

BEYOND THE GOVERNANCE MANUAL

Governance as Decision Architecture — What the Textbooks Leave Out

A Practitioner's Perspective on the Gap Between Portfolio Governance Theory and the Reality of Organisational Decision-Making

Giovanni Leonardi

February 2024



Governance as Decision Architecture — What the Textbooks Leave Out

Portfolio · No. 107

Originally written February 2024 | Re-edited and re-rendered for the WhereBeyond Edition,
2025

"The portfolio that governs well is not the one with the most comprehensive framework — it is the one where the people in the room understand what they are there to decide."

The Gap Between the Manual and the Room

Portfolio governance, as the textbooks describe it, is a rational system. Investments are prioritised against strategic objectives. Resources are allocated to the highest-value opportunities. Performance is monitored through clear metrics. Underperforming initiatives are identified and addressed. The portfolio evolves as strategy evolves, guided by a governance structure that ensures coherence, accountability, and informed decision-making.

The reality, in my experience, bears little resemblance to this description.

What actually happens in most portfolio governance forums is a ritual of reporting that has been refined over years into something approaching performance art. Each programme presents its status. The RAG ratings are noted. The risks are acknowledged. Questions are asked that are pointed enough to demonstrate engagement but rarely sharp enough to force a genuine reckoning with whether the portfolio is pursuing the right things in the right way.

The textbooks leave this gap unnamed. They describe what governance should achieve without examining the organisational forces that prevent it from doing so. They prescribe decision-making frameworks without acknowledging that the decision most portfolio boards avoid is the one that matters most: which of these things should we stop doing?

Why Stopping Is the Hardest Decision

Portfolio governance exists, in theory, to ensure that the organisation's investment in change is directed toward the outcomes that matter most. This implies a willingness to reallocate — to shift resources from initiatives that are no longer strategically aligned, that are underperforming, or that have been overtaken by changing circumstances.

In practice, stopping a programme is the decision that portfolio boards are least willing to make. Every initiative in the portfolio has a sponsor, a constituency, and a sunk cost. The political economy of the portfolio is such that cancellation carries a reputational cost for the sponsor that far outweighs the organisational benefit of releasing resources. The result is a portfolio that accumulates initiatives but rarely sheds them — a pattern I have observed across sectors, from financial services to government, from telecommunications to healthcare.

The governance framework does not cause this problem, but it consistently fails to correct it. The standard portfolio review is designed to assess the health of individual initiatives, not to challenge whether the portfolio as a whole makes sense. The question asked is "how is this programme doing?" rather than "should this programme still exist?" — and the second question, the one that would actually constitute portfolio governance, is the one that almost never gets asked.

The Decision Architecture That Is Missing

What the textbooks leave out is not the theory of decision-making but its architecture — the practical design of how decisions are structured, who is in the room when they are made, what information is available, and what happens after the decision is taken.

In the organisations where I have seen portfolio governance work effectively, three elements were present that are absent from most governance models.

First, the portfolio board had an explicit mandate to reallocate. Not merely the theoretical authority to cancel or defer programmes, but a standing expectation that reallocation was a normal part of portfolio management, not an exceptional intervention. This normalisation of stopping made the decision politically manageable in a way that it is not when cancellation is treated as a failure.

Second, the information presented to the board was structured around portfolio-level questions, not programme-level status. Instead of reviewing each programme in sequence, the board examined the portfolio through lenses that cut across programmes: strategic alignment, resource concentration, risk exposure, interdependency health. This shifted the conversation from "how is programme X doing?" to "is our investment portfolio configured to deliver our strategic intent?" — a fundamentally different question that produces fundamentally different decisions.

Third, decisions were tracked through to implementation. A portfolio board that decides to defer an initiative but has no mechanism to ensure that the resources are actually reallocated has not made a decision — it has expressed a preference. The organisations that govern well treat implementation tracking as an integral part of the governance cycle, not an afterthought.

What the Lived Experience Teaches

The textbook model of portfolio governance assumes a rational organisation — one in which decisions follow from evidence, priorities are stable enough to serve as a

framework, and leaders are willing to trade short-term political comfort for long-term portfolio health. The lived experience of governing complex portfolios suggests that none of these assumptions hold reliably.

Decisions are shaped as much by relationships, history, and organisational politics as by evidence. Priorities shift more frequently than governance frameworks can accommodate. And leaders, particularly in large organisations, are structurally incentivised to protect their own programmes rather than to optimise the collective portfolio.

This is not a counsel of despair. It is a recognition that effective portfolio governance requires designing for the organisation as it is, not as the textbooks wish it were. That means building governance structures that acknowledge the political dimension of portfolio decisions, that create safe spaces for the honest assessment of underperforming initiatives, and that equip portfolio boards with the information and the mandate to make decisions that will be uncomfortable for someone in the room.

The gap between the manual and the room is not a failure of knowledge. It is a failure of design — and closing it requires a willingness to look at what governance actually does in practice, rather than what it is supposed to do in theory.