

EXECUTIVE BRIEF

MyConsultancy

Governing Agent AI

| When AI systems move
from advice to action

Manoj Tavarajoo



THE GOVERNANCE SHIFT

**AI governance began as governance of outputs.
It has become governance of delegated operational authority.**



TRADITIONAL AI SYSTEMS

- ✓ Generated recommendations and predictions
- ✓ Produced outputs for human review
- ✓ Supported decisions already being made
- ✓ Required explicit human approval to act
- ✓ Operated within bounded, visible workflows

AGENTIC AI SYSTEMS

- + Initiate workflows and trigger actions
- + Interact across systems and tools autonomously
- + Make operational decisions without human approval
- + Coordinate multi-step processes dynamically
- + Delegate authority across agent networks

The question is no longer whether AI can generate outputs. It is what authority organisations are prepared to delegate.

**AI governance began
as governance of models.**

**It has become governance of
delegated operational
authority.**

”



Most governance models were designed for AI that advises, not AI that acts

TRADITIONAL GOVERNANCE ASSUMES

- 
- ✓ Human review before consequential action
 - ✓ Visible, bounded decision-making processes
 - ✓ Clear intervention points in every workflow
 - ✓ Stable system behaviour post-deployment

AGENTIC SYSTEMS CHALLENGE ALL FOUR

- ✗ Actions are initiated without prior human approval
- ✗ Multi-step decisions occur across opaque systems
- ✗ Intervention points are dynamic and time-critical
- ✗ System behaviour evolves through context and learning

Autonomy changes the governance problem more than intelligence does.

Not all AI systems create the same governance exposure

THE FOUR-TYPE AI GOVERNANCE TAXONOMY

AI systems differ by what they do, the role humans play, and the governance challenge they create.



1. TRADITIONAL AI

Predictive Systems

What it does Predicts, classifies, scores or detects patterns in data.	Human role Human makes the decision.
Governance focus Model accuracy, fairness and explainability.	Primary Governance Risk Incorrect or biased predictions lead to poor decisions.
Example Use Cases Fraud detection, credit scoring, demand forecasting.	



2. GENERATIVE COPILOT

Content & Recommendations

What it does Generates content, analysis or recommendations for human use.	Human role Human reviews and decides.
Governance focus Output reliability, appropriateness and safety.	Primary Governance Risk Misinformation or harmful / inappropriate output influences decisions.
Example Use Cases Drafting emails, summarising documents, generating analysis or insights.	



3. WORKFLOW AUTOMATION

Deterministic Execution

What it does Executes predefined, deterministic rules and processes.	Human role Human configures and monitors.
Governance focus Process integrity, rule configuration and operational control.	Primary Governance Risk Misconfiguration or rule failure disrupts operations.
Example Use Cases Invoice processing, approval routing, IT service workflows.	



4. AGENTIC AI


Autonomous Action

What it does Plans, initiates and acts across tools, workflows and systems on behalf of the organisation.	Human role Human delegates authority and intervenes by exception.
Governance focus Delegated authority, autonomy boundaries and accountability for action.	Primary Governance Risk Autonomous action outside approved boundary causes harm or loss.
Example Use Cases Customer issue resolution, financial operations, supply chain orchestration.	

Lower GOVERNANCE INTENSITY & COMPLEXITY Higher

KEY TAKEAWAY: As AI systems move from prediction to autonomous action, the governance challenge shifts from output quality to **delegated authority and accountability for action**.

Every autonomous system requires an operational ceiling



Every autonomous system requires clearly defined limits on delegated authority. Unbounded autonomy is not innovation. It is unmanaged delegation.

Human oversight is changing fundamentally

Human oversight used to mean...

- ✓ Reviewing every AI output before it was acted on
- ✓ Approving every workflow at defined checkpoints
- ✓ Validating every decision with a human in the loop
- ✓ Operating predictable, sequential processes
- ✓ Monitoring stable systems with visible boundaries

Governance increasingly depends on...

- 🛡 Defining the operational boundaries before delegation
- 🛡 Monitoring behavioural drift across system outputs
- 🛡 Intervening at escalation thresholds, not every action
- 🛡 Maintaining containment capability at all times
- 🛡 Governing authority structures, not individual decisions

Human oversight is not the presence of a person somewhere in the workflow. It is a designed governance control.

Intervention capability matters more than approval capability



Delegated authority does not remove accountability

Agentic systems may initiate actions, trigger decisions and coordinate workflows without human approval at every step. Accountability for those actions remains organisational.

THE AGENT MAY

- + Execute actions on behalf of the organisation
- + Interact across external systems and data
- + Trigger downstream decisions and workflows
- + Coordinate with other agents without oversight

BUT ACCOUNTABILITY REMAINS FOR

- + The authority boundaries granted to the agent
- + The design and adequacy of escalation mechanisms
- + The monitoring and intervention capability in place
- + The governance structures that govern delegation

BOARDS REMAIN ACCOUNTABLE FOR

- + Setting the autonomy ceiling approved for each system
- + Ensuring intervention frameworks exist and are tested
- + Receiving evidence that containment is operational
- + Governing the delegation of operational authority

Delegated authority changes execution responsibility. It does not remove accountability.

The governance decisions cannot wait for full autonomy

DEFINE AUTONOMY CEILINGS

Set the maximum level of authority each AI system is permitted to exercise without mandatory human intervention. Make it explicit, documented and board-approved.

ESTABLISH ESCALATION THRESHOLDS

Define precisely when, how and to whom agentic behaviour must escalate. Thresholds must exist before systems are deployed, not after incidents occur.

DESIGN INTERVENTION RIGHTS

Build the organisational capability to pause, redirect, override and contain agentic AI systems. Intervention must be tested, not assumed.

IMPLEMENT BEHAVIOURAL MONITORING

Monitor system behaviour continuously against defined boundaries. Behavioural drift is a governance event, not a technology anomaly.

GOVERN DELEGATED AUTHORITY EXPLICITLY

Treat every agentic AI deployment as a formal delegation of operational authority. The governance structures must match the authority granted.

Agentic AI is scaling faster than most governance models are evolving.

WHAT LEADERS MUST DO NOW

Five questions leaders should be uncomfortable answering




Where are AI systems already taking action without meaningful human review?

Have we defined explicit autonomy ceilings for agentic systems?

Do escalation thresholds exist for autonomous AI behaviour?

Can we intervene immediately when an AI agent behaves unexpectedly?

Are we governing AI outputs or governing delegated operational authority?



If any of these are difficult to answer, agentic AI is already ahead of your governance model.

THE DEFINING GOVERNANCE QUESTION

The defining governance question is no longer whether AI can produce outputs.


IT IS WHAT AUTHORITY ORGANISATIONS ARE PREPARED TO DELEGATE.

The organisations that govern agentic AI successfully will not be those with the most advanced agents. They will be those with the clearest operational boundaries, the most rigorous escalation design and the institutional will to intervene when agents exceed their authority.

*AI governance began as governance of models.
It has become governance of delegated operational authority.*

THE AI TRANSFORMATION FRAMEWORK

Lead AI. Organise AI. Govern AI. Execute with discipline.



ABOUT

MANOJ TAVARAJOO

Manoj Tavarajoo has spent over two decades working with boards, executives and senior leaders across enterprise, digital and AI transformation. His work sits at the intersection of strategy, operating model design, governance and execution, where transformation ambition becomes measurable enterprise capability.

He is the author of *Leading the AI Transformation*, *The AI Operating Model Playbook* and *The AI Governance Imperative*, together forming a practical leadership trilogy on how organisations lead AI, organise for AI and govern AI at enterprise scale.

Manoj leads MyConsultancy, an independent advisory and transformation practice based in Australia, partnering with boards and executive teams on enterprise, digital and AI transformation spanning strategy, operating model design, governance and execution. He also undertakes select senior embedded leadership roles where sustained executive ownership is required to translate strategy into measurable enterprise outcomes.

 [@manojtavarajoo](#)
 www.myconsultancy.com.au




THE AI TRANSFORMATION TRILOGY

Three books, one system: how to lead AI, organise for it, and govern it once it acts.

**LEAD AI**
AI Strategy and Transformation Leadership


**ORGANISE AI**
AI Operating Model and Organisational Design

**GOVERN AI**
AI Governance, Accountability and Assurance



Leading the AI Transformation
Lead AI
How Bold Leaders Unlock Value, Reshape Strategy, and Build the Future

A practitioner-focused guide for executives navigating the shift to AI-enabled organisations. Covers leadership alignment, AI strategy, governance, responsible adoption, and building sustained enterprise AI capability.



The AI Operating Model Playbook
Organise AI
Why Structure, Not Algorithms, Determines AI Outcomes

A structural guide to scaling AI in the enterprise. Examines the operating model decisions that determine whether AI initiatives scale or stall, including portfolio governance, funding logic, decision rights, team design, and process integration.



The AI Governance Imperative
Govern AI
How Boards and Leaders Govern Value, Risk, Accountability and Autonomous AI

A governance guide for boards and leaders on AI behaviour, evidence, accountability, vendor risk and autonomous systems. At its centre is the Behavioural Envelope, converting approval into a continuous evidence obligation.



Governing Agentic AI