

The Wrong Race

Why coherence, not speed, will decide who wins with AI

Australia is running the wrong race. The leadership task in this moment is not to win an AI race imported from Silicon Valley, but to conduct the tempo of adoption so that the systems on which our prosperity, legitimacy and stability depend emerge from this transition stronger rather than fractured. This essay is written for the directors, executives and policymakers who carry that responsibility.

A director leans forward at a board table in eastern Australia. The proposal in front of her is a \$400 million upgrade to the regional network the utility operates — substations, transformers, power lines. The asset will be in the ground for forty years. Operations and finance have done the modelling, the case is solid, the regulatory determination supports it.

She approves it. Then management moves to the next agenda item: the company's AI strategy. The slides describe a landscape that has shifted materially since the last board meeting, two months ago. New models. New capabilities. New cyber threats targeting industrial control systems.

Halfway through the presentation, she asks the question that has been quietly forming for months: *How do we make forty-year bets when key economic assumptions can shift between quarterly meetings?*

Her question reaches well beyond utilities. It captures something deeper about what it now means to lead anywhere in this period. AI is the most visible disturbance, but it is not the only one. Energy systems, supply chains, geopolitical alliances, the social licence to operate, the workforce contract — all are shifting at once. Most leaders are navigating this turbulence with mental models built for a calmer time. Few have lived through anything quite like it before, and none of us have a settled vocabulary for how to think and act in it.

AI looks like a race. It isn't.

Most public discussion of AI treats it as a race. Who adopts fastest, who captures advantage first, who becomes “AI-native” before the competition.

The framing is everywhere. The US Senate's May 2025 hearing on AI competitiveness was titled, without irony, *Winning the AI Race*¹. Australian boardrooms, government strategies and policy debates have largely imported the same vocabulary — speed, agility, dominance — often without pausing to ask whether it actually fits the situation we are in.

It contains a kernel of truth. Some AI capabilities are advancing at extraordinary speed. Some businesses will

be reshaped by them in months. But as a general framing for leaders, the race metaphor misleads in three ways.

It assumes a single contest with a single finish line. AI is not one technology arriving at one speed. It is many capabilities arriving at different speeds in different parts of the system with different consequences.

It treats speed as the only meaningful response. In a real race, the fastest runner wins. In a complex system, the fastest mover often pays a high price later — when regulators catch up, staff burn out, customers turn, or the assumptions that drove the early sprint turn out to be wrong. The financial markets have begun to register the same concern. Apollo's chief economist Torsten Slok captured the mood in a widely-circulated note² that invoked Robert Solow's productivity paradox of the 1980s: “AI is everywhere except in the incoming macroeconomic data.” Investment is real. The promised economy-wide productivity gains, so far, are not — and a substantial portion of current AI capital may yet be written off.

And the race framing is quietly disempowering. If the race is on and the only question is how to keep up, the harder questions — about whether we are running in the right direction, or even the right race — fall out of view. *Agency disappears under the weight of presumed inevitability.*

The argument of this essay can be put in one sentence. The leadership task is not to win the AI race. It is to conduct the tempo of adoption — so that capability, institutions, people and legitimacy do not pull apart.

Several clocks are running at different levels of the system

What we are calling AI disruption is not one force at one speed. It is several different dynamics, each running at its own pace, each operating at a different level of the system. Naming them is the first step to navigating them.

The capability clock runs fastest, at the global level. New AI models arrive within weeks of each other, with corresponding shifts in what these systems can

credibly do. This is the clock most commentary fixates on.

The geopolitical clock also runs fast at the global level — and is the one most likely to surprise leaders who haven't been watching it. AUKUS Pillar Two now explicitly covers AI, autonomy, quantum and advanced cyber. Compute itself has become a strategic resource, with US export controls on advanced semiconductors actively reshaping global supply. Authoritarian competitors are unconstrained by the expectations of public legitimacy that liberal democracies must honour. And much of the contest is happening below the threshold of declared conflict — what defence analysts call grey-zone activity. This clock is global, fast, and frequently invisible to commercial leaders until it intersects their world.

The industry restructuring clock runs slower than the headlines suggest. Capital moves over weeks; real industrial reordering takes years. The evidence is now accumulating. MIT's NANDA initiative, in its July 2025 *GenAI Divide* report³, found that only around 5 per cent of enterprise generative-AI initiatives had translated into measurable bottom-line returns — a finding that has invited methodological debate but is broadly consistent with other recent data. PwC's 2026 *AI Performance* study⁴, surveying 1,217 senior executives across 25 sectors, found that around 74 per cent of AI's economic value to date has been captured by just 20 per cent of organisations. The Reserve Bank of Australia's November 2025 firm-level survey reported AI adoption across Australian businesses as “relatively piecemeal” and often employee-led rather than employer-led. Capability is moving in months. Industries are restructuring in years.

The regulatory and legitimacy clock runs at the national level, slower again. Australia released its Voluntary AI Safety Standard in 2024 and is consulting on mandatory guardrails. The European Union's AI Act is commencing in phases through 2026 and 2027. These frameworks deliberately trail the capability clock — and the legitimacy that underpins them is slower still.

The functional adaptation clock runs at the institutional level: the speed at which an organisation can modernise its data, restructure its incentives, retrain its people and reconfigure its operating model. For most organisations this clock runs much slower than commentary suggests — until an AI-enabled cyber attack on operational systems compresses years of decision-making into hours.

The human adaptation clock runs at the individual and community level. Mostly slow — reskilling takes years, identity adjustment longer. But not always slow: COVID showed humans can shift behaviour quickly when threat is high and visible. The harder problem is involuntary adaptation. Many of the manufacturing workers displaced in the early 1990s never returned to comparable work. AI-driven labour shifts risk repeating that pattern unless social systems adapt with them.

The sensemaking clock runs across all the others, and may be the master constraint. It is the speed at which leaders, collectively, can read what is happening well enough to act wisely — not any individual leader, but the leadership ecosystem of board, executive, regulator, workforce and wider society. When this collective reading falls behind capability, organisations make confident decisions on outdated frames. The risk is not too little information. It is too little shared judgement, distributed across the people whose judgement actually has to converge.




Figure 1 — Different clocks, different levels. AI is not unfolding at one speed. It is unfolding through several different clocks running at different speeds and at different levels of the system. Sensemaking runs across them all. Where the gaps grow too wide, instability follows.

The point of the framing is not to synchronise these clocks. It is to see them. Changes at higher levels of the system — geopolitical realignments, capability discontinuities, capital flow shifts — produce ripple effects through the levels below. Strategic decisions that ignore the level at which a change is occurring will misjudge both its pace and its consequences.

The tale of two actors: a power utility and a manufacturer

Two examples will make this concrete, and reappear through the rest of the essay.

The utility. A regional electricity distribution business in eastern Australia. Around \$2 billion in assets — poles, wires, substations, customer connections. Regulated by the Australian Energy Regulator on five-year revenue cycles. A workforce of around a thousand: linesmen, control room operators, engineers, customer service teams, a small but growing data and analytics function. Some substations were commissioned in the 1970s. AI looks like both an opportunity (network optimisation, demand forecasting, predictive maintenance) and a serious risk (cyber threats to operational technology, regulatory and political fallout if

deployment goes wrong, unrest in a heavily unionised workforce).

The manufacturer. An owner-operated specialty food processor in regional Queensland. Sixty staff. Tight margins. The owner has run the business for fifteen years. AI has felt mostly like noise — until a major customer indicates that within two years it will require suppliers to demonstrate AI-enabled quality control and traceability. Suddenly AI is not optional. But the owner has no internal data team, limited capital, and no obvious path through the maze.

The argument that follows is shaped to work for both. But the right action looks very different in each case.

The deeper risk: success that fractures the system around it

In a hospital boardroom in another Australian city, the same week, a different leader is reviewing the early results of an AI-assisted triage system rolled out across the emergency department. Wait times are down. Throughput is up. The clinical-quality dashboard is the best it has been in three years. The board is pleased.

Two things have not yet appeared on the dashboard. The nursing union has begun raising concerns it has not yet formalised. The state regulator has sent two letters that have not yet been escalated. A young clinician who left last month told a colleague, on her way out, that she no longer felt she was being trained — she was being audited. None of these signals has reached the room where the dashboard is being celebrated. They will, in time. By then, some of them will be irreversible.

This why the multi-clock, multi-level framing matters. Because organisations — and the wider systems they sit inside, of regulators, customers, workforces, communities and political constituencies — can optimise one part of themselves while quietly fracturing the whole. The bit that gets faster looks like progress. The bits that don't keep up look like problems. Until something gives.

Robodebt is an obvious recent Australian example. An automated welfare compliance system was introduced to recover money faster. In narrow operational terms, it worked. The Royal Commission later found that the scheme had issued unlawful demands, caused substantial human harm, and ultimately cost the Commonwealth far more than it ever recovered. The technology, such as it was, did what it was asked to do. The surrounding system — legal, administrative, ethical, political — fractured. Function had been optimised. Coherence had been lost.

Robodebt predates the current generation of AI, which is precisely the point. The pattern it illustrates — one part of a system racing ahead of the rest, with consequences that arrive late and at scale — is the pattern AI now risks repeating across many more domains, more quickly.

Consider the utility. Suppose its analytics team deploys an AI system that materially reduces network

outages — a real operational win. But the system is opaque to the regulator, distrusted by the workforce who feel decisions are being made over their heads, and built on customer data the company has not clearly explained it is using. In raw operational terms, success. Across the wider system, an accumulation of small fractures: regulatory tension, workforce disengagement, quiet erosion of social licence. Each fracture is defensible in isolation. Together, they generate fragility.

Or take the food processor. Pressed by their customer, they buy a cheap AI quality-control system from an overseas vendor. It works well at first. But it doesn't integrate with their production systems, the staff don't understand how it makes decisions, and when it fails it fails in ways that take half a day to diagnose. The owner is now busier, not less busy. The investment has solved the customer's compliance question while creating an operational dependency the business cannot maintain. Local optimisation. System fragility.

This is the deeper risk in the AI moment. Not that AI fails. That it succeeds — locally, narrowly, impressively — in ways the surrounding system is not ready to absorb.

It also exposes a missing question. Almost no current commentary asks the question that should sit alongside *how fast?* — namely, *to what end?* Vast capital, attention and effort is now flowing toward AI without a clearly articulated picture of what we want it to add up to. The standard line is that AI will solve our biggest problems. But by what design, on what timeframe, in whose interests? In the absence of an articulated purpose, “fast” is all we get.

When coherence pays — and when speed wins

A fair challenge to all this: when slower competitors invoke “coherence”, isn't it sometimes self-serving — slowness rebadged as wisdom?

Sometimes, yes. In some markets, speed wins outright. Consumer-facing platforms with strong network effects — ride-share, social media, food delivery — reward whoever scales first. Uber's early posture in Australia could fairly be called incoherent: it ignored regulators, antagonised drivers, and disregarded local norms. It also won. By the time the system caught up, it was the system. There are markets where being early matters more than being right.

But these markets are narrower than the prevailing rhetoric suggests. They do not describe the conditions most Australian institutions actually operate in.

Coherence pays in three settings, and between them they cover most of the economy and almost all of the public sector.

The first is where assets and commitments are long-lived. The utility's substations will be in the ground for forty years. A hospital's clinical workflows take a decade to redesign. A university's research infrastructure outlasts every vice-chancellor who oversees it. Lock in commitments based on today's assumptions, and you may discover — too late — that those assumptions have

moved. Coherence here is the discipline of refusing to commit irreversibly to ground that is still shifting.

The second is where you need public trust to operate at all. Regulators, banks, insurers, public agencies, healthcare providers — and increasingly the digital platforms mediating critical infrastructure — depend on continuing social authorisation to operate. Lose that authorisation and you lose your business model, not just your reputation. Fast AI deployment that erodes the trust on which the licence to operate depends can take down the operator before the technology has matured.

The third is where complexity itself punishes premature commitment. In environments where outcomes depend on many interacting agents and second-order effects matter, the leaders who lock in early often pay the highest correction costs later. Coherence in these settings is steady direction with options preserved: moving with intent, but keeping the ability to reconfigure as understanding improves.

What must be acknowledged is the difference in agency. Different actors have very different room to move.

The utility has bandwidth, capital, regulatory weight and institutional memory. It can — if it chooses — deliberately conduct its tempo: adopt some applications quickly, hold others back, invest in workforce capability, manage regulatory engagement, influence its own operating environment. Coherence is a real choice.

The food processor cannot. Its margins are too tight, its team too small, its influence too modest. Coherence here is not about commanding the tempo. It is about staying directionally consistent under pressure: choosing carefully which fast-moving currents to ride, which to step back from, which signals to filter out so the noise doesn't drown out the owner's own intent. The work is the same in nature. The freedom of action is not.

This matters because the dominant framing is shaped by the actors with the most agency, and then exported as universal advice. A line like “every business needs an AI strategy” lands very differently on a utility CEO than on a food-processor owner-manager — but each is told to feel the same urgency.

The claim, properly stated, is not “be coherent rather than fast”. That sets up a false choice. It is that, in most contexts that matter, slowing down a little — sometimes only marginally — buys a more sustainable speed over time. *Coherence is not the opposite of speed. It is what allows speed to compound rather than self-destruct.*

The institutions on which Australian prosperity, legitimacy and stability most depend operate in conditions where coherence pays. The “AI race” framing leads them to fight the wrong fight on the wrong terrain.

The other clock: capital, concentration and sovereignty

There is one clock that's not yet been named, and it may be the one that matters most over the next decade. The structural-economic clock — the financial architecture behind the AI build-out, and the dependence it creates.

The numbers are sobering. Capital expenditure on AI infrastructure by the largest US hyperscalers and model developers ran into the high hundreds of billions of dollars through 2025 and 2026, financed increasingly through circular relationships in which chip suppliers, cloud providers, model developers and their major customers are also each other's principal investors. The Bank of England and the International Monetary Fund have both publicly flagged the concentration and valuation risk this creates in global equity markets. The investment is now substantially ahead of demonstrated enterprise returns. Slok's observation, on this reading, is not just an aside about productivity statistics. It is an early signal that the build-out and the value capture are running on different clocks — and that the gap between them is being underwritten by capital markets that may not absorb the correction quietly.

For Australia, this opens a strategic question the AI race framing systematically obscures. The question is not only what AI does inside our organisations. It is what our dependence on a small number of mostly US-headquartered providers does to our cost base, our optionality and our sovereign capability over a decade.

The dependence is forming quickly and quietly. The cloud platforms our banks run on, the cybersecurity services our utilities depend on, the foundation models our public agencies are beginning to embed into service delivery — these are increasingly supplied by the same handful of providers. Token costs that today appear incidental will become discrete line items on the P&Ls of every organisation of meaningful scale. The providers raising those costs will, in many cases, be the same providers our critical infrastructure now runs on. A nation that finds itself renting most of its strategic capability from offshore, at prices it does not set, has fewer choices than one that conducts its own tempo of dependence.

This is, at heart, the same coherence problem — only at national scale. A country can optimise the adoption of AI inside its firms and agencies while quietly fracturing its longer-term strategic autonomy. The function gets faster. The dependence deepens. The cost line that did not exist this year becomes structural by the next.

The point is not autarky (self sufficiency), which is neither possible nor desirable. The point is that the leadership task at national level is the same as at organisational level: *read the clocks, protect reversibility, sequence for compounding*. Read the structural-economic clock alongside the capability and regulatory ones. Protect reversibility in commitments — sovereign compute, critical-infrastructure dependencies, public-sector model choices — that will be very difficult to unwind once locked in. Sequence national investment so that what we build first makes what we build next more powerful, not more constrained.

The AI race framing makes this conversation almost impossible to have. It assumes the race is the only contest that matters and that running it is the only available posture. The harder and more useful question

is whether Australia will conduct its own tempo, or accept the one being conducted for it.

Coherence is conducted, not commanded

It is worth being precise about what coherence actually means, because the word can be misread in ways that convey exactly the wrong intent.

Coherence is not alignment to a single resolved plan. It is not central planning under a more fashionable name. It is not the synchronisation of every clock to a master schedule.

It is something more disciplined and more nuanced.

Coherence, in the sense used here, is *moving in the same overall direction even as different parts adapt at different speeds and in different ways*. It is holding intentions, values and strategic philosophy steady while allowing genuine learning and local variation in how those intentions get expressed. It is appropriate to the system level, the timeframe, the moral terrain and the practical realities of the organisation involved.

A coherent organisation can move fast in some places and pause deliberately in others, provided the movement and the pause make sense within the same overall frame. An incoherent organisation mistakes activity for progress.

This distinction matters because the natural reflex of many leaders facing acceleration is to reach for tighter control: more reporting, more alignment, less flexibility. In complex environments, that reflex makes things worse. It slows the parts of the system that need to adapt, while doing nothing to address the levels at which the dangerous incoherences are forming.

A more useful image is that of a conductor rather than a commander. The conductor sets the *tempo* and listens for what the system is doing; intervenes where needed; lets sections lead where they are stronger; and holds the overall shape of the piece. The work is judgement, not control.

That word — *tempo* — is worth keeping in mind. A race rewards raw speed. A complex system rewards timing, sensitivity to feedback, and the disciplined sequencing of moves so that apparent restraint at one moment generates stronger acceleration later. *Coherent actors conduct tempo. They do not merely run faster.*

THREE QUESTIONS FOR THE BOARD

1. Are we and our key stakeholders reading this disturbance the same way?
2. Where in our strategy are we committing long-lived resources to assumptions unlikely to survive the next phase of this disturbance?
3. Where in our operating environment are the clocks pulling apart in ways that could fracture our coherence?

The board's work: where to look, how fast to move

For boards, the central challenge in this moment is attention and pace. Directors cannot operate the technology. They cannot rebuild the organisation. What they can do is decide what the institution must pay attention to, how quickly it needs to move, and where the risks of moving too slowly are outweighed by the risks of moving too fast.

The Australian Institute of Company Directors and the University of Technology Sydney⁶ have rightly observed that investment in AI systems across Australian organisations has not been matched by investment in AI governance. The diagnostic questions that follow are designed to close part of that gap — not as an audit checklist, but as the live questions a board should be carrying through every relevant agenda item over the next two years.

First — are we and our key stakeholders reading this disturbance the same way? Where the answer is no, what is our plan to close the gap before it shapes outcomes for us? In the utility's case, the question is whether the AER, the workforce, the executive and the board are all reading the same disturbance — or each looking at a different part of it. Sensemaking, distributed across the people whose judgement actually has to converge, is the master constraint.

Second — where in our strategy are we committing long-lived resources to assumptions unlikely to survive the next phase of this disturbance? And what is the cost of keeping those commitments reversible? The utility's \$400 million network upgrade is the live example. The technology that runs on that asset will turn over many times during its life. The asset itself cannot. Some reversibility is cheap. Some is expensive. Both deserve a number.

Third — where in our operating environment are the clocks pulling apart in ways that could fracture our coherence? Between capability and regulation, between regulation and public legitimacy, between executive ambition and workforce capacity, between national interest and commercial logic. What is the earliest signal we would see that this is happening — and who is responsible for watching for it?

These are governance questions, not technology questions. They surface the upstream work that determines whether downstream execution is even directed at the right problem. Most current AI advice arriving in boardrooms sidesteps them entirely — heavy on adoption frameworks, light on the harder strategic question of *what* the organisation is trying to remain coherent toward.

The executive's work: what to build before what

For executives, the challenge is different. It is execution under bandwidth constraint.

Most executive teams already understand, broadly, what AI can do. What they lack is time, organisational machinery, and sometimes the political space, to do the patient work that turns AI from announcement into compounding advantage. Building data foundations. Restructuring workflows. Reskilling staff. Creating governance disciplines. Managing legacy IT and operational technology systems that were not designed for any of this.

The executive's question is therefore not "what should we attend to?" but "what must we sequence first so that later work compounds rather than collapses?"

That sequencing question — what to build before what — is the harder strategic act for most enterprises right now. Capability is evolving faster than implementation cycles. Fiscal conditions are tightening. Workforces are stretched. Almost every organisation is, in the familiar phrase, rebuilding the plane while flying it. In those conditions, the executives who deliver are the ones clearest about what *not* to do this year.

Sequencing also requires protection. Sometimes — uncomfortably — protection from boards that have themselves become impatient under external race pressure. A good board chair recognises this and creates the space for it.

There is a paradox here that experienced executives will recognise. In conditions of acceleration, the right move for some organisations is to slow down their consumption of AI rather than speed it up. Not to back away from the technology, but to take longer to decide where to apply it; to demand sharper diagnostics; to refuse premature lock-in; to invest in the sensemaking capacity that complex periods reward. The instinct to demand quick answers is, in such conditions, often a symptom of pressure rather than a sign of clarity. It substitutes activity for thought.

For the food processor, the executive question is starker. With limited bandwidth, what *one* AI investment matters most this year? Not three. Not five. One — the one that meets the immediate customer pressure while building data foundations that compound for the next decision. Sequencing under constraint is harder than sequencing with resources. But the underlying discipline is the same.

THE THREE DISCIPLINES

Read the clocks > see the disturbance

Protect reversibility > preserve choices

Sequence for compounding > choose what builds

Agency in an age of acceleration: react, or create?

There is a third posture most discussion of AI overlooks.

Most commentary assumes leaders are choosing between adopting AI faster and adopting it more slowly.

But some organisations are not adapting to AI at all. They are being built around it.

The most interesting AI-driven transformations of the next decade may not come from incumbents successfully retrofitting AI onto legacy systems. They may come from new organisations designed from the ground up around what AI now makes possible — leaner, less burdened by legacy assets, free to redesign workflows and business models from scratch. A digital-by-design organisation, by definition, doesn't face the multi-clock coherence problem in the same way: its technology, operating model and strategic intent were aligned from the start.

This needs a candid qualifier. Not every organisation can become this. The utility cannot. It owns physical assets that will outlast multiple technology generations and serves communities that depend on uninterrupted supply. Even if it could redesign itself digitally tomorrow — which it can't — its obligations would not allow it to.

But the food processor might. With sixty staff, modest fixed assets, and a customer base willing to pay for traceability and quality, it is in a position the utility is not. It could choose to adopt AI defensively, narrowly, customer-by-customer. Or it could choose something braver: redesign its operating model around what AI now enables. Real-time process control, predictive supply chains, leaner operations. In a bolder version of the question, the food processor stops thinking of itself as a food processor. AI is converging with synthetic biology, advanced materials and precision fermentation in ways that are quietly redrawing the boundary between making and designing food. Within a decade, the most successful "manufacturers" in this space may not be those who automated their factories. They may be those who reconceived what the factory is for.

For some smaller, less-burdened actors, the most coherent response to AI is not to adapt at all. It is to redesign. To create rather than to react.

This points to something larger. AI is not the only foundational technology shifting at once. Energy, advanced materials, biotechnology, transport — all are moving on their own clocks, and their convergence is part of what makes this period so disorienting. In a world where multiple foundational technologies are reshaping in parallel, the most coherent posture for some actors is not to compete on the existing terrain but to create a new one.

For most readers of this essay — directors, executives, policymakers — the answer will not be a clean choice between react and create. It will be a more careful mix of defending what must be defended, adapting what can be adapted, and creating where genuine new possibility opens. The mistake is to assume the choice is binary, or that the answer is the same for every actor.

The questions every leader should now be asking, alongside *how fast?*, are *to what end?* and *what is the right kind of organisation to pursue that end at all?* These are the questions that surface real agency. They

are exactly the questions the AI race framing systematically fails to ask.

The defining test of leadership

Imagine the same boardroom we started in, eighteen months from now. Our director is back. The next big asset decision is on the table. A different cohort of AI tools is in play. The race rhetoric is louder, not quieter. What does she do?

If she has been listening to the dominant message of the last three years, she will feel the familiar pressure to commit, to move, to be visibly acting. If she has been thinking systemically, she will do something less intuitive but more powerful. She will hold the room — not against AI, but against the framing that says only velocity matters.

She will do three things her predecessors largely didn't.

Read the clocks. Before agreeing to a tempo, she will ask whether she and the parts of her environment that matter — regulator, workforce, investors, customers, executive — are reading the disturbance the same way. Where they are not, the strategic risk is incoherence, not slowness. She will know which gaps are most likely to fracture the system, and who is responsible for watching them.

Protect reversibility. Before committing long-lived resources, she will ask which assumptions are most likely to shift, and what it costs to keep the decision reversible until those shifts become visible. She will stop confusing irreversibility with conviction.

Sequence for compounding. Before agreeing to scope, she will ask what *not* to do this year so that what is done compounds rather than collapses. She will protect the executive's bandwidth for the patient work that turns AI from announcement into advantage.

These are not technology disciplines. They are leadership disciplines, applied under harder conditions.

The leadership task, restated, is not to win the AI race. It is to conduct the tempo of adoption — so that capability, institutions, people and legitimacy do not pull apart.

The wager is not that coherence is the opposite of speed. It is that, in the contexts that matter most for Australia's prosperity, legitimacy and stability, coherence is the form of speed that compounds rather than collapses. *Fast movers may capture moments. Coherent actors are how eras get made.*

The defining test of leadership in the years ahead will not be how quickly we adopt AI. It will be whether we have the discipline and the imagination to lead through this transition so the systems we depend on come out of it stronger, not fractured.

That is a demanding ambition. It is also entirely within our power.

Written by Dr Nick Fleming, Managing Director of Innergise Pty Ltd, who is a non-executive director, educator in strategy and complex problem-solving, and specialist in systems thinking and leadership.

End Notes

1. The US Senate Commerce Committee hearing of 8 May 2025 was titled Winning the AI Race: Strengthening US Capabilities in Computing and Innovation. Hearing record: <https://www.commerce.senate.gov/meetings/winning-the-ai-race-strengthening-u-s-capabilities-in-computing-and-innovation/>
2. Apollo Global Management's chief economist Torsten Slok, *Waiting for the AI J-Curve*, Apollo Academy Daily Spark, February 2026: <https://www.apolloacademy.com/waiting-for-the-ai-j-curve/>
3. MIT NANDA, *The GenAI Divide: State of AI in Business 2025*, July 2025. The headline finding — that approximately 95 per cent of enterprise generative-AI initiatives have not yet delivered measurable bottom-line returns — has attracted some methodological commentary on sample size (300+ public initiatives, 52 interviews, 153 survey responses), but is broadly consistent with other recent enterprise data. Report: https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf
4. PwC, *2026 AI Performance Study*, released 13 April 2026, surveying 1,217 senior executives across 25 sectors and multiple regions: <https://www.pwc.com/gx/en/news-room/press-releases/2026/pwc-2026-ai-performance-study.html>
5. Reserve Bank of Australia, *Technology Investment and AI: What Are Firms Telling Us?*, RBA Bulletin, November 2025, drawing on a survey of approximately 100 medium and large Australian firms. The RBA's own characterisation of AI adoption is "relatively piecemeal" and "often employee-led rather than employer-led": <https://www.rba.gov.au/publications/bulletin/2025/nov/technology-investment-and-ai-what-are-firms-telling-us.html>
6. The Australian Institute of Company Directors and University of Technology Sydney Human Technology Institute observation on the gap between AI investment and AI governance is drawn from the AICD's *Director's Guide to AI Governance* and associated commentary by Professor Nicholas Davis: <https://www.aicd.com.au/innovative-technology/digital-business/artificial-intelligence/governance-of-ai.html>
7. AUKUS Pillar Two coverage of AI, autonomy, quantum and advanced cyber is documented in the US Congressional Research Service report *AUKUS Pillar 2 (Advanced Capabilities): Background and Issues for Congress*: <https://www.congress.gov/crs-product/R47599>

Copyright © Innergise Pty Ltd, 2026.

This document contains general information only. Before taking any action or making any decision that may affect you, your business or finances, you should consult a professional advisor. Innergise shall not be responsible for any loss whatsoever sustained by any person that relies on this communication.