MyConsultancy

Why AI Cannot Be Governed Like IT or Digital Programs

The AI Operating Model Playbook

Manoj Tavarajoo

January 2026

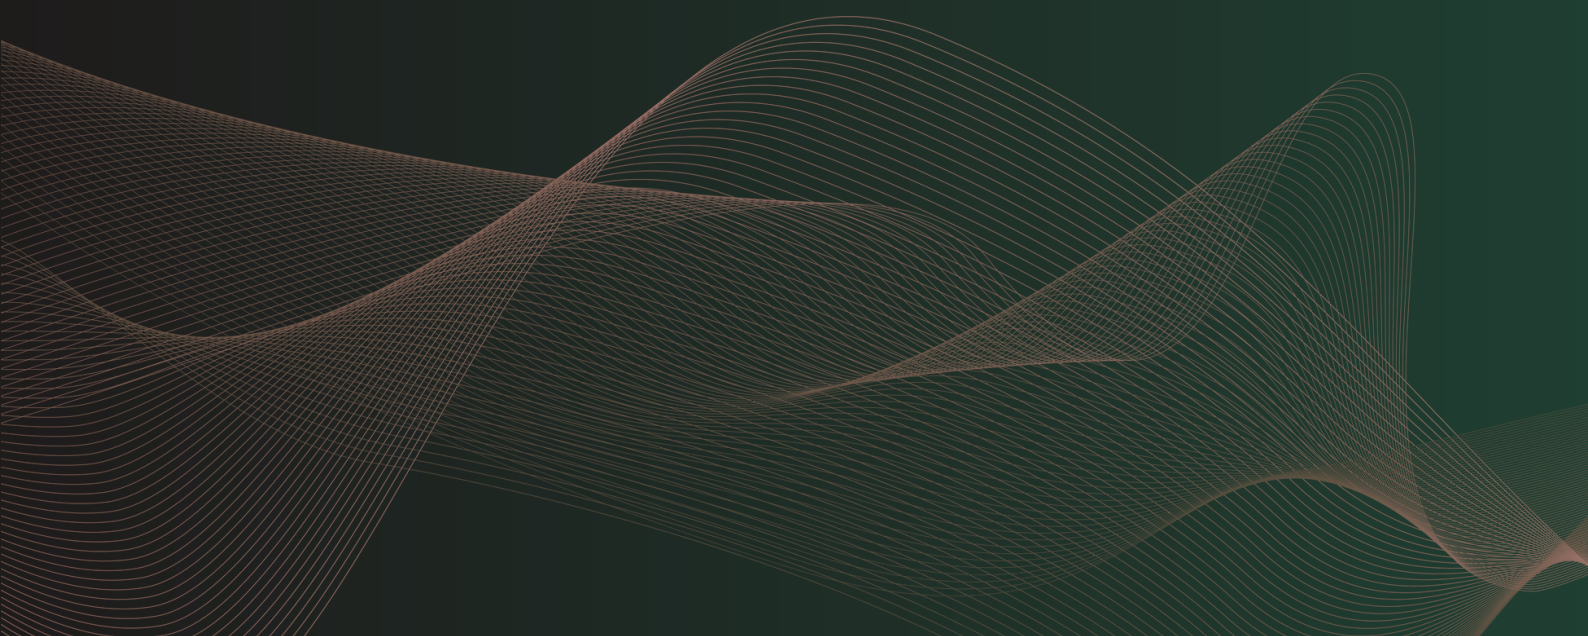

MyConsultancy

Why AI Cannot Be Governed Like IT or Digital Programs

The AI Operating Model Playbook

Manoj Tavarajoo

Opening context

When AI initiatives struggle to scale, governance is frequently identified as the bottleneck. Delivery teams cite slow approvals, unclear decision rights, and risk functions that intervene too late or too bluntly. Executives respond by adding policies, committees, and review stages in an effort to regain control.

Despite this, confidence at board level often remains fragile. Leaders sense that risks are not being managed effectively, even as governance artefacts proliferate. At the same time, AI teams become constrained by controls that appear disconnected from how AI systems behave in production.

This tension reflects a deeper misunderstanding of what effective governance looks like when systems learn and adapt over time.

Why this fails in most organisations

Most enterprises govern AI using mechanisms inherited from IT and digital program delivery. These mechanisms are built around upfront approval, static risk assessment, and centralised decision-making. They assume systems behave predictably once deployed and that risks can be identified and mitigated in advance.

AI does not conform to these assumptions. Models evolve as data changes. Performance shifts as user behaviour adapts. New risks emerge in production rather than at design time. Governance approaches that concentrate authority at the beginning of the lifecycle leave organisations exposed later, when real-world behaviour diverges from expectations.

Another common failure is the separation of governance from execution. Risk, legal, and ethics functions often operate at arm's length from day-to-day AI work. Their involvement is episodic, triggered by escalations or audits rather than embedded in operational rhythms. This weakens shared understanding of trade-offs and creates adversarial dynamics.

Finally, traditional governance structures often conflate control with restriction. In attempting to eliminate uncertainty, they suppress experimentation. This delays learning, pushes risk discovery downstream, and increases the cost of failure when issues surface at scale.

The operating model insight

AI requires governance that evolves alongside the systems it oversees.

Effective AI governance is not about preventing change. It is about ensuring learning, accountability, and oversight progress together. This requires a shift from approval-based control to ongoing stewardship.

Decision rights must be clear, but they must also sit close enough to the work for judgement to be exercised in context. Oversight must extend into production, where model behaviour, performance drift, and unintended consequences actually emerge. Governance must focus less on intent at approval and more on behaviour over time.

This reframes governance from gatekeeping to stewardship. It does not reduce the importance of risk management, but it fundamentally changes how and where it is exercised.

What this looks like in practice

Organisations that govern AI effectively behave differently at key decision moments. Instead of relying solely on upfront approvals, they establish guardrails that define acceptable risk thresholds, escalation triggers, and decision boundaries.

Risk, legal, and ethics functions are embedded into delivery forums and operating rhythms. Their role is to surface trade-offs early, not to block progress after decisions are made. This proximity improves judgement and builds trust.

At board and executive level, attention shifts from individual use case approvals to portfolio-level assurance. Leaders focus on how AI systems perform in production, how quickly issues are detected and corrected, and whether learning is being institutionalised across the organisation.

Common mistakes to avoid

Adding governance layers without changing decision logic. More committees rarely produce better outcomes.

Treating ethical principles as abstract statements rather than operational requirements. Without translation into day-to-day decisions, ethics remains disconnected from execution.

Centralising all AI decisions in the name of risk management. This slows response times and weakens accountability where AI actually operates.

Relying on documentation-heavy assurance processes that provide comfort without insight. Compliance on paper does not equate to control in practice.

What leaders must do differently

Leaders must redefine what good governance means in an AI context. This starts with accepting that not all risks can be anticipated upfront and that oversight must extend into production.

Decision rights must be explicitly designed, not assumed. Teams need clarity on what they can decide autonomously and when escalation is required. Risk and ethics functions must be empowered to influence execution, not just review outcomes.

At board level, the focus must shift from approval to assurance. The central question becomes whether the organisation can observe, respond to, and learn from AI behaviour at speed.

Conclusion

Governing AI as if it were IT or a digital program creates a false sense of control while undermining learning. The result is slower execution, hidden risk, and growing frustration.

AI demands governance that is continuous, contextual, and embedded in how the organisation operates. Without this shift, attempts to scale AI will remain constrained, regardless of intent or investment.

MyConsultancy

Building AI-Ready Organisations at Scale

www.myconsultancy.com.au