


BEYOND INTEGRATED GOVERNANCE

The Dual Mandate That Governance Frameworks Refuse to Acknowledge

On the competing claims of compliance and delivery — and the trade-off programme leaders are left to manage alone



Giovanni Leonardi

October 2006

The Dual Mandate That Governance Frameworks Refuse to Acknowledge

Programme · No. 72

Originally written October 2006 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“They are competing claims on the same finite resource: the organisation's capacity to govern.”

The Board That Governs Twice

The programme board convenes fortnightly, and every meeting follows the same arc. The first hour is given to regulatory compliance: evidence packs, audit findings, the regulator's latest letter, and the gap analysis that maps the programme's deliverables against the regulatory timetable. The remaining thirty minutes — often less, after the compliance discussion has overrun — covers delivery: milestone progress, resource conflicts, dependencies, technical risk. By the time the board reaches the decision it was convened to make — whether to proceed to the next stage — half the members have already begun checking their watches.

This is not a dysfunctional board. It is a board doing exactly what two governance frameworks, operating simultaneously, require it to do. The dysfunction lies not in the people but in the architecture — in the quiet assumption, embedded in every major governance framework, that regulatory oversight and delivery oversight are complementary perspectives on the same programme. They are not. They are competing claims on the same finite resource: the organisation's capacity to govern.

The Textbook Promise

The frameworks are careful and comprehensive. COBIT sets out the structures by which IT governance aligns technology investments with business objectives and regulatory requirements. The Combined Code prescribes the board-level oversight that ensures risk is managed and controls are effective. MSP provides the programme-level architecture — the vision, the blueprint, the tranches, the benefits — through which strategic change is delivered. Each is well-designed for its purpose.

The implicit promise is convergence. Layer these frameworks together — strategic governance at board level, programme governance at delivery level, regulatory governance threaded through both — and the result is a coherent system in which compliance and delivery reinforce one another. The regulatory requirements become deliverables in the programme plan; the programme's governance provides the evidence trail the regulator requires; the board oversees both through a single, integrated assurance structure.

In a textbook, this works. In practice, it breaks — not because any individual framework is wrong, but because the combination makes demands no real programme board can meet.

Why They Compete

The mechanism is straightforward, and every programme leader in a regulated sector has lived it.

Regulatory governance carries external enforcement. The regulator sets the timetable, defines the standard, and holds the power to fine, restrict, or revoke. The consequences of non-compliance are existential — not in the slow-burn way that failed programmes damage reputations, but in the immediate, binary sense that a regulatory breach can shut down a line of business. Delivery governance, by contrast, carries internal enforcement. The sponsor sets the targets, the programme board monitors progress, and the consequences of failure — cost overruns, delayed benefits, strategic drift — are negotiable, deferrable, and shared across too many accountable parties to land on any one of them.

When both sit on the same board agenda, the competition is decided before it begins. Regulatory items command the time, the attention, and the anxiety. Delivery items are compressed into whatever space remains. Over time, this bias becomes structural: the programme office allocates its strongest people to regulatory workstreams, the board papers grow thicker on compliance evidence and thinner on delivery insight, and the language of risk narrows until "programme risk" means "regulatory risk" and nothing else.

The result is a programme that is meticulously governed in one dimension and barely governed in another. The compliance evidence is pristine. The delivery is drifting. And because the governance architecture treats both as a single integrated system, there is no mechanism to make the drift visible until it is too late to correct.

The Cost the Frameworks Ignore

The textbooks treat governance as though it were costless — a design decision rather than a resource allocation. Add a regulatory assurance layer to the programme governance, and the total governance is the sum of the two. But governance consumes the very capacity it is meant to protect. Every hour a programme director spends preparing compliance evidence is an hour not spent managing dependencies. Every board meeting dominated by regulatory status is a meeting in which a critical delivery decision was not taken. Every member of the programme office assigned to audit response is a member not available for benefits tracking or stakeholder management.

This is not a marginal cost. In the large regulatory change programmes that dominate financial services — Basel II implementations, Sarbanes-Oxley compliance, FSA-mandated systems upgrades — the governance overhead routinely absorbs twenty to thirty per cent of the programme's management capacity. The frameworks acknowledge none of this. They assume that the organisation can add governance layers without diminishing the capacity available for the work being governed.

They also assume that programme boards can hold two fundamentally different lenses at the same time. A regulatory lens asks: are we compliant? Is the evidence sufficient? Will the regulator accept this? A delivery lens asks: are we making progress? Are the dependencies managed? Will we realise the benefits? These are not two views of the same question. They require different information, different judgement, and different expertise. A board that attempts both in a single meeting does neither well — and in practice defaults to the one with external consequences.

What Practitioners Actually Do

The programme leaders who manage this well – and some do manage it remarkably well – treat the dual mandate not as a design problem but as a genuine trade-off. They do not pretend that the right governance architecture will make regulatory and delivery oversight align. Instead, they make deliberate choices about where to spend the organisation's governance capacity.

The pattern I have observed in the programmes that navigate this tension most effectively is separation at the working level and integration only at the decision point. The regulatory governance stream operates with its own reporting, its own assurance, and its own escalation path. The delivery governance stream does the same. The two streams come together at stage gates – the points at which the programme board must decide whether to proceed – and only there.

In one arrangement that recurs across banking and insurance, the regulatory governance stream reports monthly to a compliance committee chaired by the chief risk officer, while the delivery governance stream reports to the programme board chaired by the business sponsor. The two streams share a dependency log and a change-control process, but their assurance is separate: the compliance committee never reviews delivery milestones, and the programme board never reviews regulatory evidence packs. They converge at the stage gate, where both groups sit together and the programme director presents a single recommendation – proceed, defer, or restructure.

This is not what the textbooks recommend. The textbooks recommend integration throughout: a single risk register, a single reporting framework, a single assurance plan. But integration throughout means competition throughout, and the competition is always won by the stream with external enforcement. Separation preserves the integrity of both streams. Delivery governance gets the space to function as delivery governance – to ask the hard questions about progress, dependencies, and benefits – without being crowded out by compliance.

The trade-off is real: separation means the board must work harder at the gates to synthesise two streams into a single decision. It means accepting that governance is not a unified system but two parallel systems that must be actively reconciled. It means giving up the textbook's comforting promise of coherence.

The dual mandate is not a coordination problem. It is a resource-allocation problem that governance frameworks – by treating oversight as costless – have no vocabulary to describe.

The Simplification That Matters

The governance frameworks will continue to evolve. COBIT will release new editions; the Combined Code will be revised; programme management standards will grow more sophisticated. None of this will resolve the tension at the heart of the dual mandate, because the tension is not a gap in the frameworks. It is a feature of the operating environment that the frameworks are built to deny.

Regulatory governance and delivery governance make competing claims on the same finite resources: board time, management attention, programme office capacity, and the

patience of the teams doing the work. Until the frameworks acknowledge this competition — until they stop treating governance as an infinitely scalable overlay and start treating it as a scarce resource that must be allocated — practitioners will continue to solve the problem the textbooks refuse to name. They will do it pragmatically, imperfectly, and largely without credit. But they will do it, because the alternative — following the textbook and watching delivery governance disappear beneath the weight of compliance — is a failure no programme leader can afford.