


BEYOND GOVERNANCE THEATRE

The AI Governance Committee Is the New Steering Committee — Same Theatre, Higher Stakes

Why the Enterprise's Newest Oversight Body Has Already Inherited Its Oldest Dysfunction

Giovanni Leonardi

November 2024



The AI Governance Committee Is the New Steering Committee — Same Theatre, Higher Stakes

Leadership · No. 109

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

“When twelve people approve a use case, no one has approved it.”

1. The Same Committee, Reconvened

The quarterly AI governance committee convenes for the third time. Twelve members — two from legal, three from IT, the CISO, the CDO, a representative from HR, someone from procurement, and three senior leaders whose diaries required six weeks to align. The agenda runs to fourteen items. The pack, circulated the evening before, contains forty-seven slides. By item five, the committee will approve two use cases it does not understand, defer one it should have killed, and produce minutes that record none of these failures. The members leave feeling responsible. No one is accountable.

This scene is not speculative. It is the pattern I have observed forming across every sector in the past eighteen months, as organisations rush to stand up AI governance structures at the same pace they once rushed to stand up steering committees for their transformation programmes.

Sixteen years ago, I wrote about the steering committee — the governance body that had become, in most organisations, a mechanism for distributing accountability until it weighed nothing. The pathology was specific: too many members for honest challenge, papers written for information rather than decision, and a collective rhythm that approved what was already in motion. The argument was simple. The steering committee was not governing; it was performing governance.

I had hoped the pattern would not recur. It has recurred exactly. But the poem is darker this time, and the rest of this piece is about why.

2. The Inheritance

The AI governance committee has inherited every pathology of the steering committee because it was built by the same people, using the same organisational instincts. When the board asked for "governance around AI," the corporate machine did what it always does: it formed a committee. And it formed it in the image of every committee before it —

broad membership to ensure representation, regular meetings to ensure rhythm, and a terms of reference document that confuses oversight with attendance.

Consider what the typical AI governance committee actually does in its first year. It defines an AI policy — usually by adapting a template that circulates freely among the consultancies. It establishes a use-case intake process — a form, a scoring rubric, a register. It reviews those use cases — which in practice means reading a two-page summary prepared by someone who wants approval and asking questions that sound rigorous but cannot be answered in the room. And it reports upward — a RAG-rated dashboard to the board that translates uncertainty into the soothing language of traffic lights.

None of this is governance. It is administration dressed in governance's language. And to anyone who remembers the trajectory of the steering committee, the shape of what comes next is already visible.

3. The Category Error

The steering committee's failures were expensive but bounded. A transformation programme that was inadequately governed might overrun by millions, might fail to deliver its benefits, might consume years of organisational energy. These are serious costs. But the programme itself was inert — it did not act autonomously, did not compound its own errors, did not scale its failures faster than humans could detect them.

The systems now being waved through by AI governance committees do all of these things.

The steering committee governed projects — bounded, sequential, human-paced. The AI governance committee is being asked to govern systems — unbounded, continuous, machine-paced. And it is being asked to do so with the same structure, the same cadence, and the same depth of engagement that failed for projects.

A large language model deployed into customer-facing operations does not wait for the next committee meeting to make its next decision. It is making thousands of decisions every hour, each one carrying the organisation's brand, its legal obligations, and its customers' trust. A model integrated into underwriting or credit decisioning is not a project that can be paused at a stage gate — it is a live system whose outputs accumulate in the real world, and whose consequences may not surface for months. A recommendation engine does not submit a change request before altering its behaviour; it drifts as its data drifts, and the committee that approved it in January may not learn what it has become until the regulator asks in October.

Governance theatre around a transformation programme wastes money. Governance theatre around autonomous systems compounds errors at machine speed. This is not a difference of degree. It is a difference of kind.

4. Where the Failures Compound

The specific failures are not mysterious. They are the same three that afflicted the steering committee, but each now carries consequences that scale differently.

4.1 Membership Without Literacy

The committee is assembled for representation rather than capability. The logic is familiar and, on its surface, reasonable: every function that touches AI should have a voice at the table. Representation builds buy-in, ensures no perspective is overlooked, guarantees that legal, ethical, operational, and commercial considerations are all heard. This is genuinely the strongest argument for the current model.

But representation and capability are not the same thing, and the trade-off between them is not symmetric. A committee of twelve achieves breadth at the cost of depth, and in AI governance the cost of insufficient depth is categorically different from the cost of insufficient breadth. The function whose voice was missing from the room may surface its concern at the next meeting. The technical risk that no one in the room could identify surfaces in production.

What is required is not a member from every function but a small number of people — five, perhaps six — who can read a model card, interrogate a fairness metric, understand what a confidence threshold means in operation, and recognise when they are being told a system is "low risk" by someone who has defined risk to exclude everything that might slow the deployment.

4.2 Cadence Without Continuity

The committee meets quarterly, or monthly at best. The systems it governs operate continuously. A quarterly meeting is not governance of a live system; it is a post-mortem dressed as oversight. The committee reviews what has already happened, approves what is already built, and the gap between the speed of the technology and the rhythm of the oversight body widens with every cycle.

The redesign here is not simply to meet more often — meeting fatigue solves nothing — but to abandon the assumption that governance happens primarily in meetings at all. Effective oversight of live systems requires continuous monitoring against defined parameters, threshold-based escalation protocols, and decision rights that can be exercised between meetings by named individuals with the authority to pause, modify, or retire a system when predefined conditions are met. The meeting becomes the review of how those authorities were exercised — not the moment of decision itself.

4.3 Collective Approval Without Individual Accountability

In the steering committee, accountability was distributed across the membership until it dissolved entirely. The same dynamic is already visible. When twelve people approve a use case, no one has approved it. When the committee "endorses" a risk assessment, no individual has staked their professional judgement on its adequacy. And when the system subsequently fails — when it produces biased outputs, when it hallucinates in a context where hallucination carries legal consequences, when it drifts beyond its validated

parameters — the committee's minutes will show a collective decision that no individual can be held to account for.

The fix is direct and uncomfortable: named individual accountability for every AI system in production. Not the committee. Not the function. A person — with their name against the system, with the authority to make decisions about it, and with the exposure that comes from both. This is the practice that financial services learned, painfully, about model risk over the past decade. It is the practice the AI governance committee must adopt before it learns the same lesson at greater cost.

5. The Window Is Narrowing

The regulatory environment is making this redesign unavoidable, even for organisations content with theatre. The EU AI Act, now moving from text to implementation, places specific obligations on deployers of high-risk AI systems — obligations that require documented human oversight, not documented committee attendance. The sector-specific guidance emerging from financial regulators, healthcare authorities, and data protection bodies assumes a level of technical governance capability that the typical committee structure cannot deliver.

Organisations that wait for the regulator to tell them their governance is inadequate will find themselves restructuring under external pressure and public scrutiny. Those that recognise the pattern now — that see the AI governance committee already becoming what the steering committee became — have perhaps twelve months to redesign before the gap between their governance structures and their regulatory obligations becomes indefensible.

6. The Redesign That Is Already Known

The pattern rhymes, but the stakes do not. The steering committee's failure was a failure of organisational discipline — expensive, frustrating, ultimately survivable. The AI governance committee's failure, if left unreformed, is a failure of oversight over systems whose errors compound at a speed and scale that makes retrospective governance meaningless.

The redesign is neither radical nor unknown. Fewer members with genuine technical literacy, chosen for capability rather than representation. Continuous monitoring with threshold-based escalation, not periodic reviews. Named individual accountability for every system in production. Decision authorities that operate between meetings, not only within them. Failure scenarios rehearsed before they are needed, not reconstructed from minutes after the fact.

Every element of this redesign was already understood when the steering committee was failing. The organisations that fixed their programme governance did so by applying these principles. The question is not whether we know what effective AI governance looks like. The question is whether we will build it before the systems we are supposed to govern have outrun the theatre we are performing in their name.