NEURAL SEVEN
ADVISORY

AI, To use or Not to use: The questions you need to be

Artificial intelligence has arrived in the boardroom wearing two faces. One is the polished promise of productivity, growth and strategic edge. The other is a less comfortable reality: variable outputs, hidden costs, weak governance and risks that do not fit neatly inside yesterday's control frameworks.

That tension is now one of the defining leadership issues of the decade. For senior executives and directors, the challenge is not whether AI matters. It does. The challenge is knowing when it makes commercial sense, when it creates more risk than value, and how to approach governing AI systems that evolve before the organisation learns the hard way.

There is no shortage of AI commentary. There is a shortage of practical judgment. Plenty has been written about transformation, disruption and competitive advantage. Much less has been said about the basic questions that should come first: what problem is AI solving, what does it actually cost, how will inaccurate outputs be detected, and who is accountable when the system changes over time?

This is where AI street smarts matter. Not hype, not fear, and certainly not a procurement strategy driven by vendor demos and hype. Just sharper questions, stronger governance and a realistic view of where AI helps and where it plainly does not.

The questions leaders should be asking

The first question should not be “How do we use AI everywhere?” It should be “Where would AI create durable value, and can we govern it properly once it is live?” Australian director guidance has been explicit that boards need guardrails that allow AI’s benefits while managing risk, and that requires clarity in governance structures, controls, monitoring and accountability.

That matters because AI systems are not as static as traditional enterprise software often is. Their behaviour can shift with retraining, prompt changes, new data, vendor updates or integrations with external tools. Governing AI systems that evolve means thinking beyond launch approvals and pilot milestones. The

model that performed acceptably in March may behave differently in September unless monitoring and review have been designed in from the start.

The second question is about truth. Leaders should ask where an AI error would actually matter. If a tool produces a weak first draft for an internal workshop, the risk may be modest. If the same system is summarising compliance obligations, shaping customer responses or informing board decisions, the risk changes immediately. Public guidance on misinformation and hallucinations now treats these issues as operational and governance concerns, not just technical quirks.

The third question is cost. Not the subscription fee on a vendor slide, but the actual cost of responsible, enterprise-grade use. AI does not become economical simply because the first demo looked cheap. In many cases, the largest expenses appear later through integration, monitoring, human assurance, vendor dependency and usage patterns that no one modelled properly in the first place.

The hidden costs of AI

One of the more persistent myths in the market is that AI is inexpensive because access can start with a modest monthly fee or a pay-as-you-go API. That is only partly true. Public pricing pages show that AI services often charge separately for input tokens,

output tokens, and cached inputs, while other features can incur additional costs for retrieval, search, and storage.

On paper, token pricing can look harmless. In practice, it compounds quickly. A single use case might involve a user prompt, a large context window, file retrieval, multiple revisions, external search and a lengthy response. Repeat that across teams, workflows and business units, and the spend curve looks very different from the original pilot.

There is a deeper issue. Token costs are only visible on the meter. The larger line items often sit elsewhere:

· Data preparation and remediation, because internal data is often fragmented, duplicated or poorly tagged.

· Integration into identity systems, repositories, workflow platforms and security controls, which is where many projects become slower and more expensive than expected.

· Monitoring, testing and evaluation, because models need to be checked for drift, bias, hallucinations and changing performance over time.

· Human review, because in high-stakes settings AI outputs still need subject matter experts to validate them.

· Governance overhead, including policy work, reporting, risk assessment, auditability and third-party oversight.

For executives and boards, the useful question is not “What does the model cost?” It is “What does safe, governed, scalable use cost over the next 12 to 24 months?” That is the more honest number. It includes the token bill, but also the infrastructure, controls, assurance work and people required to stop a fast-moving system from becoming a quiet source of strategic risk.

A simple way to explain it is this: token pricing is the taxi meter, not the full trip budget. The full budget includes the vehicle, the insurance, the route planning, the driver and the cost of ending up in the wrong suburb.

Productivity: where the real value sits

Most AI business cases talk confidently about time saved. Far fewer ask what happens next. That gap matters because time saved is not the same as productivity created. If AI helps someone complete a task in half the time but the recovered hours are not redirected into more valuable work, the organisation may have changed the shape of the day without changing the economics of the business.

Recent research supports that distinction. The Federal Reserve Bank of St Louis found that workers using generative AI reported saving 5.4 per cent of their work hours, equivalent to a 1.1 per cent increase

in aggregate productivity, and estimated that workers were about 33 per cent more productive in the hours when they used the tools. McKinsey's work makes a similar point at a broader level, arguing that generative AI's productivity effect depends heavily on how effectively worker time is redeployed into other activities.

That is where the false economy appears. If saved time is absorbed by low-value admin, duplicated checking or simply disappears into unmeasured slack, the headline efficiency gain is misleading. Workday's research found that nearly 40 per cent of AI time savings were being lost to rework, including correcting errors, rewriting content and verifying outputs that were not reliable enough the first time.

For leaders, this changes the conversation. AI should not be judged only by whether it helps people work faster. It should be judged by whether the saved time is intentionally reinvested into higher-value work: better analysis, stronger customer engagement, innovation, risk reduction, improved decision-making or capability building.

That means executives need to ask two linked questions every time an AI productivity claim is made:

· What measurable time is genuinely being saved?

· What are people doing with that time, and how much of it is being lost to checking or rework?

If there is no good answer to the second question, the organisation may be looking at a false economy. Faster task completion on its own is not transformation. It is only valuable if the business deliberately captures and redirects that capacity.

AI for executives

For the C-suite, the issue is not whether AI should appear in the strategy. It already has. The sharper question is how to turn AI into an operating capability without producing a mess of disconnected tools, duplicated spend and unmanaged risk.

1. Set an explicit AI posture

Executives should define the organisation's AI posture in plain terms. Is the business aiming to be a cautious adopter, a selective fast follower or a category leader in a narrow set of use cases? Many organisations behave as though they can be all three at once.

A clear posture improves investment discipline, gives managers a practical boundary for experimentation and reduces the habit of approving isolated pilots without a shared logic. It also helps separate strategic use cases from "AI for AI's sake".

2. Put data and controls before scale

Many AI efforts stumble because leaders approve front-end experiences before addressing back-end realities. If the data is weak, permissions are unclear, and lineage cannot be trusted, AI simply amplifies confusion more efficiently.

Strong data governance improves output reliability, reduces remediation later and makes it easier to defend decisions when regulators, customers or auditors ask how a system works. In practical terms, this is often the difference between a robust platform and a fragile demo.

3. Manage AI as a portfolio, not a collection of experiments

Executives should avoid letting each business unit procure and test AI in isolation. A portfolio view with common oversight, defined value measures, and clear exit criteria is more effective than scattered experimentation.

A portfolio model makes it easier to compare use cases, stop weak initiatives early and focus resources where AI has a credible path to measurable value. It also reduces overlap, shadow AI and the multiplication of tools that all claim to solve roughly the same problem.

4. Track value capture, not just activity

Executives should measure whether AI is creating usable capacity and whether that capacity is being reinvested into work that matters. That means looking beyond volumes, turnaround times and number of prompts processed.

This is how leadership avoids the false economy of “faster but no better”. If time savings are not being converted into improved decisions, reduced risk, stronger service or more strategic output, the AI program may be busy without being valuable.

AI for board members

Boards have a different role. They are not there to design models or choose prompt templates. They are there to ensure the organisation’s use of AI aligns with its strategy, risk appetite, obligations, and values. That means oversight without illusion.

1. Clarify where AI oversight sits

Boards should make AI oversight an explicit governance responsibility rather than letting it drift between audit, risk, technology and strategy committees. Public Australian guidance points to the need for clarity in roles, governance structures and accountability for AI decision-making.

If no one owns AI oversight, it will be reported inconsistently, challenged weakly and escalated too late. Clear ownership improves

discipline and makes it much more likely that AI risks and opportunities appear regularly in formal board agendas.

2. Demand an adaptive governance framework

Boards should insist on a governance model that can cope with AI systems that evolve. Static policies are not enough. Directors need assurance that management has controls covering procurement, deployment, monitoring, incident response and periodic review.

An adaptive framework recognises that AI risk changes over time. Vendor models change, internal usage expands, regulations develop and user behaviour shifts. Governance has to move with those changes, not lag behind them.

3. Ask for AI-specific reporting

General technology updates are not enough once AI begins influencing material operations or decisions. Boards should expect reporting on performance, incidents, drift, stakeholder impacts, control effectiveness, third-party dependencies and whether promised productivity gains are actually being captured.

Why it is recommended:
Board-level visibility creates accountability. It also helps directors test whether management is genuinely governing AI systems that evolve, or simply assuming the controls will hold as use expands.

4. Build board capability deliberately

AI for board members should include structured education, not just occasional briefings after an incident or a polished innovation presentation. Directors need enough understanding to challenge assumptions, question metrics and recognise where AI risks intersect with reputation, conduct, compliance and long-term resilience.

Boards that lack AI fluency tend to swing between overconfidence and unnecessary fear. Neither is useful. Deliberate capability building helps directors exercise judgment grounded in governance rather than novelty.

When AI makes sense

The best use cases for senior leaders usually involve decision support rather than decision replacement. AI can help summarise large volumes of material, surface patterns in operational data, draft first-pass analyses and assist with triage across complex information flows. Used properly, it can reduce low-value cognitive load and free people to focus on judgment, exceptions and strategy.

This tends to work best where the domain is reasonably bounded and the source material is well controlled. Internal policy libraries, service knowledge bases, procurement analytics and large-scale

document review are all examples where AI can be useful if controls are strong and outputs are checked.

AI can also help with monitoring AI itself. As organisations expand their use of models, automated oversight tools can support drift detection, anomaly monitoring and pattern analysis that would be difficult to manage manually at scale. That does not remove the need for human accountability, but it does reflect operating reality.

When AI does not make sense

AI is a weak fit where truth, explainability and context matter more than speed, especially when the system is drawing from messy or uncontrolled information sources. It is risky to use generative AI for high-stakes regulatory interpretation, sensitive board communication or formal external statements unless there is rigorous human review and disciplined source checking.

It is also a poor substitute for strategic thinking. AI can generate scenarios and surface patterns, but it cannot define the organisation's appetite for risk, set its values or carry fiduciary accountability. Those remain human leadership tasks.

There are also situations where the economics simply do not stack up. A use case that appears efficient in a demo can become expensive in production once token consumption, retrieval activity, integration effort and quality assurance are included. If the task is

stable, rules-based and already handled well by existing software, AI may add novelty more than value.

Real-world lessons

A practical lesson emerging from current guidance is that boards and executives need to move AI out of the innovation sideshow and into ordinary management discipline. Australian governance resources frame AI as a matter of roles, structures, principles, controls, stakeholder impacts, infrastructure and reporting, which is exactly how mature organisations should approach it.

There is also growing recognition that misinformation and hallucinations are not just technical defects. They become governance problems when organisations deploy systems without data boundaries, verification pathways and human oversight. Public guidance recommends concrete measures such as improving data quality, grounding outputs in reliable sources and maintaining processes that detect inaccurate responses before they create harm.

The organisations applying these lessons well are not treating AI as magic. They are treating it as a capability with real economics, real controls and real accountability. They test before scaling, monitor after deployment and remain willing to stop a use case that no longer fits the value or risk case.

What comes next

NEURAL SEVEN
ADVISORY

The next phase of AI adoption will not belong to the loudest organisations. It will belong to the ones that can make sensible decisions under uncertainty. That means knowing where AI has a credible role, where it should remain tightly constrained, and how to put governance around systems that will not stay still just because the project team has moved on.

For executives, the practical next step is to define an AI posture, review the top use cases through a commercial and governance lens, and put hard controls around data, oversight, productivity measurement and spend. For board members, the next step is to ask whether the current governance model is genuinely fit for governing AI systems that evolve, and whether reporting is strong enough to give directors confidence rather than vague reassurance.

Bring AI onto the formal agenda for the next executive and board cycle. Review existing use, identify hidden cost drivers, test whether time savings are becoming real value, define ownership and require a governance approach designed for change rather than comfort. That is the beginning of AI street smarts for the AI era.