


BEYOND THE MODEL

When the Spreadsheet Lied

How the 2008 Crisis Exposed the Governance Gap Between Risk Models and Executive Accountability



Giovanni Leonardi

March 2009

When the Spreadsheet Lied

Transformation · No. 158

Originally written March 2009 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“The spreadsheet answered the question before it was asked — and the spreadsheet lied, not through malice, but through the structural impossibility of a backward-looking model predicting a forward-looking catastrophe.”

1. The Comfortable Fiction

By the spring of 2007, the global financial system had perfected a remarkable trick: it had learned to hide catastrophic risk inside the appearance of mathematical certainty.

Risk models — Value at Risk calculations, Monte Carlo simulations, correlation matrices built on a decade of benign market data — had become the lingua franca of board-level assurance. When a non-executive director asked whether the institution's exposure was manageable, the answer came back as a number. A clean number. A number with decimal places. A number that implied a precision so fine-grained that questioning it felt like questioning arithmetic itself.

This was the comfortable fiction. Not that the models were wrong — models are always wrong to some degree — but that the models had become a substitute for judgement rather than an input to it. The spreadsheet had stopped being a tool and had started being an authority.

What the 2008 crisis exposed was not a failure of mathematics. It was a failure of governance. The gap was not between the model and reality — that gap is permanent and well understood by any competent quant. The gap was between the model and the executive who relied on it, between the output and the accountability that should have surrounded it.

2. Three Ways the Governance Failed

2.1 The Delegation of Understanding

The first failure was structural. As risk models grew more sophisticated through the early 2000s, the people accountable for risk decisions — board members, executive committees, senior leadership teams — progressively delegated not just the modelling but the understanding of what the models meant.

This is a pattern that recurs far beyond financial services. In any organisation where technical complexity increases faster than leadership literacy, a dangerous gap opens. The decision-maker sees the output. The decision-maker does not see the assumptions. The decision-maker cannot challenge the methodology because they no longer speak the language.

In regulated financial institutions before 2008, this manifested as boards that received risk reports showing green across every metric — because the metrics had been designed to show green under the conditions that had prevailed for the previous decade. The models were calibrated on a world that had not yet broken. When the world broke, the models broke with it. But the governance structures that should have asked "what happens if the world breaks?" had long since stopped asking, because the spreadsheet said everything was fine.

2.2 The Confusion of Precision with Accuracy

The second failure was epistemic — a category error that pervaded not just risk management but programme governance, portfolio assurance, and benefits realisation across every sector.

Precision is the number of decimal places. Accuracy is whether the number is pointing at the right thing. A model can be extraordinarily precise — Value at Risk to the nearest thousand pounds, updated daily, reported to four significant figures — and simultaneously miss the entire category of risk that will destroy the institution.

Before the crisis, the industry had built an elaborate infrastructure of precise measurement. Correlation assumptions were specified to basis-point accuracy. Liquidity risk was modelled on historical trading volumes. Counterparty exposure was netted and collateralised according to formulae that assumed continuous market function. Every number was precise. Almost none were accurate in the scenario that mattered.

The governance failure here was not that boards were given bad numbers. It was that boards were given numbers without context. Without the question: "What are the three things this model cannot see?" Without the challenge: "Under what conditions does this entire framework become meaningless?"

This is not a failure of risk management. It is a failure of the governance layer above risk management — the layer whose job is not to run the model but to interrogate it.

2.3 The Accountability Vacuum

The third failure was political. When risk became a technical function — staffed by specialists, reported through dashboards, governed by committees that reviewed metrics rather than challenged assumptions — accountability became diffuse to the point of absence.

Who was accountable for the risk position of a major bank in 2007? The Chief Risk Officer would say the models were sound. The CFO would say the positions were within limits. The CEO would say the board had reviewed and approved the risk appetite. The board would say management had assured them. Each layer pointed to the layer that had provided the number. Nobody pointed to themselves.

This is the accountability vacuum that governance structures are supposed to prevent. The entire purpose of a governance framework — whether in financial services, programme delivery, or portfolio management — is to ensure that someone with authority is asking the right questions and taking ownership of the answers. When the governance framework instead becomes a mechanism for distributing accountability so thinly that nobody holds it, it has failed at its most basic function.

3. The Deeper Pattern

What makes the 2008 crisis instructive beyond financial services is that the governance failure it exposed is not unique to banking. It is a universal pattern that appears wherever three conditions converge:

Technical complexity exceeds leadership comprehension. When the people making decisions cannot independently assess the tools informing those decisions, they become dependent on assurance from specialists who have no incentive to highlight the limitations of their own work.

Measurement replaces judgement. When dashboards, RAG statuses, and quantified risk metrics become the primary mechanism of governance, the governance layer stops thinking and starts reading. The meeting becomes a review of numbers rather than a challenge of assumptions.

Accountability is structural rather than personal. When governance is designed around committees, terms of reference, and escalation procedures rather than around named individuals who own specific decisions, the system optimises for process compliance rather than outcome ownership.

These three conditions exist today in programme governance, portfolio assurance, benefits realisation, and digital transformation oversight across every sector. The models are different — Earned Value rather than Value at Risk, benefits maps rather than correlation matrices, RAG statuses rather than credit ratings — but the governance failure mode is identical.

A programme board that reviews a green RAG status without asking "what would turn this red?" is making the same error as a bank board that reviewed a VaR number without asking "what would make this meaningless?" The spreadsheet is different. The lie is the same.

4. What Good Governance Actually Requires

The lesson of 2008 is not that models are dangerous. Models are essential. The lesson is that governance must be designed to sit above and around the model — not downstream of it.

This means three things in practice.

First, the governance layer must own the assumptions, not just the outputs. Every model, every dashboard, every RAG status is built on assumptions. The governance layer's job is to surface, challenge, and periodically stress-test those assumptions. If a board or steering committee cannot articulate the three most important assumptions underlying the numbers they are reviewing, they are not governing — they are spectating.

Second, precision must be accompanied by limitation statements. Every quantified risk, every forecast, every benefits projection should come with an explicit statement of what it cannot see. Not buried in methodology appendices — visible, prominent, discussed. "This VaR calculation assumes liquid markets and stable correlations. It does not model systemic contagion or market seizure." "This benefits forecast assumes full adoption by month six. It does not account for organisational resistance or parallel change fatigue."

Third, accountability must be personal and named, not structural and distributed. Committees do not make decisions — people do. Governance frameworks that attribute decisions to committees rather than individuals create the accountability vacuum that allows catastrophic risks to be collectively owned and individually ignored.

5. The Question That Wasn't Asked

In every post-mortem of the 2008 crisis — the Turner Review, the Walker Report, the US Financial Crisis Inquiry Commission — the same finding recurs: the information was available. The risks were not invisible. The models had known limitations that were documented, sometimes extensively, in technical appendices that governance layers never read.

The failure was not information. The failure was interrogation. Nobody with the authority to act asked the question that would have surfaced the risk. Not because the question was difficult. Because the governance structure had evolved to make it unnecessary. The spreadsheet answered the question before it was asked. And the spreadsheet lied — not through malice, but through the structural impossibility of a backward-looking model predicting a forward-looking catastrophe.

The uncomfortable question for every organisation running complex programmes, portfolios, and transformation initiatives today is not whether their models are accurate. It is whether their governance is designed to catch the moment when the models stop being accurate — or whether, like the banking sector in 2007, they have built a system that will tell them everything is fine right up until the moment it isn't.