


BEYOND OVERSIGHT

Governance as Enabler: The Difference Between Deciding and Watching

Programme governance has become an apparatus of assurance — but the bodies that actually work are the ones that still decide.



Giovanni Leonardi

September 2007

Governance as Enabler: The Difference Between Deciding and Watching

Transformation · No. 124

Originally written September 2007 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“Oversight quietly redefined itself as documentation.”

1. The board that met, and decided nothing

Picture a programme board on a Thursday afternoon. The pack ran to forty pages and every one of them was immaculate: a milestone plan with the critical path picked out in bold, a risk register sorted by exposure, a set of RAG indicators sitting reassuringly green save for a single amber against third-party integration. The senior responsible owner welcomed everyone, the programme manager walked the deck, three directors nodded, and a decision that had been waiting six weeks — whether to commit funds to a second delivery partner before the design was frozen — was noted, discussed, and held over for further analysis. The meeting finished on time. Everyone agreed it had gone well.

I have sat in that room more times than I can count, and for a long time I mistook it for governance working. It is not. It is governance performing. The pack was a proof of diligence; the meeting was a ritual of assurance; and the one thing a governing body exists to do — take the decision that no one below it can take — was quietly postponed for another cycle. Six weeks became ten. The partner, when finally appointed, joined too late to shape the design they were meant to build against, and the integration risk that had sat amber turned red soon after. Nobody was negligent. Every box was ticked. And the programme was materially worse off for the diligence.

We have built, across the last decade, an enormous apparatus of programme governance, and we have built most of it facing the wrong way.

2. How oversight came to mean paperwork

It is worth being honest about where the apparatus came from, because the instinct behind it was sound. The corporate failures at the turn of the decade, the arrival of Sarbanes-Oxley for anyone touching a US listing, the tightening of the Combined Code and the Turnbull guidance on internal control, the long grind of Basel II working its way through every bank's change portfolio — all of it taught boards that they would be held to account for things done several layers beneath them. The response was assurance: more reporting, more evidence, more gates, more sign-off. If you could not prove you had overseen something, you had not overseen it.

That instinct produced real goods, and a Perspective that pretended otherwise would be arguing against a straw man. Stage-gate reviews and the Gateway process have caught programmes that should never have passed go. A disciplined risk register is a better thing than a shared sense of unease. None of this should be thrown away.

But something happened on the way from principle to practice. Oversight quietly redefined itself as documentation. The question a governance body asked itself shifted

from are we making the right calls? to can we show we were watching? Those are not the same question, and an organisation can score full marks on the second while failing the first completely. The board in my opening had watched everything and decided nothing, and by the standard it was actually being held to, that counted as a success.

*Governance measured by the assurance it can evidence will optimise for evidence.
Governance measured by the decisions it unblocks will optimise for decisions. Most of ours
is measured by the first and then blamed for the second.*

3. The objection worth taking seriously

The strongest case against everything I am about to argue is not naive, and it deserves stating at full strength. It runs like this: control exists because judgement fails. The codes and the gates and the sign-offs were written in the wreckage of programmes that ran on charisma and optimism until the money was gone. Loosen the apparatus in the name of enablement and you do not get braver, faster organisations — you get the return of the very failures the apparatus was built to prevent, now dressed in a more fashionable justification. In a regulated business the argument is sharper still: the regulator does not accept "we moved quickly" as a defence, and an enabling governance that cannot produce an audit trail is simply a fine that has not happened yet.

I take this seriously because it is largely true. But the answer is not less control. It is that control and enablement are not opposites, and the belief that they are is the actual disease. A governing body that clears a decision in four days rather than deferring it for four weeks is not less rigorous; it can be far more rigorous, because it has forced itself to know what it actually needs to decide rather than hiding behind another cycle of analysis. Speed and assurance trade off only when governance has been lazy about the difference between the decisions that carry real risk and the ones that carry only paperwork.

4. What it looks like when it works

I have also, less often, sat on bodies that governed rather than watched, and the difference is not subtle once you have seen it. Three things mark them out.

The first is that they treat their own latency as a number they are accountable for. On one large regulatory programme the steering group began recording, for every decision it was asked to make, the date the question arrived and the date an answer left the room. The first month's figures were ugly: a median of thirty-one days, with the worst item sitting untouched for nine weeks. Simply publishing that number changed behaviour. Within a quarter the median was under a week — not because anyone worked harder, but because the group stopped using "let us take this offline" as a way of feeling careful. A decision deferred is a cost incurred, and the board that measures the cost soon stops incurring it so casually.

- Decision latency is treated as a governed metric, not an accident of the diary.
- Risk is absorbed upward, not merely reported upward — the body takes on what delivery cannot carry alone.
- The default is to clear the path, and the burden of proof sits with whoever wants to add a control, not remove one.

The second mark is that the body absorbs risk rather than admiring it. A reporting board receives a red risk, notes it, and asks the programme what it intends to do — which is to hand the risk straight back to the people who raised it precisely because they could not resolve it alone. An enabling board asks a different question: what can we, from where we sit, do about this that you cannot? Sometimes the answer is a quiet word with a peer director that frees a shared resource in a day. Sometimes it is the willingness to say, on the record, "proceed on this assumption and we will carry it if it proves wrong" — which is the single most valuable sentence a governing body can utter, and the one the assurance reflex most reliably suppresses.

The third mark is a disposition rather than a process: the default is yes, and the onus is reversed. In most programmes any new checkpoint, any additional sign-off, any extra column on the report arrives with the presumption that more control is prudent, and only a reckless person argues against it. Enabling governance flips the burden. A control has to earn its place by naming the decision it improves; if it cannot, it is overhead wearing the costume of diligence. This is not deregulation. It is precisely the discipline the codes already ask of the internal-control framework — proportionate, risk-based, regularly reviewed — turned honestly on governance's own machinery, which is the one place we almost never think to point it.

*"A governing body earns its existence not by the risks it records
but by the decisions no one below it could have taken."*

5. The test

There is a simple test I have come to use, and it cuts through most of the debate about frameworks and maturity models. Take any governance body in your programme and ask what would change if it stopped meeting. If the honest answer is that some reporting would go uncollated and some assurance undocumented, the body is a reporting function and should be run — cheaply — as one. If the honest answer is that real decisions would go untaken, risks would sit unowned, and the programme would slow, then it is governing, and it is worth every hour it consumes.

Most of our bodies fail that test, and the failure is not a shortage of frameworks. We have PRINCE2 and Managing Successful Programmes and the Gateway process and more maturity models than anyone can name. What we are short of is the willingness to say plainly that a governance meeting which decided nothing was not a success merely because it was well documented. The paperwork is real work, and the codes that demand it were written for good reason — but the paperwork is the evidence of governance, not the act of it. The act is the decision.

The organisations that have this right are not the ones with the thickest governance manuals or the greenest reports. They are the ones where a programme director, asked what the steering group is for, does not describe a meeting. They describe the last three decisions it took that they could not have taken alone, and how quickly it took them. That is governance as enabler, and from the inside it looks less like control and more like the quiet, unglamorous business of clearing the way.