

Risk management as decision support in insurance

By Nicolai von Rummell and Darshan Purmessur

Presented at the Actuarial Society of South Africa's 2026 Convention
Cape Town International Convention Centre, 7–8 October 2026

ABSTRACT

Insurance is a continuous sequence of decisions made under uncertainty, yet for many insurers the risk function is built mainly for governance and regulatory assurance. That work is necessary, but it does not by itself improve the decisions that determine performance. This paper argues that the function should also act as decision support: supplying independent, structured analysis at the point where material choices are made, while ownership stays with the business. We frame the argument with four roles: Architect, Controller, Integrator and Enabler. Comparing these with South Africa's Prudential Standard GOI3 shows that regulation specifies the first two well but leaves enterprise aggregation and decision-focused challenge largely to the firm.

Drawing on decision science, we argue that material decisions should be assessed on the range of outcomes, the downside tail, dependence with existing exposures, and balance-sheet capacity, not on point forecasts and confident narrative, and we set out where in the business this input matters most and how to tell whether it is embedded. We then ask why risk leaders are so often absent from those decisions, and draw on discussions with senior risk leaders at eight of the largest South African insurance groups to illustrate the mechanisms in practice: late involvement, abstract risk appetite, thin quantitative capability, and the way another function fills the gap when the risk function does not. The discussions suggest that influence rests on credibility, early involvement and usable risk appetite rather than on formal authority. We close with what the Enabler role requires in practice, including the option of locating the capital model in the risk function, and with simple measures of whether decisions have become more disciplined.

KEYWORDS

ERM; risk appetite; decision-making under uncertainty; CRO; risk governance

CONTACT DETAILS

Nicolai von Rummell, School of Statistics and Actuarial Science, University of the Witwatersrand, Johannesburg; Email: nicolai.baronvonrummell@wits.ac.za

Darshan Purmessur, QED Actuaries and Consultants, Mauritius

1. INTRODUCTION

1.1 This paper argues that a compliance-only conception of risk management is incomplete in an industry whose core activity is pricing and aggregating uncertainty. South Africa's Prudential Standard Governance and Operational Standards for Insurers (GOI3) establishes a necessary governance baseline, but it does not by itself ensure that risk analysis informs strategic choices. Where firms treat regulation as the endpoint, risk work concentrates on documentation and assurance, while material decisions proceed with little explicit analysis of risk or understanding of the dependence between exposures.

1.2 The argument proceeds in five steps. First, we clarify what a modern risk function is responsible for by organising its work into four roles: Architect, Controller, Integrator and Enabler. These are positioned on two dimensions: strategic versus operational, and control and compliance versus decision support. Widely used international frameworks such as COSO enterprise risk management (ERM) (Committee of Sponsoring Organizations of the Treadway Commission, 2017) and ISO 31000 (ISO 31000, 2018) hold that risk management should be embedded in decision-making and not treated only as a control mechanism. This taxonomy provides the labels used throughout the paper and a way to distinguish what regulation specifies clearly from what it leaves open.

1.3 Second, we compare GOI 3 (Prudential Authority (South Africa) 2018a) with this framework. The alignment is strongest for the Architect and Controller roles, which cover system design, appetite, monitoring and reporting. It is weaker on the aggregation of exposures across business units and on the support of consequential choices that define the Integrator and Enabler roles. GOI 3 is therefore a floor rather than a blueprint for a function that consistently supports decision-making, and firms that read it as the whole of the task will underuse the function.

1.4 Third, we explain why the less specified parts of the function matter. Section 3 sets out why material decisions require an explicit treatment of uncertainty and why unaided judgement is unreliable. The paper then identifies where embedded risk input is most needed across the business and asks why, in many insurers, the Chief Risk Officer (CRO)

is often absent from key strategic and capital decisions (Mattone et al., 2024). We set out the mechanisms in four groups: how the function is positioned, when it is brought into the process, what it contributes once involved, and the feedback loop that entrenches the pattern. At these points we draw on discussions with senior risk leaders at eight of the largest South African insurance groups, which illustrate how these mechanisms appear in practice.

1.5 Fourth, we use innovation decisions as a focused case study in Section 4. The section identifies five recurring obstacles: procedural delay, undue conservatism under ambiguity, communication gaps, weak strategic alignment and siloed decision-making. It then describes how earlier and more substantive participation can address them while leaving decision ownership with line management.

1.6 Fifth, we set out what must change for the function to perform the Enabler role consistently: incentives and expectations, skills and capabilities, performance measurement, and the credibility of risk advice. The aim is that material choices are made with explicit recognition of risk, uncertainty, capacity, and aggregation, so that like risks are treated alike and trade-offs are made on all the available information.

2. THE ROLE AND RESPONSIBILITIES OF THE RISK MANAGEMENT FUNCTION

This section sets out the responsibilities of the risk management function, organised into four roles: Architect, Controller, Integrator, and Enabler. To make the discussion concrete, we position them on two dimensions in Figure 1: strategic (longer-term) versus operational (day-to-day), and control and compliance versus decision support. The same responsibilities appear in many firms under different labels.

2.1 Architect

The Architect designs and maintains the structure and processes of ERM: system and process design, governance, risk policies, methodologies, and documentation. It provides the framework against which risk-taking can be assessed consistently. Without it, the function lacks consistency and credibility.

2.1.1 Methodologies and tools

Develop, maintain, and improve methodologies for identifying, assessing, quantifying, and managing risks. This includes consistent models, stress-testing approaches, and scenario modelling.

— *If missing:* risk assessments become subjective and inconsistent, and different business units evaluate similar risks differently, which leads to misaligned decisions. At Long-Term Capital Management, sophisticated models were not matched by methods capable

of testing the correlation and liquidity assumptions on which they rested (President's Working Group on Financial Markets, 1999). Scenario suites can also fall out of date: before 2020, pandemic scenarios were considerably less developed than the catastrophe and market scenarios run alongside them.

2.1.2 *System design and governance*

Design, implement, and evolve the enterprise risk management system, including governance structures, committees, policy frameworks, and decision rights.

— *If missing:* fragmented or siloed governance weakens control and creates gaps, decision rights are unclear at the points where risk is taken, and an ineffective structure tends to produce reactive risk management.

2.1.3 *Documentation and institutional knowledge*

Maintain clear documentation of risk frameworks, policies, models, decisions, and governance arrangements. This preserves institutional memory and supports auditability.

— *If missing:* the organisation struggles to respond to regulatory inquiries, restructuring, or emerging risks. At Equifax, the absence of comprehensive, documented patch-management processes and accountability contributed to a major data breach and a loss of public trust (U.S. House Committee on Oversight and Government Reform, 2018).

2.2 **Controller**

The Controller applies the risk management framework in day-to-day operations: identifying risks, assigning ownership, defining appetite, monitoring exposures, and setting priorities. It also ensures that risks are reported to decision-makers in a timely and usable form.

2.2.1 *Risk ownership and accountability*

Ensure that all material risks are identified and that each has a designated owner responsible for managing it within agreed limits.

— *If missing:* risks are left without an owner, with no one accountable for keeping an exposure within limits and no consolidated oversight to constrain it early.

2.2.2 *Risk appetite framework*

Develop clear and consistent risk-appetite statements for all material risk types, aligned with how those risks are assessed, combining qualitative statements with quantitative thresholds.

— *If missing:* risk-taking becomes inconsistent across the group and limits stop constraining strategic choices. The Financial Services Authority's (FSA's) review of the failure of the Royal Bank of Scotland found that the board had not tested the ABN AMRO Bank's acquisition against a clear statement of risk appetite, and that the transaction proceeded

on limited due diligence and a capital position that left little margin for a downturn (Financial Services Authority, 2011).

2.2.3 *Risk monitoring and assessment*

Monitor risks across the business, assess whether they remain within appetite, and establish processes for escalation and remediation.

- *If missing:* breaches and deteriorating positions persist without timely escalation or corrective action. In its review of Silicon Valley Bank, the Federal Reserve attributed the failure to mismanagement of basic interest-rate and liquidity risk, including weaknesses in liquidity stress testing, limits, and contingency funding plans (Board of Governors of the Federal Reserve System, 2023).

2.2.4 *Materiality and prioritisation*

Direct attention and resources to material risks, while remaining aware of aggregation from minor exposures and avoiding over-analysis of trivial items, so that the function's resources reflect the materiality of the risks it manages.

- *If missing:* the risk team is consumed by minutiae, devoting attention to easily measured but immaterial exposures while larger ones go under-examined.

2.2.5 *Clear and actionable reporting*

Deliver risk reports that are timely, focused, and usable, allowing decision-makers to respond to current and emerging threats.

- *If missing:* decision-makers receive incomplete or irrelevant information. Carillion, a British outsourcing and construction firm, collapsed suddenly in 2018. Weak corporate governance and reporting that obscured its financial position were among the causes identified (House of Commons Work and Pensions Committee & House of Commons Business, Energy and Industrial Strategy Committee, 2018).

2.3 *Integrator*

The Integrator looks at risk across business lines and units: breaking down silos, looking across portfolios, fostering risk culture, and keeping a forward-looking view. Done well, it lets the organisation learn from a failure in one business unit before it is repeated in another, and strengthens resilience to adverse events (International Association of Insurance Supervisors, 2019).

2.3.1 *Risk aggregation and correlation*

Aggregate risks across business units and silos to identify hidden concentrations, correlations, and systemic vulnerabilities.

- *If missing:* the organisation may believe it is diversified when its exposures are in fact concentrated, so that a single event produces losses across many positions at once.

2.3.2 *Forward-looking risk view*

Anticipate future risks through emerging-risk identification, scenario analysis, and stress testing across the planning horizon and beyond (European Insurance and Occupational Pensions Authority, 2015).

- *If missing*: the firm reacts to events rather than preparing for plausible shifts, and unintended exposures accumulate quietly. Lloyd's required policies to state explicitly whether cyber losses were covered or excluded specifically to reduce unintended non-affirmative cyber exposure (Lloyd's, 2019).

2.3.3 *Risk culture and awareness*

Foster a culture in which risk awareness and ethical behaviour are expected at all levels, and in which accountability does not collapse into blame.

- *If missing*: risk becomes someone else's problem. At Wells Fargo, aggressive sales targets and a weak risk culture led employees to open fraudulent accounts, harming customers and triggering regulatory sanctions (Office of the Comptroller of the Currency, 2016).

2.4 **Enabler**

The Enabler embeds risk management where decisions are made. By providing independent analysis and challenge, supporting risk-return trade-offs, and informing capital decisions, it helps the organisation pursue opportunities with discipline.

2.4.1 *Independent decision support*

Provide objective risk input to major business decisions while accountability remains with the business owners.

- *If missing*: risk input arrives too late to shape options, or becomes contested late in the process, which reduces its influence. At Lehman Brothers, risk professionals flagged excessive real-estate exposure but lacked the influence to alter material decisions (Valukas, 2010).

2.4.2 *Constructive challenge function*

Question assumptions, test alternative scenarios, and challenge strategic and operational decisions on the basis of risk insight.

- *If missing*: groupthink prevails and an untested belief persists unchallenged until an event exposes it: that controls are adequate, or that a market will behave as it has.

2.4.3 *Capital efficiency and optimisation*

Support capital allocation by bringing risk-adjusted performance measures into planning and investment decisions, and provide oversight of value-based management that weighs profit against risk.

- *If missing*: capital is deployed inefficiently, trapped in low-return areas or under-allocated

to promising but misunderstood ventures. The monoline insurers MBIA and Ambac expanded into structured finance without pricing the risk appropriately, and suffered large losses in the global financial crisis (International Monetary Fund, Monetary and Capital Markets Department, 2008).




FIGURE 1. Four roles of risk management positioned on two dimensions

2.5 The rest of the paper uses these labels as shorthand. The next section overlays GOI 3 on the framework to assess which responsibilities the standard covers, where it is silent, and where interpretation would be needed for a fully embedded function.

2.6 GOI3: The regulatory definition of the risk function

2.6.1 To understand the regulatory conception of risk management in South African insurance, this section compares the responsibilities set out in Section 2 with how they are reflected in the Prudential Standard GOI 3 (Prudential Authority (South Africa), 2018a; 2018b). Appendix A gives the detailed mapping of all fourteen responsibilities; the aim is to identify where a broader interpretation may be required for insurers that want to embed risk more deeply into strategy.

2.6.2 From the comparison, several patterns emerge. GOI 3 is strongest on the

Architect and Controller roles, setting clear expectations for system design, documentation, appetite, monitoring, and board-level reporting, consistent with a compliance-based governance model; even there, the requirements for methodology (Responsibility 1) and for decision-relevant reporting (8) are stated only in general terms, and materiality-based prioritisation (7) is not addressed at all. The Integrator role is addressed unevenly: cross-silo aggregation (9) is only partly required, while forward-looking work through Own Risk and Solvency Assessment (ORSA) (10) and risk culture (11) are mandated. The Enabler role is least covered: independent decision support (12), constructive challenge (13), and capital efficiency (14) are only partially addressed or not required at all.

2.6.3 GOI 3 thus provides a foundation for governance, documentation, and control, and it states clearly the need for an independent risk management function. But it places that function mainly in a supporting and assurance role, helping the board and senior management maintain a compliant and effective system. That framing is necessary, but it may not be enough for a competitive insurance market in which risk must also support decision-making and innovation. A function designed only to comply with GOI 3 would likely fall short of the Integrator and Enabler roles described above.

2.6.4 Most insurers should treat the regulation as a minimum rather than a specification (CRO Forum, 2012). The task is to expand the function's scope deliberately: to clarify its role in materiality-driven prioritisation, cross-silo aggregation, independent decision support, constructive challenge, and capital efficiency, the areas that GOI 3 leaves mostly unspecified. Where an insurer bases its understanding of risk management on GOI 3 and the earlier Board Notice BN 158 guidance (Registrar of Long-term Insurance and Registrar of Short-term Insurance, 2014) alone, the resulting function tends to be strong on governance but light on quantitative analysis and commercial engagement.

2.6.5 The next section examines one element of the Enabler role, independent decision support (Responsibility 12), and argues why structured risk input at the point of decision matters for strategy, capital efficiency, and long-term performance.

3. EMBEDDING RISK IN DECISION-MAKING

3.1 Section 2 defined the Enabler role as providing independent, methodical, risk input where material decisions are made, while ownership and accountability stay with the business. This section develops what that means for insurers. The core claim is that decisions based mainly on point forecasts and confident narrative underweight extreme outcomes, correlation, and capital and liquidity strain, so a firm can look profitable on its base case while becoming fragile under stressed conditions.

3.2 We treat embedded risk as a problem of decision quality before it is one of governance. First, we set out what decision-relevant analysis must add beyond a business case: the distribution around the expected outcome, the scenarios in the tail, correlation with other existing exposures, and the capacity of the balance sheet to absorb a bad result.

Second, we summarise why unaided judgement and organisational dynamics distort risk perception and suppress timely challenge, especially when uncertainty is high. We then turn to operating practice: a map of the processes where structured risk input matters most, tests for whether the application of risk appetite to decisions is working, and a diagnosis of why the risk leader is often brought in late.

3.3 Parts of the argument are checked against discussions with senior risk leaders at eight of the largest insurance groups operating in South Africa, conducted between January and March 2026. The discussions were semi-structured and covered how the risk function becomes involved in decisions, how risk appetite is used in practice, and what gives the function influence. They offer a practitioner check rather than a representative survey, and are used here to illustrate points that would otherwise remain abstract. All references to them are anonymised.

3.4 Why every decision requires risk analysis

3.4.1 Every strategic or operational decision in insurance rests on a projection of expected outcomes, usually a business case, an investment return, or a financial model. That is only half of the question. The other half is the distribution around that expectation: how far outcomes can swing, and whether the downside compounds with exposures the firm already holds.

3.4.2 A business case may show positive expected value, but if the downside tail is large and potentially correlated with existing exposures, or hard to assess, a rational decision-maker may decline it (CRO Forum, 2023b; Slovic, 1987). An overstated downside, equally, can cause a firm to pass up a good opportunity. Risk appetite frameworks are meant to make this trade-off explicit, yet they are often left out of day-to-day and even longer-term decisions (CRO Forum & CRO Council, 2013) How decision-makers perceive risk shapes the choices they make, and perception can diverge from the underlying risk, particularly without structured input.

3.4.3 A stylised coin-flip: why expected value is not enough

3.4.3.1 Consider three fair coin-flip games with identical positive expected value but very different decision implications:

- **Game 1 (many small bets):** Flip a fair coin repeatedly (e.g., 10^9 times). Receive R2 on heads and pay R1 on tails each round. The per-flip expected value is $E[X] = 0.5 \times 2 - 0.5 \times 1 = 0.5$, so over 10^9 repetitions the expected outcome is +R500 million. The standard deviation of the total is about R47 000, so the realised outcome is R500 million give or take a rounding error; for a normally capitalised insurer this is almost appetite-agnostic.
- **Game 2 (one big bet):** Flip once. Pay R1 billion on tails; receive R2 billion on heads. The expected value is still +R500 million, but the distribution is bimodal (either –R1 bn or +R2 bn). Whether to play depends on *risk appetite* and *risk capacity*: can a R1 bn loss be

absorbed without breaching solvency, covenants, or stakeholder tolerances? If not, the rational decision is ‘no’ despite the positive mean.

- **Game 3 (one big bet while already deeply underwater):** Same payoffs as Game 2, but start R1.2 billion in debt with no realistic path to repay. Here the *state* matters: ruin is already effectively realised, and any further loss falls on creditors rather than on the decision-maker. The payoff facing that decision-maker is therefore truncated below and open above, and the gamble becomes attractive even to an agent who is risk-averse over outcomes they would actually bear. When distress is severe, preferences can turn risk-seeking, because the convex payoff dominates the mean.

3.4.3.2 These games make three points for insurance decisions. Expected value is necessary but not sufficient. Perception of risk and dependence on existing exposures determine acceptability. Risk appetite is *state-dependent*: as capital buffers fall, some choices become unacceptable (Game 2) even if the expected outcome is high, yet when situations are desperate the same gamble can become attractive (Game 3), which is why formal appetite statements, triggers for breaches, and governance structures matter. Repeated small-stake decisions differ from one-shot, high-stake ones: the first is governed by pooling across repetitions, the second by the size of the loss relative to capacity.

3.4.3.3 The coin-flip games are examples of *decision-making under risk*, where the distribution of outcomes is known and stable enough that the mean, volatility, and tail are well-defined. Many insurance decisions are not in that position. Insurers operate across a spectrum, from settings where experience data and credible models make the distribution reasonably tractable to settings dominated by *uncertainty*, where the structure of outcomes is changing, relevant data are scarce, and probabilities are not well understood.

3.4.3.4 If risk input is to be part of a decision, it has to enter the decision criterion in a way that makes the trade-off explicit. One practical route is a value-based framing: evaluate a decision not on expected profit alone, but on profit after charging for the capital that must support the risk profile. At a high level, Economic Value Added (EVA) can be written as (Chmielowiec-Lewczuk, 2015):¹

$$\text{EVA} = \text{Profit} - (\text{Cost of Capital}) \times \text{Capital}.$$

3.4.3.5 Under *risk*, the distribution supports direct quantification: expected profit can be set against a tail-based capital requirement, such as a percentile or tail value at risk (TVaR) measure drawn from the same distribution of changes in own funds. Under *uncertainty*, the framing still holds but the inputs are approximate, with scenario analysis

1 This expression is intentionally high level. In implementation, firms must define the profit measure (e.g. net operating profit after tax (NOPAT) or expected change in own funds), the relevant capital basis (e.g. regulatory, rating, or internal economic capital plus buffers, and the treatment of excess capital), and a cost-of-capital rate consistent with their financing and performance rules.

and ranges used to bound plausible outcomes and to locate where tail and correlation risk may emerge. Either way, the purpose is the same: to make the value, risk, and capital trade-off explicit, comparable across options, and testable against appetite and capacity.

3.4.3.6 In insurance and finance, exposures that look attractive in expected terms can still destroy value once risk and capital intensity are made explicit (Kraus, 2013). Sometimes the distribution is tractable and the failure is one of modelling: hidden correlations, aggregation, or regime shifts that matter only in the tail. In other settings the decision is characterised by uncertainty, where thin data and shifting underlying risks mean the apparent stability of the business case reflects a lack of knowledge and not real robustness. A value-based lens forces the same questions either way: what is the expected benefit, what capital must support the downside, and what value remains after charging for that capital? Two cases illustrate the point.

— **AIG and the 2008 credit default swap crisis**

Before the global financial crisis, AIG Financial Products, a subsidiary of AIG, expanded heavily into selling credit default swaps. The expected return looked attractive: premium income with little short-term capital required. The perceived chance of a broad housing collapse was low, so hedging and capital buffers were thin. When U.S. housing markets fell, losses across the contracts arrived together; AIG had not modelled this systemic risk, could not meet the resulting collateral calls, and was rescued with government support that ultimately reached about \$182 billion. The business case was compelling on paper, but the underlying risk was never identified (U.S. Department of the Treasury, 2012; Financial Crisis Inquiry Commission, 2011).

— **Hurricane Andrew and geographic concentration**

In the early 1990s, hurricane exposure in Florida was widely treated as diversified across policyholders. Premium growth was strong and the expected annual loss looked manageable. Many insurers had written heavily in the same coastal zones without measuring the resulting concentration. When Hurricane Andrew struck in 1992, losses exceeded \$15 billion (1992 USD) and bankrupted eleven insurers. The risk had been understated because aggregation and correlation were ignored, and enterprise-level exposure analysis was absent from the decisions that built the books (Insurance Information Institute, n.d.; AIR Worldwide & Artemis.bm, 2017; Kunreuther et al., 1998).

3.4.3.7 Without quantitative risk input produced by a consistent method, decision-makers fall back on heuristics, intuition, or incomplete data. An effective risk function ensures risk is incorporated in a decision-relevant way: quantified through distributions and capital impacts where decisions are made under risk, and bounded through scenarios where uncertainty dominates. The next section turns to why the risk function is well placed to provide that input.

3.5 Limits of unaided risk judgement

3.5.1 The task is not only to form a view on the expected outcome but to form a credible view on downside tail exposure and capital required, the two inputs that have to be weighed together if risk is genuinely part of the decision. A large body of work in behavioural economics and cognitive psychology shows that people lean on simplifying heuristics in complex, uncertain choices; the shortcuts are often useful, but they introduce distortions, so risk perceptions depart from the underlying risk in systematic ways unless a consistent method is imposed.

3.5.2 Several patterns recur in insurance contexts (Kunreuther et al., 2013). Salient or recent events loom larger in judgement (availability); comfortable narratives can feel truer than base rates (conjunction, base-rate neglect); initial numbers anchor later estimates even when conditions have changed (anchoring); confidence in one's own resilience tends to be too high (optimism); and small probabilities are systematically overweighted while large ones are underweighted (probability weighting). Tversky & Kahneman (1974) documented that availability and anchoring skew probability and numerical estimates even among trained participants; prospect theory shows that outcomes are valued against a reference point rather than in absolute terms, and that probabilities enter the valuation in distorted form, so that two prospects of identical expected value can be treated very differently (Tversky & Kahneman, 1974; Kahneman & Tversky, 1979). Slovic's work shows that risk judgements track attributes such as dread and unfamiliarity rather than actual frequency, which explains overreaction to vivid hazards alongside neglect of diffuse, chronic ones (Slovic, 1987). Together these effects bias underwriting, pricing, and capital decisions toward what is vivid, familiar, or anchored in history rather than what is most probable and consequential.

3.5.3 The failure modes differ across the risk-dominant and uncertainty-dominant settings distinguished in Section 3.4. Under risk, biases distort probability and impact judgements even when data and models exist, which leads to underestimation of tails, correlations, and accumulations. Under uncertainty, decision-makers default to coherent narratives and single-point forecasts, and the common failure is false precision about parameters that are not well understood. The remedy is therefore not always more modelling, but decision support matched to the context: quantified analysis where it is credible, and scenario analysis with explicit bounds where it is not.

3.5.4 These predictable judgement errors call for institutional safeguards, and the risk management function is well placed to provide them. The function applies documented techniques for measurement and quantification; it relies on data and statistical models rather than intuition alone; it is structurally independent from the commercial incentives that colour first-line judgements; it holds an enterprise view that reveals correlations and accumulations across silos; and it has a mandate to offer challenge and a second view. In practice this can run through standard processes and checkpoints that require explicit statements of expected benefit, tail exposure and capital intensity, and

the implied value after a capital charge. Where uncertainty dominates, a range bounded by scenarios serves the same purpose.

3.5.5 This is why the first responsibility matters: develop and maintain methods for identifying, assessing, quantifying, and controlling risks that are sound, documented, and repeatable. The aim is consistency across business units, time, and teams with different experience and incentives, so that like risks are treated alike and the differences that matter for a decision are visible. Where these tools are absent or unevenly applied, firms fall back on subjective judgement, and outcomes then turn on who is in the room rather than on the underlying risk. Used this way, risk management can be the analytical discipline that complements a financial forecast with an explicit treatment of risk and uncertainty, improving decision quality without displacing first-line accountability. The next section turns to the organisational constraints that keep this discipline from being applied consistently even where the technical tools exist.

3.6 Organisational and systemic constraints on decision-making

3.6.1 Cognitive biases do not fully explain the difficulty organisations have in managing risk. Decision-making is also shaped by organisational and systemic factors: structure and processes, cultural norms, incentives, the complexity of coordination, and external conditions. These constraints bind more tightly as uncertainty rises. During strategic shifts, innovation or external shocks, ambiguity widens the scope for narrative-driven decisions, misalignment across units, and delayed escalation of emerging exposures.

3.6.2 Organisational structures often fragment information flows and produce siloed decisions, which limits the integration of different risk perspectives (Dacin et al., 1999). The risk function counters this by operating across silos, using shared methodologies and governance structures to build an enterprise-wide view (Committee of Sponsoring Organizations of the Treadway Commission, 2017), so that risk information is consolidated and presented in a consistent, decision-relevant form. Cultural and social norms can suppress dissent and discourage challenge (Schein, 2010); the function's challenger mandate and its accountability structures are meant to make open discussion of uncertainty safer, and over time to support a culture in which challenge is expected (Power, 2004). Incentives frequently work against prudent risk-taking, rewarding volume, growth, or short-term accounting outcomes while treating tail exposure and capital consumption as secondary (Kirkpatrick, 2009). A value-based framing helps here, by making the capital cost of risk explicit in how options are compared, so that decisions are judged on expected benefit and capital intensity together.

3.6.3 The complexity of coordinating across decentralised units is a further barrier, which the function addresses by standardising measurement and aggregation so that concentrations and systemic risks become visible. Decisions taken in response to regulatory change, technological disruption, or market volatility (Fligstein, 2001) are typically uncertainty-dominant, and point estimates there create false confidence;

horizon scanning, scenario analysis, and stress testing provide the foresight that individual decision-makers lack.

3.6.4 What these constraints show is that improving decisions under risk and uncertainty takes more than correcting individual biases. Embedding the function at decision points, with its independence, methods, challenger mandate and enterprise view, is what addresses the constraints that biases alone cannot explain.

3.7 Risk-sensitive business processes

3.7.1 This section looks at the business processes that expose an insurer to material risk and therefore call for structured involvement from the risk function. They span strategic, technical, financial, operational, and innovation-related domains. The aim is to map where risk input matters most and how to tell whether it is embedded.

3.7.2 Each process has a typical risk profile and specific points where objective risk input is most useful. A further question is whether the decision is typically *risk-dominant*, *uncertainty-dominant*, or mixed. In risk-dominant settings the structure of outcomes is stable enough that probabilities and loss distributions can be treated as reliable for the purpose at hand. In uncertainty-dominant settings the structure is changing, data are thin, and the task shifts from estimation to bounding and learning. Many processes fall between the two, with some components modelled credibly and others structurally uncertain (Knight, 1921; Zweifel, 2021).

3.7.3 What counts as useful risk input follows from this classification. In risk-dominant processes the function adds most through quantified distributional analysis, correlation and accumulation assessment, and capital and liquidity impact estimation. In uncertainty-dominant processes it adds most through scenario framing, explicit safe-to-fail bounds, identification of non-obvious failure modes, and explicit ranges for expected benefit, tail exposure, and capital intensity.

3.7.4 Embedding risk is therefore more than one technique applied everywhere (Kaplan & Mikes, 2012). It depends on the mechanics of each process and on the economics that process determines, so the function has to combine technical capability—measurement, aggregation, scenario design—with knowledge of the process itself: benefit design and demand elasticity in product development, peril geography and reinstatement terms in reinsurance, cash-flow timing in application lifecycle management (ALM). That knowledge improves with experience, as the function learns where decisions are actually made, which assumptions dominate outcomes, and where control or conduct failures tend to arise. Because the approach differs by process, whether risk input is embedded is best tested in context. For any given process, four questions are useful.

- Are there decision checkpoints where proposals are reviewed before commitments are made?
- Is risk appetite made explicit at those checkpoints, as limits, ranges, or safe-to-fail bounds relevant to the decision? (CRO Forum & CRO Council, 2015)

- Does routine reporting cover changes in that process, such as new products, channels, geographies and counterparties, rather than only business as usual?
- And does the material at the checkpoint make the value-risk trade-off explicit?

3.7.5 Positive answers indicate that the function is present at the point where a decision is made, that acceptability is tested against declared tolerances, and that reporting captures forward exposure.

3.7.6 The grouping in Figure 2 is intended to help prioritise where further work on embedding will do most. Priority usually lies with processes that create concentrated exposures: product launches, reinsurance programmes, and major technology and operating-model changes.




FIGURE 2. Risk-sensitive process map: five areas and their core processes

3.7.7 Decisions depend on perceived risk; unaided perception is unreliable; and embedding therefore requires both method and knowledge of the process at the point of decision. Yet survey evidence suggests the CRO is fully involved in strategic decisions in only about half of firms in Europe (Mattone et al., 2024), with no comparable figure published for South Africa. The discussions we held with the CROs we spoke to are consistent with a similar gap in the large South African groups: even where the function was well regarded and routinely present at executive level, its involvement in material decisions often began only once a proposal was substantially formed, usually at the point it reached the executive committee, and the use of risk appetite at the point of decision varied widely between firms. The next section considers why this gap persists.

3.8 Why the CRO is often absent from strategy

3.8.1 Industry surveys and case descriptions suggest that risk leaders are not consistently involved in high-stakes decisions, and that their influence varies even when they attend the relevant meetings. The internal reasons are rarely visible from outside, so the discussion below sets out a range of mechanisms; the discussions described above gave concrete instances of several, which we draw on at the relevant points. The mechanisms fall into four groups. Three concern the function itself: how it is positioned, when it is invited into the process, and what it contributes once involved. The fourth is a feedback loop that can entrench the other three, and because it does, it helps to start there.

3.8.2 A common starting point is a function designed mainly to demonstrate compliance, where most effort goes to policies, controls, formal reports, and second-line assurance. That work is necessary, but it does not by itself produce information that helps decision-makers compare options, quantify trade-offs, or understand the consequences of uncertainty. A function engaged mainly to satisfy governance tends to be involved late and to speak in the language of compliance (Power, 2004; 2007). Decision-makers then experience risk input as a constraint that arrives after choices have been shaped; they reduce early engagement to avoid delay, and seek informal advice elsewhere. Less engagement gives the function fewer chances to learn the decision context, which weakens its contribution further, and the cycle becomes self-reinforcing: late involvement lowers decision relevance, low relevance reduces involvement, and the loss of access stalls the development of capability. The clearest instance in our discussions was a function that had long had limited credibility, where first-line ownership dominated and, over time, internal audit had become the stronger practical voice of independent challenge. That is the end state of the loop. A function can exist on paper and meet regulatory requirements while remaining peripheral to the decisions that shape the firm's risk profile; when it does not occupy the challenge role, another function does, and the compliance framing is entrenched rather than displaced.

3.8.3 Positioning can make the cycle more likely. Some firms read regulation, and 'independence', as a reason for the function to stay at a distance from strategy (CRO Forum, 2021), treating independence as non-involvement rather than independence of judgement. Committee terms of reference may give the risk leader no defined role in shaping decisions, or confine participation to formal approval points. Visibility is a related problem: a function not told early about coming products, partnerships, investments, or operational changes cannot shape them, because it sees the proposal only when the documentation is nearly final. Even good positioning does not settle the matter. One large group we spoke to had strong relationships across the business and a standing seat on its executive committee, so new initiatives reliably reached the function; involvement still tended to begin once an idea was mature enough to be tabled, rather than while it was being formed. Relationships may bring change to light early in smaller organisations, but become less reliable as the organisation grows and activity increases. More formal processes are

then needed to identify risk areas earlier. One example is a simple triage step at the annual budgeting stage, through which planned projects and material expenditure are flagged for a short risk and governance assessment before reaching the executive committee.

3.8.4 When the function is invited in is often decisive. If engagement comes mainly at the end, its role narrows to acceptance or rejection, because there is little time to adjust the design, and input that raises real uncertainty reads as obstruction. Worries about role boundaries reinforce late engagement: decision-makers may avoid early collaboration for fear that the function will be seen as co-owning the decision, or that involvement will create accountability disputes later. Without clear rules separating advice from ownership, the function is kept at arm's length, which again raises the chance of late intervention. One group's experience showed both the value and the cost of timing. Its risk function sat on the steering committee of a project that had begun to go wrong, and intervened, setting risk-appetite conditions that had to be met before each phase could proceed. The intervention worked and the project ran more smoothly afterwards. But it was reactive: the function stepped in once the project was already deteriorating. The episode supports appetite expressed as operational limits, while showing that limits introduced late and from outside the team can be resisted even when they work.

3.8.5 What the function contributes once involved also matters. A team without enough technical depth, commercial understanding, or modelling capability cannot turn uncertainty into the quantities and scenarios a decision needs (Mikes, 2009), so its contributions stay generic: restating policy, or describing risks at a high level without connecting them to outcomes, variability, tail exposure, or failure modes. The capability this reflects, and what it takes to build, are taken up in Section 5. Other limits compound the problem. Reporting built on backward-looking indicators, issue registers, or qualitative heat maps does not answer the questions decision-makers face at the point of choice. Abstract risk-appetite statements that are never translated into operational thresholds turn disagreements into debates about interpretation rather than tests against a shared reference; in one function, appetite was used mainly to detect breaches after the fact rather than to test a decision before it was taken, which left the function with little influence at the point of decision. And where analysis cannot be produced quickly enough, decision-makers conclude without it and treat the later review as a formality.

3.8.6 Across these discussions, formal authority and veto rights were a poor guide to influence. The functions that shaped decisions most did so through credibility, early involvement, and usable risk appetite, and leaned least on the power to stop things; one function with comparatively strong formal powers was among the more compliance-bound. We would not draw a general rule from a small number of discussions, but the pattern fits the broader argument that formal authority is not what gives the function influence over decisions.

3.8.7 The mechanisms point to a simple implication. The absence of the risk leader from decision-making is usually a result of governance design, sequencing, capability,

and information quality, not personality. A function built mainly to satisfy compliance tends to sit downstream of decisions, to arrive late, and to produce outputs poorly matched to what decision-makers need. Changing that takes work in several places at once: clarifying early involvement while preserving independence of judgement, ensuring visibility of the decision agenda, translating appetite into thresholds for decisions, improving analytical capability and turnaround, and integrating risk views across categories so that options can be judged against a coherent picture of exposure.

4. INNOVATION AS A TEST CASE FOR EMBEDDED RISK

4.1 Innovation is a good setting in which to test whether risk management is embedded. It concentrates uncertainty, compresses timelines, requires coordinated decisions across underwriting, distribution, operations, finance, and technology, and carries reputational exposure if execution fails. If the function can add decision-relevant analysis here without slowing delivery, the same behaviours should work in more stable settings (Tushman & O'Reilly, 1996).

4.2 By innovation we mean material changes in what the firm offers, how it offers it, or with whom it operates: new products and variants, new channels or partnerships, and significant changes to operating models or core systems. The argument does not depend on the product; the details of a particular contract or treaty are not the point, and the focus is on how uncertainty is made explicit and bounded while choices are being made. Where innovation is under way, risk management is often seen as a source of delay or complexity, particularly where its involvement comes late or is framed as assurance.

4.3 The section proceeds in two parts. First it identifies a few recurrent obstacles that degrade decision quality in innovation contexts; the list is not exhaustive, but the obstacles follow from the diagnosis above and lie within the remit of the function. Second it describes how the function can be present at the relevant decision points in a way that keeps first-line accountability while supplying analysis that aligns choices with appetite and capacity. The role of risk in innovation is neither to slow change nor to sponsor it, but to make explicit the conditions under which a change is acceptable and the signals that would warrant adjustment (CRO Forum, 2023a).

4.4 Recurrent obstacles in innovation

4.4.1 The first obstacle is procedural delay, caused by a mismatch between how innovation work proceeds and how governance is structured (CRO Forum, 2017). Where engagement is built around end-stage approvals and standardised documentation packs, proposals reach formal forums largely fixed in scope, time pressure compresses the discussion to a yes-or-no decision, and there is little room left to shape the design. Risk input is then read as late and obstructive even when the issues it raises are material.

The underlying problem is usually not disagreement about appetite but the absence of an earlier point at which uncertainty could be raised while options were still open (Cooper & Sommer, 2016).

4.4.2 Ambiguity also invites over-conservatism. Where risk appetite has not been translated into limits for the decision at hand, uncertainty tends to be resolved by caution. Projects are deferred or truncated not because they breach a declared tolerance but because no tolerance has been translated into thresholds, ranges, or safe-to-fail bounds that would permit controlled experimentation, and small, reversible trials get treated as one-way, high-stakes commitments. The opposite can occur under stress, where a loss the firm cannot absorb makes risk-seeking choices attractive in a way that arguments based on expected value do not capture. In both cases, the lack of decision-relevant appetite leaves judgement to circumstance rather than to agreed limits (McGrath & MacMillan, 1995; Snowden & Boone, 2007; Thomke & Manzi, 2014). What matters is not only whether appetite is concrete, but whether the business helped shape it early enough to treat it as its own.

4.4.3 There is often a gap between technical risk analysis and what sponsors need from it. Innovation decisions are made under time pressure, and rely heavily on heuristics and narrative. If risk material is written in the language of control, or is written for completeness rather than for the decision at hand, the forum will not find what it needs at the point of choice. Decision-makers need compact statements of what could happen, how sensitive the outcome is to a few dominant assumptions, and where exposures accumulate if the initiative succeeds.

4.4.4 Analysis may also be disconnected from what the initiative is meant to achieve and from the firm's capacity to absorb loss. Technically sound work is still unhelpful if it is not expressed against both: the outcome sought, whether that is growth in a target segment, a better combined ratio or resilience of service levels, and the volatility the firm can absorb. The same gap appears when routine reporting dwells on retrospective variance and issue counts, leaving the forum without forward-looking scenarios and trade-offs. The analysis then reads as an external compliance check rather than an input to the choice.

4.4.5 Finally, decision-making is often siloed. Innovation typically spans underwriting, pricing, distribution, operations, technology, and finance, and if engagement is fragmented by risk type or organisational boundary, no single view captures the dependencies across the proposal. A vendor choice can shift operational and conduct risk; a product design can change correlation in the book; a core-system change can affect resilience and claims leakage. Without an enterprise lens, local optimisations dominate and the exposures at the seams go unseen. The absence of cross-silo aggregation does not remove the aggregation; it removes the visibility.

4.4.6 This list is not exhaustive, and the obstacles co-exist. Between them they explain why, where uncertainty is concentrated and interdependencies are material, risk input can be present yet have little bearing on the decision.

4.5 Addressing the barriers to innovation

The aim here is to show how the function can take part in innovation work without creating predictable friction. The discussion groups the five obstacles into three responses: earlier engagement, explicit boundaries with the literacy to work inside them, and an enterprise view anchored in the firm's objectives.

4.5.1 *From procedural delay to co-creation*

4.5.1.1 The function should engage early enough to shape assumptions rather than review a near-final proposal. That means rethinking how and when risk management engages with innovation work.

4.5.1.2 Instead of arriving at the end to approve or veto, the team should be involved early to help set the boundaries within which a venture can move: thresholds for capital-at-risk, loss-ratio volatility, or model uncertainty. These limits define where teams can proceed without further sign-off, and where they cannot.

4.5.1.3 This does not require risk professionals in every sprint or meeting; innovation stays business-led. But the process can include predefined checkpoints at which risk input is sought:

- **Initial scoping:** defining the initiative's risk boundaries and safe-to-fail limits in line with appetite.
- **Mid-project check-ins:** reviewing prototypes, pilot results, or early market feedback to test key assumptions and adjust for emerging risks.
- **Pre-launch review:** confirming that residual risks are within tolerance and that scenario analysis has covered the worst cases.
- **Post-launch reflection:** comparing outcomes against the original assumptions and recalibrating for future initiatives.

4.5.1.4 Making risk input part of the process rather than a final hurdle makes the function part of a feedback loop rather than a control gate. The difference is earlier and more specific input: instead of enforcing compliance after the fact, the function helps shape the proposal as it forms, translating appetite into initiative-specific guidance and flagging blind spots without taking ownership away from the business. The question changes: not why an idea should be refused, but what would make it acceptable within the firm's tolerance. Early involvement also tends to speed innovation, since it heads off the late-stage surprises and last-minute vetoes that derail projects shortly before launch.

4.5.2 *Safe-to-fail boundaries and risk literacy*

4.5.2.1 Innovation always involves uncertainty, but not all uncertainty is intolerable risk (Ellsberg, 1961). Safe-to-fail boundaries—a maximum loss, a capital-at-risk cap, a volatility range—mark off tolerable uncertainty from the kind that could threaten solvency, reputation, or long-term value. Within them, business units can experiment with

autonomy, knowing they are inside agreed tolerances; the boundaries guide rather than block.

4.5.2.2 For the boundaries to work, teams need enough risk literacy to manage within them without constant oversight. That involves practical risk scoping (templates or checklists to identify the relevant factors early), simple sensitivity analysis, and an understanding of how the firm's appetite statements translate into limits for their specific project. The function's role shifts toward that of an educator and translator, and the result is better collaboration and less of the mistrust that surrounds the function's involvement.

4.5.2.3 With that understanding in place, the function can contribute as a design partner, refining assumptions, identifying breakpoints and defining mitigations, rather than reviewing a near-final proposal, and can step in early when a team approaches a boundary.

4.5.2.4 As projects push against those limits, scenario-based techniques help teams plan for adverse outcomes:

- **Reverse stress testing:** asking what would have to go wrong for the initiative to fail, identifying the breakpoints, and assessing how plausible they are.
- **Key-assumption sensitivity:** challenging the critical assumptions in the business case. If a product is priced to a 60% loss ratio, what happens at 70% or 75%, and would it still be acceptable?
- **Option mapping:** setting out contingency or pivot strategies if key risks materialise. If the primary approach hits a given risk, what is the fallback?
- **Scenario overlays:** testing how the initiative would hold up under an economic downturn, a different regulatory regime, or extreme climate events.

4.5.2.5 These techniques move teams past base-case optimism and force explicit attention to tail risks without stalling delivery. By testing worst cases, identifying failure modes, and specifying mitigations, downside exposure becomes clearer and the resulting go/no-go decision is easier to defend. Without that perspective, decision-makers tend either to overstate the risk and pass up good opportunities, or to understate it and press ahead blindly. A capable risk function offers a third option: taking the risk deliberately, based on evidence.

4.5.3 *Strategic integration and enterprise risk perspective*

4.5.3.1 Risk analysis is less useful when it is not anchored in the firm's objectives and risk-bearing capacity. Assessment should start from the initiative's purpose and the binding constraints—growth aims, customer outcomes, capital and liquidity limits, stated non-financial commitments—and express results as impacts on those constraints, so that decision-makers can compare options against the firm's capacity rather than in the abstract.

4.5.3.2 The anchor has to be paired with an enterprise view of the risk. Innovation may cut across many areas of the business and exposures may aggregate across

those areas, so silo-by-silo engagement misses the dependencies. A cross-silo view reveals the accumulations a proposal creates: concentration in a segment, reliance on a shared service, coupled failure between a platform change and a product launch. The mechanism is straightforward: convene the relevant specialists and business owners around a single view of how group exposure changes if the initiative proceeds, so that it is judged against the firm's total risk profile rather than within one silo (Star & Griesemer, 1989; Perrow, 1999).

4.5.3.3 The common thread across the three is early engagement. The function's job is not to find reasons an initiative cannot proceed, but to establish what would make it acceptable within appetite, and its input should match the pace of the work through ongoing discussion and light tools such as assumption testing and pre-mortems rather than static assessments. None of this relaxes controls or governance; it changes when the function contributes and what it contributes. Delivering it consistently requires change in how the function works and in what it is able to do, which the next section turns to.

5. TRANSFORMING THE RISK MANAGEMENT FUNCTION

5.1 Risk functions are expected to support innovation but often fail to in practice. Many risk teams have been built around governance skills rather than commercial understanding, which limits their ability to shape solutions. Without a grasp of how processes create value, how operational trade-offs work, and what drives change, a team falls back on standard processes and compliance artefacts, and so confirms the perception of risk as a late-stage constraint. The problem is less one of intent than of skills, mindset, and functional design. Sections 2–3 established that GOI3 provides a sound baseline for system design, documentation, appetite, monitoring, and reporting, but does not prescribe the integrative, decision-focused work this section addresses.

5.2 What a function needs beyond the baseline falls into five areas: mindset and culture, skillset and capabilities, where the firm's models sit, performance measurement, and reputation. Where it is useful, we take each through three questions: where to start, what to sustain, and what resistance to expect. Change attempted everywhere at once rarely works; in each area only a few activities move the function toward earlier engagement, and each predictably raises anxieties about encroachment and role confusion that are better anticipated than handled defensively.

5.3 Mindset and culture

5.3.1 The first change required is one of posture: from policing to enquiry. The traditional model casts the risk professional in a control-and-assurance role, focused on adherence to policies and limits. That stance invites disengagement, and the function ends up being consulted only once a solution is fully designed. The more useful first question is how something works, rather than whether it is compliant or too risky. Approaching an

initiative to understand its intent, its mechanics, and the uncertainties around it turns the function from a judge into a partner in shaping the solution.

5.3.2 Cultural change of this kind has to be anchored at the top. A CRO repositioning the function needs visible backing from the CEO and the chair of the risk committee, stating that earlier and more substantive involvement is expected. Support at that level is what makes the expectation credible further down. It helps to make the direction explicit rather than let it emerge informally. A short charter can state what will change, what will not, in particular that decision ownership stays with the business, and how colleagues should engage. The credible place to start is narrow: one or two processes, such as product design, claims, underwriting, or a single major project, where the function can embed meaningfully and the business is likely to engage. The purpose is to show in a contained setting that early involvement improves feasibility, alignment to appetite, and delivery. Success in one area creates advocates, and their account of how the function clarified trade-offs and reduced late-stage rework carries more weight than any mandate.

5.3.3 The resistance is twofold and predictable. Within the team, colleagues who prefer structured review work may feel exposed in advisory discussions and need coaching and role clarity. Among business leaders, earlier involvement can read as surveillance, especially at the ideation stage. The combination produces friction: the function intends to help but is experienced as intrusive, while the business wants speed and underrates the cost of late surprises. The way through is demonstration rather than argument: consistent behaviour and a few early successes establish that earlier engagement neither shifts accountability nor slows delivery.

5.4 Skillset and capabilities

5.4.1 Enquiry is not enough. To move past generic advice, the function needs business fluency: a working understanding of how processes create value, how operational decisions are made, and what constraints and pressures shape them. Risk professionals are often expert in frameworks, governance, and compliance but have little exposure to commercial decision contexts, which confines their input to procedural enforcement. This fluency is built less by training than by exposure: secondments, shadowing, rotations through underwriting, claims, reinsurance and technology, and a seat in the meetings where trade-offs are actually made.

5.4.2 Fluency alone is also insufficient. The function has to assess risk in the same units and tolerances used to express appetite. Ordinal '1–5' ratings are useful for triage but rarely adequate for insurer decisions that turn on capital, tail risk, and accumulation; assessment should use the quantitative terms the business already works in: loss ratio, reserving and pricing uncertainty, capital-at-risk, tail quantiles, solvency coverage and liquidity. Because modelling underpins pricing, reserving, and capital, the function must engage with actuarial work closely enough to understand model structure, assumptions,

data lineage, and limitations, and to translate firm-level appetite into project- and portfolio-level limits. The core skillset is quantitative risk management across insurance, credit, and market risk: model construction and review, stress and scenario design, dependency and aggregation methods, and interpretation of outputs for setting and monitoring appetite. The cost of lacking it is visible in practice. In one of the functions we spoke to, staffed largely from compliance and internal-audit backgrounds and reliant on the actuarial team for quantitative work, risk assessments took narrative form or amounted to legal and compliance feedback on an initiative, and the function could rarely express a decision as a range of outcomes, tail exposure, or capital impact. A shortage of quantitative skills is partly a structural feature of the South African market, where skilled actuarial resources are scarce and expensive. Capability is also built by doing. One large group has built deep understanding of its products because the risk function does the modelling itself. Setting the assumptions for a new product requires understanding the product, its distribution, and the existing book, and shadowing alone does not produce that.

5.4.3 The credible starting point is an honest inventory: plot the team on two dimensions, qualitative versus quantitative strength and depth of business understanding. That usually shows strength in governance and assurance, and weakness in interpreting models, in scenario thinking and in operational understanding. What deserves sustained attention is range rather than uniformity: all staff conversant at a conceptual level with the firm's core models, with some roles going further to quantify scenarios and challenge assumptions. Hiring follows from this. Where risk teams recruit mainly from audit, compliance, and technical-risk backgrounds, they reinforce the very limits that keep the function from partnering with the business; the profile should widen to include people with frontline experience in product, claims, operations, technology, or strategy, and development paths for existing staff should run through operational rotations, aiming at fluency in the language of both risk and the business. The resistance to expect is a concern about boundaries: that a more numerate risk function becomes a second actuarial team. The balance is for the risk function to be numerate enough to interrogate assumptions and translate findings into decision language while the actuarial function keeps primary modelling, which preserves independence and avoids duplication.

5.5 Where the models sit

5.5.1 A function's capability and its access to decisions both depend on where the firm's models are owned. In most insurers the models that matter for decisions—pricing, reserving, capital—sit in actuarial or business teams, and the risk function consumes their outputs. The arrangement is workable, but it leaves the function dependent on analysis prepared elsewhere, and it means the business does not need to involve the risk function to get the numbers a decision requires.

5.5.2 One of the groups we spoke to takes a different approach. Its risk function owns the firm's principal models, so that a new product, a reinsurance structure,

or a change in strategic asset allocation cannot be evaluated without coming through the risk function for the figures. Ownership of this kind turns involvement from something that has to be arranged into something structural: the business comes to the function first because it holds the tools that produce the decision-relevant numbers, and the function builds deep knowledge of the business through the modelling itself. The cost is a conflict of interest, since a function cannot independently review models it has built. The conflict is real, but third-line and external review can address it.

5.5.3 For firms unwilling to undertake a change of that scope, a narrower version captures much of the benefit at far lower cost. Moving the capital model into the risk function, rather than every model, places in the function the one model that nearly every material decision depends on, since the capital consequences of a product, a reinsurance structure, or an investment strategy are what most often determine whether it is acceptable. The conflict that comes with full model ownership is largely avoided, because the independent review of the capital model stays with the head of the actuarial function, as the Prudential Standard requires. For the model that matters most, then, building and review are already separated by regulation, so the risk function can own the model while assurance over it sits elsewhere. Owning the enterprise capital and aggregation view this way also supports the two roles the standard leaves least specified: it puts the function at the centre of risk aggregation, the Integrator role, and gives it a structural seat at material decisions, the Enabler role.

5.6 Measuring risk function performance

5.6.1 How the function is measured shapes how it behaves. A function judged only on compliance and process metrics will default to a policing role, so sustaining the shift toward decision support requires measures that reflect it. Useful indicators capture whether the function engages at the right time, whether it contributes to decisions within agreed boundaries, and whether it is building the fluency and credibility to take part meaningfully: whether it was involved before options hardened, whether trade-offs between risk, return, and capital were made explicit, whether decisions aligned with appetite, and whether the function is embedded across the processes in Section 3.7 with clear appetite for each. These indicators depend on feedback: post-decision reviews, and short structured reflections from executives after major initiatives. These often reveal whether the function clarified the uncertainty and helped prioritise mitigations, or merely confirmed what was already known.

5.6.2 The credible place to start is small: a limited set of indicators applied to a single project or decision stream, used to observe when the function becomes involved, how that is perceived, and whether it shapes the options put to decision-makers. Output counts—reports issued, policies updated, issues tracked—remain necessary to show the governance machinery works, but they do not describe contribution to decision quality and should not be mistaken for it.

5.6.3 What deserves sustained attention is restraint about attribution. Decision outcomes depend on many variables and show up only over long horizons, so trying to attribute success or failure directly to risk interventions invites spurious precision, while falling back on purely administrative measures abandons the point of the exercise. The pragmatic stance is to accept approximate indicators that support learning rather than judgement. Over time a pattern tends to emerge: where the function is invited earlier, discussion quality improves; where measurement stays compliance-oriented, engagement plateaus. Treating measurement as an evolving tool lets the function refine its contribution without over-claiming its impact.

5.7 Reputation and perception

5.7.1 Reputation shifts only when behaviour and experience change together: it depends on the timing of engagement, the quality of the analysis, and the practicality of the controls proposed, shown consistently day to day. Relationship-building is part of the job. Understanding the pressures and constraints operational teams face, and being present where ideas are shaped, is what lets the function contribute as a credible source of independent analysis while keeping independence of judgement. What sustains the shift is consistency. When it arrives informed, frames uncertainty as choices, and avoids reflexive vetoes, advocates emerge who speak from experience, and their voices carry further in leadership forums than abstract arguments about embedding risk.

5.7.2 The resistance to expect is that early missteps confirm the stereotype the function is trying to shed: that risk management delays progress and offers generic caution disconnected from operational reality. Informal narratives travel fast and credibility is harder to rebuild than to establish, which makes sequencing matter: choose allies carefully, prepare thoroughly, and communicate outcomes well, so that as successful cases accumulate, involving the function earlier becomes something leaders choose to do because it visibly reduces avoidable failure.

6. CONCLUSION

6.1 This paper has argued for a clear and practical role for risk management as decision support in insurance, with the aim of improving decision quality by making uncertainty, downside outcomes, and key assumptions explicit at the decision point.

6.2 The mechanisms that deliver this are set out above: decision checkpoints that ask the same questions each time, risk appetite translated into limits and review triggers for experimentation, and measures that track the timing and influence of risk input rather than the volume of documents produced. The contribution of the paper is to connect these mechanisms to the regulatory baseline set by GOI 3, and to locate decision support in the Integrator and Enabler work that the standard leaves least specified.

6.3 For the role to be real rather than nominal, it has to be visible in how the organisation behaves. The relevant evidence is not the volume of reports and risk committee packs but the timing and quality of support to decisions: earlier engagement on material decisions, clearer statements of uncertainty and downside, fewer late surprises, and better-documented reasons for accepting or rejecting a risk. Short feedback tools, such as post-decision reviews and focused surveys of decision owners, can test whether the input was understandable, timely, and changed the decision in a useful way.

6.4 The value of risk management as decision support comes from structured checkpoints, clear limits for experimentation, and simple measures of whether decisions became more disciplined. With those in place, risk management becomes a practical support to decisions rather than a record of them.

DISCLAIMER

The views expressed are those of the authors and not of their employers.

APPENDIX A
COMPARISON OF RESPONSIBILITIES WITH GOI 3

Table 1 compares each of the 14 responsibilities of the risk management function with GOI3’s provisions and assesses the degree of alignment. The role column follows the grouping used in Section 2.

TABLE 1. Responsibilities of the risk management function compared with the requirements of GOI 3

#	Role	Responsibility	GOI 3 Coverage Summary	Assessment	GOI 3 Ref.
1	Architect	Methodologies and tools	Requires a documented risk management system, but does not prescribe specific quantitative or qualitative tools such as rating scales, stress tests, or risk models.	Partially: Methodology requirement is vague; no expectations for consistency or back-testing of assessments.	5.3(e), 11.2(c)
2	Architect	System design and governance	Governance and structure of the risk system are emphasised, particularly board oversight.	Sufficient: Explicitly described responsibilities and expectations in multiple sections.	2.1, 5.1, 11.2(a)
3	Architect	Documentation and institutional knowledge	Requires policy documentation and system architecture records, but not detailed logs of rationale, decision trails, or institutional memory preservation.	Partially: Missing expectations for retaining decisions, insights, and risk learning over time.	5.3(e), 11.2
4	Controller	Risk ownership and accountability	Mentions roles and responsibilities but leaves accountability frameworks for specific risks to interpretation.	Partially: Missing clear risk ownership assignment to business units and specific accountability mechanisms.	5.3(b), 11.1
5	Controller	Risk appetite framework	Requires a documented risk appetite statement with quantitative thresholds for each risk type and regular review.	Sufficient: Comprehensive requirements for risk appetite development and maintenance.	5.2(a), 5.4
6	Controller	Risk monitoring and assessment	Board-level oversight of risk exposures and escalation of breaches is clearly outlined.	Sufficient: Monitoring frequency, responsibilities, and thresholds are appropriately defined.	5.5, 11.2(b)
7	Controller	Materiality and prioritisation	Material risks are referenced but there is no formal requirement for a prioritisation framework or for de-prioritisation of immaterial risks.	Not explicit: Lacks structure for materiality-based risk selection or filtering.	—
8	Controller	Clear and actionable reporting	Reporting to the board and executive is required but lacks detail on clarity, format, or actionability.	Partially: Management information system requirements exist but lack guidance on decision relevance or escalation content.	5.5, 11.2(b)

9	Integrator	Risk aggregation and correlation	Requires enterprise-wide identification of risks but not specific aggregation across silos or correlation analysis.	Partially: Concentration risk addressed but lacks systemic aggregation methodologies.	5.3(a), 5.4(b)
10	Integrator	Forward-looking risk view	ORSA requirements address future outlooks, scenario analysis, emerging risks, and solvency impact.	Sufficient: Strong future-oriented requirements within the ORSA framework.	7.2
11	Integrator	Risk culture and awareness	GOI 3 mandates promoting a sound risk culture and awareness, including training and communication.	Sufficient: Clear expectations for embedding risk culture across the organisation.	5.3(f)
12	Enabler	Independent decision support	The risk function is to 'assist' senior management and the board; involvement in strategic decisions is not prescribed.	Not explicit: No guidance on supporting business-line or product-level decision-making.	—
13	Enabler	Constructive challenge function	No language on formal challenge rights or independent challenge processes.	Not explicit: The risk function has no defined role in testing or contesting assumptions.	—
14	Enabler	Capital efficiency and optimisation	Capital planning is addressed in the ORSA, but optimisation and risk-adjusted allocation roles for the risk function are absent.	Not explicit: No role for risk management in guiding risk-adjusted capital allocation or capital efficiency.	7.2 (indirect)

REFERENCES

- AIR Worldwide & Artemis.bm (2017). Andrew-Like Hurricane Could Cause \$6.3bn Cat Bond Loss. <https://www.artemis.bm/news/andrew-like-hurricane-could-cause-6-3bn-cat-bond-loss-air-worldwide/>
- Board of Governors of the Federal Reserve System (2023). Review of the Federal Reserve's Supervision and Regulation of Silicon Valley Bank. <https://www.federalreserve.gov/publications/files/svb-review-20230428.pdf>
- Chmielowiec-Lewczuk, M (2015). Value Based Management in Insurance Companies. *Wiadomości Ubezpieczeniowe*, no. 4: 17–27
- Committee of Sponsoring Organizations of the Treadway Commission (COSO) (2017). Enterprise Risk Management—Integrating with Strategy and Performance
- Cooper, RG & Sommer, AF (2016). Agile–Stage–Gate Hybrids: The Next Stage for Product Development. *Research-Technology Management*, 59(1), 21–29. <https://doi.org/10.1080/08956308.2016.1117317>.
- CRO Forum (2012). CRO Forum Paper on the Own Risk and Solvency Assessment (ORSA): Leveraging Regulatory Requirements to Generate Value. <https://thecroforum.org/cro-forum-paper-on-the-own-risk-and-solvency-assessment-orsa/>
- CRO Forum (2017). A Guide to Defining, Embedding and Managing Risk Culture. <https://thecroforum.org/a-guide-to-defining-embedding-and-managing-risk-culture/>
- CRO Forum (2021). The Three Lines Model. <https://thecroforum.org/the-three-lines-model/>
- CRO Forum (2023a). Embedding Digitalization into Governance, Risk Management and Compliance Activities. <https://thecroforum.org/embedding-digitalization-into-governance-risk-management-and-compliance-activities/>
- CRO Forum (2023b). ORSA Stress and Scenario Testing – Best Practice for Assessing Risks. <https://thecroforum.org/orsa-stress-and-scenario-testing-best-practice-for-assessing-risks/>
- CRO Forum & CRO Council (2013). Establishing and Embedding Risk Appetite: Practitioners View. <https://thecroforum.org/establishing-and-embedding-risk-appetite-practitioners-view-2/>
- CRO Forum & CRO Council (2015). Risk Appetite – Survey Results. <https://thecroforum.org/risk-appetite-survey-results/>
- Dacin, MT, Ventresca, MF & Beal, BD (1999). The Embeddedness of Organizations: Dialogue & Directions. *Journal of Management*, 25, 1–7
- Ellsberg, D (1961). Risk, Ambiguity, and the Savage Axioms. *Quarterly Journal of Economics*, 75(4), 643–69. <https://doi.org/10.2307/1884324>
- European Insurance and Occupational Pensions Authority (2015). Guidelines on Own Risk and Solvency Assessment (ORSA). EIOPA-BoS-14/259 EN. <https://op.europa.eu/en/publication-detail/-/publication/354pCFYpI2WT>
- Financial Crisis Inquiry Commission (2011). The Financial Crisis Inquiry Report. U.S. Government Printing Office. <https://fraser.stlouisfed.org/title/financial-crisis-inquiry-commission-4967/financial-crisis-inquiry-report-draft-530180>

- Financial Services Authority (2011). The Failure of the Royal Bank of Scotland: Financial Services Authority Board Report. <https://www.fca.org.uk/publication/corporate/fsa-rbs.pdf>
- Fligstein, N (2001). *The Architecture of Markets: An Economic Sociology of Twenty-First-Century Capitalist Societies*. Princeton University Press
- House of Commons Work and Pensions Committee & House of Commons Business, Energy and Industrial Strategy Committee (2018). Carillion (Second Joint Report of the Work and Pensions and Business, Energy and Industrial Strategy Committees). HC 769. UK Parliament. <https://publications.parliament.uk/pa/cm201719/cmselect/cmworpen/769/769.pdf>
- Insurance Information Institute (n.d.). Hurricane Andrew Facts & Statistics. <https://www.iii.org/fact-statistic/facts-statistics-hurricanes>
- International Association of Insurance Supervisors (2019). Insurance Core Principles (ICP)—Notably ICP 8 and ICP 16. <https://www.iaisweb.org/publications/insurance-core-principles-icps>
- International Monetary Fund, Monetary and Capital Markets Department (2008). Global Financial Stability Report, April 2008: Containing Systemic Risks and Restoring Financial Soundness. <https://doi.org/10.5089/9781589067202.082>
- ISO 31000:2018 Risk Management—Guidelines (2018). <https://www.iso.org/standard/65694.html>
- Kahneman, D & Tversky, A (1979). Prospect Theory: An Analysis of Decision Under Risk. *Econometrica*, 47(2), 263–91. <https://doi.org/10.2307/1914185>
- Kaplan, RS & Mikes, A (2012). Managing Risks: A New Framework. *Harvard Business Review*, 90(6), 48–60
- Kirkpatrick, G (2009). The Corporate Governance Lessons from the Financial Crisis. *OECD Journal: Financial Market Trends*, 2009: 61–87
- Knight, FH (1921). *Risk, Uncertainty and Profit*. Houghton Mifflin Company
- Kraus, C (2013). EVA/RAROC vs. MCEV Earnings: A Unification Approach. *The Geneva Papers on Risk and Insurance – Issues and Practice* 38(1), 113–36. <https://doi.org/10.1057/gpp.2012.5>
- Kunreuther, HC, Pauly, MV & McMorrow, S (2013). *Insurance and Behavioral Economics: Improving Decisions in the Most Misunderstood Industry*. Cambridge University Press. <https://doi.org/10.1017/CBO9781139050319>
- Kunreuther, H, Burby, RJ, Black, RG et al. (1998). *Paying the Price: The Status and Role of Insurance Against Natural Disasters in the United States*. National Academies Press. <https://nap.nationalacademies.org/read/5784/chapter/6>
- Lloyd's (2019). Market Bulletin Y5258: Providing Clarity for Lloyd's Customers on Coverage for Cyber Exposures. <https://assets.lloyds.com/assets/y5258-providing-clarity-for-lloyd-s-customers-on-coverage-for-cyber-exposures/1/Y5258%20-%20Providing%20clarity%20for%20Lloyd%E2%80%99s%20customers%20on%20coverage%20for%20cyber%20exposures.pdf>

- Mattone, D, Pancaldi, L, Jurisic, M & Kaposztas, D (2024). Elevating the Risk Function in Insurance: Building a Strategic Advantage. McKinsey & Company. <https://www.mckinsey.com/capabilities/risk-and-resilience/our-insights/elevating-the-risk-function-in-insurance-building-a-strategic-advantage>
- McGrath, RG & MacMillan, IC (1995). Discovery-Driven Planning. *Harvard Business Review*, 73(4), 44–54
- Mikes, A (2009). Risk Management and Calculative Cultures. *Management Accounting Research*, 20(1), 18–40
- Office of the Comptroller of the Currency (2016). OCC Assesses Penalty Against Wells Fargo, Orders Restitution for Unsafe or Unsound Sales Practices. <https://www.occ.gov/news-issuances/news-releases/2016/nr-occ-2016-106.html>
- Perrow, C (1999). *Normal Accidents: Living with High-Risk Technologies*. Updated. Princeton University Press
- Power, M (2004). *The Risk Management of Everything: Rethinking the Politics of Uncertainty*. Demos
- Power, M (2007). *Organized Uncertainty: Designing a World of Risk Management*. Oxford University Press
- President's Working Group on Financial Markets (1999). Hedge Funds, Leverage, and the Lessons of Long-Term Capital Management. <https://www.treasury.gov/resource-center/fin-mkts/Documents/hedgefunds.pdf>
- Prudential Authority (South Africa) (2018a). GOI 3: Risk Management and Internal Controls for Insurers. <https://www.masthead.co.za/wp-content/uploads/2019/01/Prudential-Standard-GOI-3-Risk-Management-and-Internal-Controls.pdf>
- Prudential Authority (South Africa) (2018b). GOI 3.1: Own Risk and Solvency Assessment (ORSA) for Insurers. <https://www.masthead.co.za/wp-content/uploads/2019/01/Prudential-Standard-GOI-3.1-ORSA.pdf>
- Registrar of Long-term Insurance and Registrar of Short-term Insurance (2014). Governance and Risk Management Framework for Insurers, 2014 (BN 158). Financial Services Board (South Africa). https://www.saflii.org/za/legis/consol_reg/garmffi2014498/
- Schein, EH (2010). *Organizational Culture and Leadership*. 4th ed. Wiley
- Slovic, P (1987). Perception of Risk. *Science*, 236(4799), 280–85. <https://doi.org/10.1126/science.3563507>
- Snowden, DJ & Boone, ME (2007). A Leader's Framework for Decision Making. *Harvard Business Review*, 85(11), 69–76
- Star, SL & Griesemer, JR (1989). Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907–39. *Social Studies of Science*, 19(3), 387–420. <https://doi.org/10.1177/030631289019003001>
- Thomke, S & Manzi, J (2014). The Discipline of Business Experimentation. *Harvard Business Review*, 92(12), 70–79

- Tushman, ML. & O'Reilly, CA (1996). The Ambidextrous Organization: Managing Evolutionary and Revolutionary Change. *California Management Review*, 38(4), 8–30
- Tversky, A & Kahneman, D (1974). Judgment Under Uncertainty: Heuristics and Biases. *Science*, 185(4157), 1124–31. <https://doi.org/10.1126/science.185.4157.1124>
- U.S. Department of the Treasury (2012). AIG Program Status. <https://home.treasury.gov/data/troubled-assets-relief-program/aig/status>
- U.S. House Committee on Oversight and Government Reform (2018). The Equifax Data Breach. <https://oversight.house.gov/sites/democrats.oversight.house.gov/files/documents/2018-12-10.EE.%20Equifax%20Report.pdf>
- Valukas, AR (2010). Report of Anton R. Valukas, Examiner: *In Re Lehman Brothers Holdings Inc., et al.* (Chapter 11 Case No. 08-13555), Volume 1. United States Bankruptcy Court, Southern District of New York. <https://web.stanford.edu/~jbulow/lehmandocs/VOLUME%201.pdf>
- Zweifel, P (2021). Bridging the Gap Between Risk and Uncertainty in Insurance. *The Geneva Papers on Risk and Insurance—Issues and Practice*, 46(2), 200–13. <https://doi.org/10.1057/s41288-021-00220-y>