



BEYOND THE EVIDENCE BASE

Evidence Versus Power

A Practitioner's Perspective on the Collision Between Evidence and Organisational Power

Evidence Versus Power

Portfolio · No. 66

Originally written January 2016 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“The decision to continue a failing programme is almost never made in a single moment. It is made incrementally, through a series of small accommodations, each of which is individually defensible and collectively catastrophic.”

1. The Moment Nobody Wants to Name

Every practitioner who has spent time in complex portfolio environments has witnessed the same scene. A programme is failing. The evidence is clear — the milestones are slipping, the benefits case has eroded, the technical approach has been compromised to the point where the original objectives are no longer achievable. The data, examined dispassionately, points in one direction: stop. Redirect the investment. Accept the sunk cost and reallocate the resources to initiatives that still have a viable path to value.

And yet the programme continues. It continues because a board member sponsored it. Because a restructuring was justified on the basis of its delivery. Because a supplier contract would be expensive to terminate. Because admitting failure would trigger questions about who approved it in the first place. The data says stop, but the politics say go — and in the collision between the two, politics wins with a regularity that should trouble anyone who believes in evidence-based governance.

This is not a story about bad people making bad decisions. It is a story about rational actors operating within systems that make the rational choice — stopping — personally and politically catastrophic.

2. The Anatomy of Continuation

The decision to continue a failing programme is almost never made in a single moment. It is made incrementally, through a series of small accommodations, each of which is individually defensible and collectively catastrophic.

The pattern typically begins with scope adjustment. When the original objectives become unachievable, the programme does not stop. Instead, the objectives are quietly redefined. The transformation programme becomes an "improvement programme."

The enterprise-wide rollout becomes a "phased implementation" with the most difficult phases deferred indefinitely. The business case is rewritten to reflect the reduced ambition, and the programme is presented as still on track against its revised targets. Technically, this may be true. Strategically, it is a fiction.

Next comes narrative management. The language used to describe the programme shifts from aspiration to resilience. It is no longer "delivering transformation" but "navigating complexity." Problems are reframed as "challenges" and then as "learnings." The programme's continued existence becomes its own justification: we have invested too much to stop now, the argument goes, and any future investment is measured not against the original business case but against the cost of writing off what has already been spent.

Finally, there is accountability diffusion. The people who originally sponsored the programme move on — promoted, transferred, retired. The institutional memory of why the programme was launched, what it was supposed to deliver, and who committed to what dissolves. The programme becomes an orphan, sustained not by sponsorship but by inertia. Nobody owns it strongly enough to champion it, and nobody owns it strongly enough to kill it.

The most dangerous phase in a failing programme's life is not when things go wrong. It is the period immediately after, when the organisation decides to redefine success rather than acknowledge failure.

3. Why the Portfolio Fails to Intervene

In theory, this is precisely the problem that portfolio management exists to solve. The portfolio function is supposed to take the enterprise view — to assess each programme not in isolation but as one of many competing claims on finite resources, and to reallocate investment from lower-value to higher-value initiatives. Portfolio governance boards exist to make the difficult decisions that individual programme boards cannot make for themselves.

In practice, portfolio governance is rarely equipped to do this. The reasons are structural, not personal.

Information asymmetry is the first obstacle. The portfolio board relies on information provided by the programmes themselves. Each programme has every incentive to present its status in the most favourable light possible, and the mechanisms described above — scope adjustment, narrative management, accountability diffusion — ensure that the information reaching the portfolio level has already been processed to minimise alarm. The portfolio board is making decisions about programme continuation on the basis of data that has been curated to support continuation.

The absence of a credible alternative is the second. Stopping a programme is only rational if the released resources can be deployed to better effect elsewhere. But most portfolio governance processes are better at approving new investments than at identifying them. The portfolio board may recognise that Programme A is failing, but

unless it has a clear, funded, ready-to-start Programme B waiting to absorb the resources, the argument for continuation — however weak — has no competitor. The default wins because the alternative has not been articulated.

Political entanglement is the third and most intractable. In large organisations, programmes are not independent investments. They are nodes in a network of political relationships, strategic commitments, and organisational narratives. Stopping a programme is never just a financial decision. It is a political act that redistributes power, undermines narratives, and creates winners and losers. Portfolio governance bodies are composed of the same senior leaders whose political interests are entangled with the programmes under review. Asking them to make dispassionate, evidence-based decisions about termination is asking them to act against their own institutional positions.

4. The Sunk-Cost Fallacy at Organisational Scale

The pattern is recognisable to any student of behavioural economics as the sunk-cost fallacy — the tendency to continue investing in a losing proposition because of what has already been spent rather than what is likely to be gained. At the individual level, this is a well-documented cognitive bias. At the organisational level, it is amplified by structures that make it rational.

Consider the incentives facing a portfolio board member who suspects a programme should be stopped. If they advocate for termination and are overruled, they have spent political capital for nothing. If they advocate for termination and succeed, they have created a visible failure that will be investigated, generating risk for everyone associated with the original approval. If they say nothing and the programme eventually fails on someone else's watch, they bear no personal cost. The rational strategy, given these incentives, is silence.

This is not cowardice. It is a perfectly logical response to a badly designed system. The organisation has created a structure in which the personal cost of advocating for the right decision exceeds the personal cost of allowing the wrong one to continue. Until that calculus changes, the data will continue to lose to the politics.

The sunk-cost fallacy is not a cognitive error in organisational settings. It is a rational response to incentive structures that punish the people who name failure more than the people who sustain it.

5. What Would It Take to Change This?

The honest answer is that changing this pattern requires interventions that most organisations will resist, because they challenge the distribution of power and accountability that senior leaders have spent their careers constructing.

Independent assurance with authority. Portfolio governance needs access to information that has not been filtered through programme leadership. This means investing in an independent assurance function — not an audit team that reviews after the fact, but a real-time capability that can examine programme data at source, form its own assessment, and present it directly to the governance board without intermediation. This function must report outside the line management of the programmes it reviews, and it must have the organisational authority to be heard even when its message is unwelcome.

Pre-committed decision rules. The most effective mechanism I have observed for overcoming the politics of continuation is the establishment of decision rules before the political stakes become clear. At the point of programme approval, the governance board agrees the conditions under which the programme will be paused, rescope, or terminated — and commits to a process for reviewing those conditions at defined intervals. The decision rules are not negotiable at the point of review. They were agreed in advance, when the political cost of honesty was low, and they bind the board to act on evidence rather than negotiating around it.

Normalising termination. Organisations that successfully manage portfolios treat programme termination as a routine governance outcome, not as a crisis. They build it into their language, their processes, and their incentive structures. A programme that is stopped because the evidence no longer supports it is not a failure of leadership — it is evidence of governance working as intended. The organisations that cannot stop programmes are the ones that treat termination as shameful, because shame creates the political dynamics that make honest assessment impossible.

Separating the approval decision from the continuation decision. The people who approved a programme should not be the same people who decide whether it continues. The conflict of interest is too direct: the decision-maker who approved the investment has a personal stake in its continuation that no governance process can fully neutralise. Rotating board membership, introducing independent non-executive challenge, or delegating continuation decisions to a different governance body are all imperfect solutions, but each is better than asking the same people to objectively evaluate a decision they personally championed.

6. The Deeper Question

Beneath the structural analysis lies a question about organisational character. Do we, as leaders and practitioners, genuinely want evidence to drive portfolio decisions? Or do we want evidence to play a supporting role in decisions that are fundamentally political?

The honest answer, for most organisations, is the latter. Evidence is welcomed when it supports the direction already chosen. It is tolerated when it suggests minor adjustments. It is resisted when it demands fundamental change. This is not hypocrisy — it is the reality of how power operates in complex institutions.

The practitioner's task is not to pretend this reality does not exist. It is to build systems that make the collision between data and politics visible, that create the structural

conditions for evidence to be heard even when it is unwelcome, and that normalise the act of stopping as a legitimate and valuable governance outcome.

The data will never fully win the argument against the politics. But it can be given a louder voice, a more protected platform, and a more direct route to the people who need to hear it. That is not an ideal outcome. It is, however, a realistic one — and in the world of portfolio governance, realism is a virtue that the profession has too long neglected.