RISK & RESILIENCE

When the *Map* Breaks

Why classical risk models were not designed for polycrisis — and what the profession must build next

Katlego Majola · February 2026 · KM Nala Advisory

There is a particular kind of silence that falls in a boardroom when a risk professional presents a sensitivity analysis. It is not the silence of concentration. It is the silence of discomfort — a collective, unspoken awareness that the numbers on the slide are structured guesses wearing the clothes of precision, and that everyone in the room knows it but no one is prepared to say so.

That discomfort is rational. And it is becoming urgent.

The tools that the risk profession built over the past four decades — Value at Risk, Probability of Default, Loss Given Default, Expected Credit Loss, Basel capital ratios — were designed for a specific kind of world. A world where the past was a reasonable guide to the future. Where individual risks could be isolated, measured, and priced. Where the system you were modelling would still be recognisable when your forecast period arrived.

That world is gone. What has replaced it is something the profession does not yet have adequate instruments to measure.

01

The Distinction We Stopped Making

Before we can discuss what the new toolkit requires, we need to be precise about what the old one actually does — because most of the confusion in risk practice stems from conflating three fundamentally different things.

Modelling is the architecture of causality. A model says: here are the variables I believe drive the outcome, and here is how they relate to each other. It is a structured map of beliefs about how the world works. A DCF, a credit model, a capital adequacy framework — these are not forecasts. They are formalised assumptions. The mistake is treating the output as the answer rather than as a structured question.

Judgement is what populates a model when reliable data does not exist. What is our realistic churn rate in year three? Will the regulator respond aggressively? Will this sovereign honour its obligations? Judgement is where humans are irreplaceable — and where most of the real risk actually lives. The failure mode is outsourcing judgement to the model: letting the spreadsheet

decide by plugging in comfortable numbers. That is how you launder bad assumptions into defensible outputs.

Probabilistic framing is the honest acknowledgement that you do not know the answer — only a range of plausible answers with different likelihoods. It is not pessimism. It is intellectual honesty. The boardroom cringe happens here, because expressing uncertainty feels like admitting ignorance rather than demonstrating analytical rigour. In reality, a point estimate in a complex environment is the least honest output a risk function can produce.

> *A sensitivity analysis does not reveal weakness in your model. It reveals the assumptions your model cannot see — which is precisely the information a decision-maker needs.*

When these three things are kept distinct, risk management is coherent. When they collapse into each other — when the model substitutes for judgement, when the point estimate replaces the distribution, when precision is mistaken for accuracy — the function fails quietly, long before the dashboard turns red.

02

What Polycrisis Does to the Toolkit

The word polycrisis, developed by complexity theorist Edgar Morin and brought into mainstream policy discourse by historian Adam Tooze, describes something specific. Not many crises occurring simultaneously. Something structurally different: crises that are causally entangled — that interact with each other, amplify each other, and produce emergent dynamics that none of them would generate in isolation.

The 2020s have been a laboratory for this phenomenon. COVID-19 disrupted supply chains. Disrupted supply chains drove inflation. Inflation forced central banks to raise rates. Rising rates stressed sovereign balance sheets already weakened by pandemic fiscal responses. Geopolitical fracture accelerated energy transition but also weaponised energy dependency. Climate stress compounded food insecurity. Institutional delegitimisation made coordinated responses harder precisely when coordination was most needed.

Each of these was hard to model individually. Their interactions were, in the classical sense, unmodelable — because the feedback loops were non-linear, the system actors were adapting in real time, and the causal arrows ran in multiple directions at once.

This is the move from complicated to complex systems. And the classical risk toolkit was built for complicated problems.

CLASSICAL ASSUMPTION	WHAT POLYCRISIS DOES TO IT	THE RESULTING FAILURE MODE
Historical data predicts distributions	Distributions shift mid-model; regimes change	VaR understates tail risk by design
Risks are separable and priceable	Risks become correlated in crisis; hedges fail together	Diversification disappears when needed most
The system recovers	Disruption can be permanent;	Mean-reversion assumptions

after disruption	new equilibria form	embed false comfort
The institution survives to implement the plan	Absorption capacity itself becomes the critical variable	Risk plans that assume a functioning responder

That last row is the one the profession has been slowest to confront. Classical risk management implicitly assumes that the institution assessing risk will retain the capacity to act on its assessments. Polycrisis is precisely the condition that exhausts that capacity — not by delivering a single catastrophic blow, but by delivering compounding, simultaneous demands on attention, capital, and decision-making bandwidth that degrade response quality progressively and invisibly.

> *The most dangerous scenario is not the one that destroys your institution. It is the one where things feel okay, where each quarter you tell yourself you are almost through it — and quietly, the window for effective response narrows.*

03

The Governance Problem Inside the Analytical Problem

There is a second layer to this that pure methodology cannot resolve: the governance architecture of risk management was not designed for structural signals that arrive before the established metrics confirm them.

Consider the scenario that exposes this most sharply. A structural diagnostic — call it any forward-looking framework that measures system pressure rather than current state — flashes amber. But the traditional ratios are green. The credit committee is bullish. The regulator has no immediate objection. The CEO wants to proceed.

A Director who votes to constrain growth in that room needs more than an elegant framework. They need fiduciary cover — something as legally and institutionally defensible as a capital adequacy breach. Because analytical courage without institutional protection is, in practice, just career risk.

This is the last-mile problem in risk governance. It is not a failure of intelligence or analysis. It is a failure of the institutional architecture that translates structural insight into sanctioned action. The profession has spent decades hardening its models. It has spent far less time hardening the governance mechanisms that determine whether model outputs actually change decisions.

Under IFRS 9, the mechanism already exists in embryonic form: judgmental overlays, applied when standard models fail to capture the current environment. Every institution that applied a COVID overlay in 2020 was, implicitly, acknowledging that its model could not see what its analysts could. The missing piece is a structured, auditable, evidence-based foundation for those overlays — so they move from expert intuition to structural diagnostic, and from discretionary adjustment to governed decision trigger.

04

A New Toolkit: Four Additions, Not

Replacements

The classical toolkit retains value in stable environments. The argument is not for replacement but for supplementation — four additions that extend the profession's reach into complex, entangled, non-stationary conditions.

ADDITION 01

Resilience Modelling Over Optimisation Modelling

The question stops being *what is the optimal path?* and becomes *which paths keep us viable across the widest range of futures?* This is a shift from maximising expected value to maximising robustness. Mathematically, it looks less like a discounted cash flow and more like options thinking — preserving flexibility has explicit value even when it costs short-term efficiency. An institution that is slightly suboptimal in normal conditions but retains response capacity in stressed conditions will consistently outperform one that optimises for the base case and fails when the base case does not arrive.

Shift: From maximising expected value → to maximising viable range

ADDITION 02

Absorption Capacity as an Explicit Balance Sheet Item

Absorption capacity — the ability of an institution to take a hit and continue functioning — is almost never modelled explicitly. It is assumed. The assumption is precisely what polycrisis exhausts. Modelling it requires asking: how many simultaneous stresses of what magnitude can this institution absorb before its response quality degrades? What are our liquidity buffers, operational redundancies, decision-making bandwidth constraints, and stakeholder trust reserves — and at what point does their depletion become self-reinforcing? This is not a soft question. It has hard, quantifiable components. The profession's failure to model them explicitly is a gap, not a limitation.

Shift: From avoiding bad outcomes → to maintaining capacity to respond to any outcome

ADDITION 03

Structural Scenario Planning — Done Properly

Not the three-box base/bull/bear structure that is probabilistic thinking in disguise. Genuine scenario planning constructs structurally distinct futures — worlds with different causal logics, not just different parameter values on the same model. The point is not to predict which scenario arrives. It is to develop strategic flexibility and early warning indicators that remain valid across scenarios, so that when the world begins to reveal which path it is on, the institution recognises the signal 18 to 30 months before the GDP data or balance sheet confirms it. The scenarios should be uncomfortable enough to be useful and specific enough to generate observable indicators.

Shift: From predicting outcomes → to developing indicators that reveal which future is arriving

ADDITION 04

Pre-Mortem and Inversion as Structured Judgement

Rather than asking *what happens if X?*, ask *we are destroyed in three years — what killed us?* This bypasses the optimistic bias that contaminates forward-looking models and forces judgement to operate in a register that surfaces systemic vulnerabilities invisible to sensitivity analysis — because they are not variables in the model. They are assumptions the model cannot see. The pre-mortem is not a pessimism exercise. It is the most rigorous stress test available for the model's own architecture. Institutions that do this well tend to discover their most dangerous risks are the ones their models are structured to ignore.

Shift: From stress-testing variables → to stress-testing the model's assumptions

05

The Epistemological Shift Underneath the Methodology

All four additions point to the same underlying reorientation. The goal of risk analysis is shifting — it must shift — from **prediction** to **preparedness**. The question is no longer primarily *what will happen?* It is *who do we need to be, and what do we need to hold in reserve, to navigate what we cannot predict?*

This is harder to sell to a board that wants a number. It requires a different kind of analytical courage — not the courage to commit to a forecast, but the courage to present a range and defend the process that produced it. Not the courage to say *we know*, but the courage to say *here is what the structure of the uncertainty looks like, and here is how we are positioned within it*.

The risk profession's credibility was built on precision. In a polycrisis environment, that credibility now depends on the profession's willingness to be honest about what precision cannot do — and to replace false certainty with something harder and more valuable: structured, defensible, institutionally actionable uncertainty.

PROOF OF CONCEPT

South Africa as the Live Case Study

In 2010, South Africa presented every structural indicator of an emerging market with genuine fiscal optionality. Debt-to-GDP at 31.5% — lower than Germany, the United Kingdom, and the United States. Functioning institutions. A renewable energy programme ready to scale. A national development blueprint in draft. The established metrics were green.

What followed is documented in the KM Nala 2010–2025 Counterfactual Backtest. The governance variable deteriorated. The lagging indicators — credit ratings, GDP per capita, debt trajectory, rail volumes, energy reliability — confirmed that deterioration years after the structural signals were visible. By the time the metrics turned red, the absorption capacity required to respond effectively had itself been eroded.

This is the sequence that classical risk models are structurally incapable of capturing: the degradation of response capacity as a consequence of compounding governance failures, long before the balance sheet reflects it.

The forward model maps four structurally distinct scenarios to 2035. The base case — which we assign 42% probability — is not catastrophe. It is the Muddle Republic: a world where things feel approximately okay, where the dashboard stays amber, and where the window for structural recovery narrows quietly each year without a moment of obvious crisis to force corrective action. That is the polycrisis signature. Not the dramatic failure. The slow exhaustion of optionality in the absence of a trigger visible enough to justify acting against prevailing sentiment.

The 2026 Budget confirmed the turning point on debt — 78.9% stabilising and beginning to fall. The credit rating upgrade was the first in 16 years. Eskom's Energy Availability Factor is holding above 65%. These are genuine Phoenix signals. GDP growth at 1.6%, less than half the global average, confirms the structural work is incomplete.

The early warning indicators that will reveal which trajectory is actually unfolding are not the lagging macro metrics. They are the municipal water non-revenue rate (currently 47%, against a peer average of 30%), the skills pipeline into artisan and technical trades, coalition durability at local government level, and the integrity of the revenue collection ecosystem. These structural diagnostics will tell the story 18 to 30 months before the GDP data confirms it.

DEBT-TO-GDP · 2026	GDP GROWTH · PROJECTED	NON-REVENUE WATER	ESKOM EAF · YTD 2026
78.9%	**1.6%**	**47%**	**65%+**
Stabilising and beginning to fall. First genuine turning point. Fiscal optionality not yet restored.	Less than half the global average of 3.3%. Phoenix requires sustained 3%+. Muddle Republic persists.	Against a peer average of 30%. A structural signal the macro dashboard does not capture.	Strongest Phoenix signal. Sustained above 67% would confirm a trajectory shift.

What makes a system — institutional, sovereign, or organisational — choose the long game when the short game is more legible, more defensible, and more immediately rewarded?

That question does not have a clean analytical answer. But the discipline of asking it seriously — with data, with structured uncertainty, with explicit attention to what the model cannot see — is where the next generation of risk practice begins.

The map broke. That is not a failure of the cartographers. It is a change in the territory. The profession's task now is to build instruments adequate to the terrain that actually exists — one where the most dangerous risks are structural, entangled, and visible only to those willing to look before the ratios turn red.

The tools for that exist, in nascent form, across complexity science, scenario planning methodology, and the emerging literature on systemic resilience. What has been missing is their translation into the governance architecture, institutional language, and fiduciary frameworks that allow them to function as decision tools rather than commentary tools.

That translation is the work. It is the most important risk management project of this decade.

Katlego Majola

Independent Risk Expert & Board Member
KM Nala Advisory · February 2026

The SA Counterfactual Backtest and 2025–2035 Forward Model referenced in this article are available at kmnala.co.za