That risk manager is not decades away. In the best operations, the role is being written now.*

Traditional risk management processes — largely mandatory, complicated and inflexible — have reached a turning point. Listed companies and the financial sector recognised the change through horizon scanning; after major economic crises, they strengthened their systems and built more robust structures led by competent risk owners. The mining sector and material risk owners are now following, as they embrace critical control management and more sophisticated processes. The one remaining hurdle is to move from a deeply entrenched culture of compliance towards a professional culture of risk, and from a narrow focus on preventing accidents towards a broader mission of protecting value — of which people are the most precious part.

From Looking Back to Looking Ahead

Where earlier risk functions analysed past events, existing controls, their cost and the efficiency of the process, the new direction is forward-looking: future controls and contingency. The signals are clear — a rapidly changing global economy, conflict, and natural perils intensifying over short periods. Risk management in mining cannot remain the work of a small group producing risk assessments, a risk profile and a mitigation plan separate from the value chain. Protecting people's lives is as important as protecting the business, and the two are inseparable.

In a world shaped by information technology and constantly evolving business strategies, risk management will focus on decisions that allow companies to respond to change efficiently. Identifying change on the horizon is different from identifying material events. Where baseline and issue-based assessments traditionally identified loss, the new era requires the identification of opportunity — not only opportunities to improve existing controls, but opportunities to anticipate and manage the change itself.

Digital by Design

Digitisation is not simply moving paper records onto a screen. Future risk managers will develop and implement digital systems that transform information within a risk decision framework. Algorithms will offer risk owners options for decisions, and responsive systems will automate actions when mitigation begins to fail. Digital twins of shafts, plants and haulage systems will allow controls to be simulated before they are built. Regulatory and shareholder pressure will continue, and trust between the workforce, management and shareholders will always need

attention. Every fatality in mining reveals where systems can improve. The day the industry fully accepts that risk management should never have to be written with the blood of the injured is the day prevention becomes the norm.

The Evolution of the Safety Officer

The first significant change on the horizon concerns the status of the risk practitioner. The natural evolution is for the roles of safety officer and chief safety officer to grow into those of risk officer and risk manager, with the competencies to match. The industry has already taken the first step by assigning critical control verification to line managers, supervisors and control owners — freeing practitioners to move from inspection towards advice, analysis and assurance.

The future risk officer will provide consultation and advisory services. The challenge is to build the right skill set, deepening existing skills while acquiring new ones. The role should be connected to authoritative bodies such as the Institute of Risk Management, with competencies defined by those institutions. Professional affiliation gives access to developing knowledge and makes the practitioner a partner in the transition to a digital, inventive and agile world (Governance Institute of Australia, 2020).

Statistical evidence confirms that neither safety officers nor regulators, working alone, can bring fatalities to an end. The risk function of the future needs a team of multiskilled professionals equipped for a fast-moving world and for the uncertainty of risk and impact.

The Competency Profile

The future risk professional will offer:

- **Clarity of thought and digital affinity** — the mental discipline to frame complex problems and the confidence to use digital tools.
- **Technical expertise** in a designated element of risk management, including artificial intelligence, augmented reality, modelling and simulation, together with knowledge of international risk management standards and the ability to apply them practically.
- **Technological awareness**, including cyber risk and emerging threats.
- **Transparent, confident communication**, using risk information as the core of every negotiation.
- **Curiosity** — the instinct to question the value of change and to seek out the opportunities within it.
- **Data capability** — the skills to secure systems for digitisation, automation, data exploration and modelling, including data science and the use of virtual reality for risk.
- **Natural peril expertise** — as natural hazards increasingly become disasters affecting people, economies and industries, demand for specialists in this field will grow.

- **Board-level fluency** — the ability to help executives and directors apply risk management principles and structures, moving the board's risk function from reporting and compliance to performance.
- **Professional integrity** — observing all relevant laws, regulatory requirements and mandatory codes of practice and conduct, and working within the spirit as well as the letter of the law.

Companies that want to thrive in the face of future risk will build strong risk functions, continually develop their skill sets for new and emerging risks, and appoint risk managers who act with skill, care, diligence, fidelity, honesty and prudence.

The Landscape Ahead

Future risk will shape the future risk professional. Trends, informed speculation and expertise together form a moonshot that lights the way ahead. Investors increasingly consider a company's social responsibility and the social stability of the countries in which it operates. Attention has turned to how companies discharge their responsibilities in labour-sending areas, how they influence environmental outcomes internationally, and how they support cultural well-being and growth. Rising geopolitical uncertainty and threats to biodiversity affect global stability, and the effects of conflict and environmental change will be felt for generations.

The future risk professional will treat the legacy impacts on water, air, soil and nature as pre-conditions to be managed. And through all of this, one constant will remain: the health and safety of people. Any company that uses people to perform hands-on tasks, to take actions that affect others, or to operate equipment in any form will remain responsible for managing the health and safety risks of its activities. Whatever the future holds, health and safety will remain a defining risk, and risk professionals must be equipped to integrate employee protection into the overall risk management strategy.

For global, systemic risks such as climate change, the next generation will add adaptation and resilience to prevention — managing a diverse range of risks that earlier generations helped to create. Scientific research and future studies will convert horsepower into brainpower for the next level of risk managers.

A Challenge for the Next Generation

In closing, the author leaves future risk managers with a challenge to resolve as they prepare for what lies ahead:

> **The challenge:** Our controls will be tested most severely at the moments when human psychology is under the greatest pressure. The next generation's task is to design risk systems that remain effective precisely then — systems in which the evidence of control is so clear, so immediate and so trusted that people, under pressure, choose the safe path because it is also the obvious one.

Key Takeaways

- Move from compliance culture to professional risk culture, and from preventing accidents to protecting value — people first.
- Look forward: identify change and opportunity, not only loss.
- Build digital, data-driven and automated risk systems, including simulation and digital twins.
- Evolve the safety officer role into a professional risk officer role linked to recognised professional bodies.
- Develop multiskilled risk teams with technical, digital, communication and board-level capability.
- Health and safety will remain the constant at the centre of every future risk landscape.

The Conclusion that follows draws the threads of this book together into an agenda for action.

Conclusion: The Groundbreaking Agenda

This book set out to capture the realities of risk management in the mining industry, and to celebrate the groundbreaking progress made by fellow risk managers, engineers and production managers throughout the author's career. It has drawn on practical examples from shafts, pits, smelters, plants and boardrooms, many of them real events whose lessons were paid for dearly.

Some of its ideas will challenge established practice. That is their purpose. A turning point in the author's career came with the realisation of how many different viewpoints, interpretations and strongholds had grown up around the discipline. The profession now stands at a threshold. Its direction should increasingly come from professional risk managers who are qualified to advise, who uphold the highest ethical and professional standards, and who take every opportunity to grow their professional capability.

The message of the book can be gathered into ten actions that any mining organisation can begin tomorrow.

- **Step one to the left.** Move the focus of risk management from proof after the event to vulnerability, pre-conditions and release mechanisms before it.
- **Use the matrix to prioritise, not to prove.** Describe residual risk in words; let measured control effectiveness provide assurance.
- **Name the energy.** Make the hazard the energy source, not the failure of a control, and place the bow tie's top event where control of that energy is lost.
- **Measure every critical control.** Track how often each control is challenged and how it performs, using the five attributes of reliability, availability, survivability, dependability and compatibility.
- **Engineer out the burden.** Shift controls from administrative towards engineered, monitored and automated solutions, so that the machine — not the person under pressure — determines the safe state.
- **Design for human nature.** Build habits, clear rules and potential error analysis into every task, and treat human error as the beginning of an inquiry, never its conclusion.
- **Investigate before the event.** Give hazards in the risk assessment the depth we give events in investigations, and raise investigations to a forensic standard of evidence and causation.
- **Prove readiness.** Extend bow ties past the consequence, populate response and contingency controls, and simulate them.

- **Turn appetite into indicators.** Express material risks as appetite, capacity and a mitigation gap with defined limits, and integrate financial and material risk into one framework.
- **Invest in the profession.** Develop professionally registered, digitally capable, multiskilled risk teams with a voice in the boardroom.

The mining industry has already achieved what many once thought impossible. The next achievement — preventing the event we have never yet experienced — is within reach. It belongs to the organisations willing to look ahead with the same determination with which they have learned from the past.

AC van der Vyver

References

Beck, U. (2016). *The metamorphosis of the world: How climate change is transforming our concept of the world*. Cambridge: Polity Press.

Bernstein, P. L. (1996). *Against the gods: The remarkable story of risk*. New York, NY: John Wiley & Sons.

Binet, A. (1907). *The mind and the brain*. London: Kegan Paul, Trench, Trübner & Co.

Chen, B. H. (2022). *A design of a digital lockout tagout system with machine learning* (Master's thesis). California Polytechnic State University, San Luis Obispo.

Cirillo, P. (2018, May 5). Of risk, fences, and unavoidable falls. [Online essay].

Committee of Sponsoring Organizations of the Treadway Commission (COSO). (2013). *Internal control — Integrated framework*. Durham, NC: AICPA.

Cullen, W. D. (1990). *The public inquiry into the Piper Alpha disaster*. London: HMSO.

Dionne, G. (2013). *Risk management: History, definition and critique*. Montréal: Canada Research Chair in Risk Management, HEC Montréal, CIRRELT and CIRPÉE.

Douglas, J., & Anderson, G. E. (2015). *Leveraging COSO across the three lines of defense*. Committee of Sponsoring Organizations of the Treadway Commission / The Institute of Internal Auditors.

Feynman, R. P. (1986). Personal observations on the reliability of the Shuttle. In *Report of the Presidential Commission on the Space Shuttle Challenger Accident* (Vol. 2, Appendix F). Washington, DC: US Government Printing Office.

Global Tailings Review. (2020). *Global industry standard on tailings management*. London: ICMM, UNEP and PRI.

Governance Institute of Australia. (2020, December). *Future of the risk management professional*. Retrieved from https://governanceinstitute.com.au/app/uploads/2023/11/govinst_fotrmp_nov2020-1.pdf

Harper, D. (2018). Danger; Risk; Threat. *Online Etymology Dictionary*. Retrieved 8 September 2022 from https://www.etymonline.com

Hering-VPT. (2022, March 10). *Understanding dielectric strength in transformers*. Retrieved from https://www.hering-vpt.com/article/understanding-dielectric-strength-in-transformers/

International Civil Aviation Organization. (2020). *Annex 13 to the Convention on International Civil Aviation: Aircraft accident and incident investigation* (12th ed.). Montréal: ICAO.

International Council on Mining and Metals. (n.d.-a). *Health and safety critical control management: Good practice guide* (latest version). London: ICMM. Retrieved from https://www.icmm.com

International Council on Mining and Metals. (n.d.-b). *Critical control management implementation guide* (latest version). London: ICMM. Retrieved from https://www.icmm.com

International Electrotechnical Commission. (2020). *IEC 60296: Fluids for electrotechnical applications — Mineral insulating oils for electrical equipment*. Geneva: IEC.

International Organization for Standardization. (2009). *ISO Guide 73:2009 Risk management — Vocabulary*. Geneva: ISO.

International Organization for Standardization. (2018). *ISO 31000:2018 Risk management — Guidelines*. Geneva: ISO.

International Organization for Standardization. (2019). *IEC 31010:2019 Risk management — Risk assessment techniques*. Geneva: ISO/IEC.

Israel, G. D. (1992). *Sampling the evidence of extension program impact* (Program Evaluation and Organizational Development, PEOD-5). Gainesville, FL: IFAS, University of Florida.

Kaplan, R. S., & Mikes, A. (2012). Managing risks: A new framework. *Harvard Business Review*, *90*(6), 48–60.

Kloman, H. F. (2009, August 20). *A brief history of risk management* (Chapter 2). Lyme, CT: Seawrack Press.

Minerals Council South Africa. (2026). *Facts and figures 2025*. Johannesburg: Minerals Council South Africa.

Reason, J. (2000). Human error: Models and management. *BMJ*, *320*(7237), 768–770.

Rittenberg, L., & Martens, F. (2012). *Enterprise risk management: Understanding and communicating risk appetite*. Committee of Sponsoring Organizations of the Treadway Commission.

Taleb, N. N., Goldstein, D. G., & Spitznagel, M. W. (2009). The six mistakes executives make in risk management. *Harvard Business Review*, *87*(10), 78–81.

Van der Vyver, A. C. (2015, Winter). [Article on risk culture and governance in South African mining]. *RM Professional*. Institute of Risk Management.

Wilde, G. J. S. (1998). Risk homeostasis theory: An overview. *Injury Prevention*, *4*(2), 89–91. https://doi.org/10.1136/ip.4.2.89

Wilson, J. Q., & Kelling, G. L. (1982, March). Broken windows: The police and neighborhood safety. *The Atlantic Monthly*, 29–38.

Acknowledgements

This book is the product of many years of work alongside remarkable people, and the author is grateful to all of them.

To the mineworkers, artisans, supervisors and managers who shared their knowledge underground, in the plants and in countless risk assessments — thank you for teaching far more than any textbook could.

To colleagues at the University of Pretoria's Faculty of Mining Engineering and at Enterprises University of Pretoria, and to every delegate on the c-MiRM course, whose questions sharpened many of the ideas in these pages.

To the SheRisk team in Bronkhorstspruit, Riversdale and Kathu, and to the clients who trusted SheRisk with their most important risks since 1997.

To the professional engineers, rock engineers, geotechnical specialists and investigators whose work underpins several of the case studies in this book.

And to family and friends, for their patience and support throughout the writing.

About the Author



Albert C van der Vyver has 44 years of experience in risk management, much of it in and around the mining industry as an employee, supervisor, manager, consultant and risk owner.

In March 1997 he founded SheRisk, a risk management training and consultancy practice that serves the mining, construction, industrial and petrochemical sectors. SheRisk operates from Bronkhorstspruit in Gauteng, serving mining and manufacturing clients, and Riversdale in the Western Cape, serving the petrochemical sector, and is opening an office in Kathu in the Northern Cape to support the mining industry there.

Albert has lectured in risk management at the University of Pretoria's Faculty of Mining Engineering and helped to develop and present the c-MiRM risk management course through Enterprises University of Pretoria. He has facilitated bow tie analyses, critical control programmes and risk assessments across a wide range of commodities, and has investigated many of the kinds of events described in this book. His conference work includes a 2011 paper on the black swan metaphor in mining, and his writing has appeared in the Institute of Risk Management's *RM Professional*.

He is passionate about building a profession of qualified, ethical and forward-looking risk managers — and about an industry in which every worker returns home safely at the end of every shift.

