



★ S2 – Issue One, Operational Risk Appetite — Start with Basic Economics

Series Two: Improving Decision Quality in Uncertain Environments

Many organisations struggle to articulate their appetite for operational risk losses.

Workshops are held. Heat maps are debated. Risk scales are endlessly recalibrated. Eventually, a framework emerges that appears sophisticated — but proves difficult to explain, implement or defend consistently.

Over time, operational risk appetite begins to resemble a black art.

And that is where many organisations go wrong.

If we accept that operational risk appetite is primarily concerned with loss avoidance — while strategic risk appetite considers both opportunity and uncertainty — the concept becomes far simpler than many organisations assume.

At its core, operational risk management is simply economics applied to uncertainty.

Core Insight

Operational risk loss appetite should be anchored in basic economics — not abstract theory.

At its simplest, risk management is about deciding when it makes economic sense to mitigate a risk and when it is rational to live with the exposure. Once this principle is accepted, operational risk appetite becomes far easier to construct and defend.

The starting point is expected loss:

Expected Loss = Consequence × Frequency

At its core, operational risk appetite is simply an optimisation exercise between expected loss and mitigation cost.

The benefit of mitigation is the avoided expected loss. The cost is the investment required to identify, monitor and mitigate the exposure. This includes controls, governance processes, reporting, technology investment and second line oversight.

Importantly, these estimates do not need to be precise. Expected operational risk loss outcomes typically follow broad probability distributions, and the magnitude of the outcome will shape the decision.

If the risk analysis is initially restricted to direct financial loss outcomes, this creates what can be described as a risk-neutral risk matrix.

If expected loss exceeds mitigation cost, investment is economically rational. If mitigation costs exceed expected loss, the organisation can rationally tolerate the exposure.

This is the purely economic view of risk — the position a machine might take if it evaluated decisions solely on expected value. It represents the point before human judgement, behavioural bias and organisational risk aversion are introduced.

This becomes the anchor point for the framework.

From there, three additional refinements are required to construct a practical operational risk appetite framework.

- Step two is calibrating other forms of consequence.
 - Step three is separating operating constraints.
 - And the final step is adjusting for risk aversion.
-

Calibrating Consequences

Operational risk frameworks rarely consider financial impact alone. Most organisations also assess customer impact, regulatory consequences, reputational damage, employee harm and operational disruption.

This creates an important challenge.

Different consequence types must ultimately compete for management attention and mitigation investment. If they are not calibrated consistently, risk appetite quickly becomes subjective and difficult to apply across the organisation.

Boards therefore need to undertake an equivalence exercise.

For each consequence category, organisations should define objective and measurable impact levels aligned to equivalent levels of financial impact. If consequence definitions rely primarily on qualitative language, appetite will inevitably be interpreted differently across business units and leadership teams. This is one of the primary reasons operational risk appetite often becomes inconsistently applied across large organisations.

Once consequence equivalence has been established, the original risk-neutral matrix can be applied consistently across all consequence types — not just direct financial loss.

Separating Operating Constraints

As organisations work through the above calibration exercise, there is often a moment where someone says:

“You simply can’t put a number on that.”

This does not invalidate the approach. It simply means some outcomes are not acceptable to the organisation.

The death of an employee, deliberate misconduct or knowingly breaching legal obligations are not conventional risk-return trade-offs. They are operating constraints. They sit outside conventional economic optimisation.

And operating constraints should be treated differently.

Rather than forcing them awkwardly into a risk matrix, they should be explicitly articulated within the Risk Appetite Statement as hard boundaries or non-negotiable obligations.

Adjusting for Risk Aversion

This is where operational risk appetite becomes more nuanced.

Consider a simple example.

If someone offered to flip a coin with you for one dollar, most people would take the bet. The expected value is neutral — there is an equal probability of winning and losing.

But what happens if the stakes increase to \$100? Or \$10,000? Or \$1 million? The mathematics remain unchanged. The probability distribution is identical. Yet most people become increasingly uncomfortable as the stakes rise.

Why?

Because humans are generally risk averse.

The pain associated with large downside outcomes begins to outweigh the equivalent upside benefit. Behavioural economics describes this as asymmetric utility — where losses are experienced more intensely than gains.

Organisations behave similarly. The difference is that organisational risk aversion is ultimately expressed through board decisions, capital allocation and control investment.

Most mature organisations are also risk averse. As potential consequences become more severe, boards will often choose to invest disproportionately in mitigation, even where the expected loss calculation alone may not justify it.

Therefore, the operational risk appetite also needs to capture where organisational risk aversion begins to meaningfully influence decision-making.

This is where genuine operational risk appetite begins to emerge.

What This Means for Executives

For executives, this approach delivers several important benefits.

First, operational risk appetite becomes commercially grounded rather than theoretically constructed. Discussions move away from abstract heat maps toward practical economic trade-offs and consequence management. This materially improves the quality of executive and board risk discussions.

Second, the framework becomes easier to explain and apply. Managers can more clearly understand why some exposures warrant aggressive mitigation while others can rationally be tolerated.

Third, it improves consistency across the organisation.

When consequence categories are calibrated properly and risk aversion is explicitly understood, decision-making becomes more aligned across the organisation.

Most importantly, the framework begins to reflect the board's genuine tolerance for risk and consequence — rather than generic compliance language.

When operational risk appetite is constructed this way, it stops being a governance artefact.

It becomes a practical decision-making tool for running the organisation.

Practical Takeaways

Executives seeking to strengthen operational risk appetite should focus on four core principles.

Start with the economics of risk rather than the mechanics of frameworks. Operational risk management is fundamentally about understanding when mitigation preserves more value than tolerance.

Calibrate consequence types carefully. Ensure financial, customer, regulatory, reputational and operational consequences are objectively defined and consistently equivalent across the framework.

Separate operating constraints from risk trade-offs. Some exposures are simply not acceptable and should be treated as hard boundaries rather than risks to optimise.

Finally, recognise that operational risk appetite is ultimately shaped by risk aversion. Boards need to understand where potential downside begins to justify investment beyond expected loss calculations alone.

Once these elements are in place, operational risk appetite becomes far simpler than many organisations assume.

It is not a black art.

It is simply economics applied to uncertainty.