Getting to better, faster, and safer

Deploying AI safely while capturing the productivity upside

A presentation from Strategic Design by Argon & Co

November 2026

Strategic Design
by Argon&Co formerly Bevington Group

© Copyright Bevington Consulting Pty Ltd. All rights reserved

Pressure to move faster

The false velocity of moving faster at higher risk

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

Australians use AI widely, yet their trust in AI ranks among the lowest worldwide

SHARE OF AUSTRALIAN ADULTS WHO AGREE WITH EACH STATEMENT

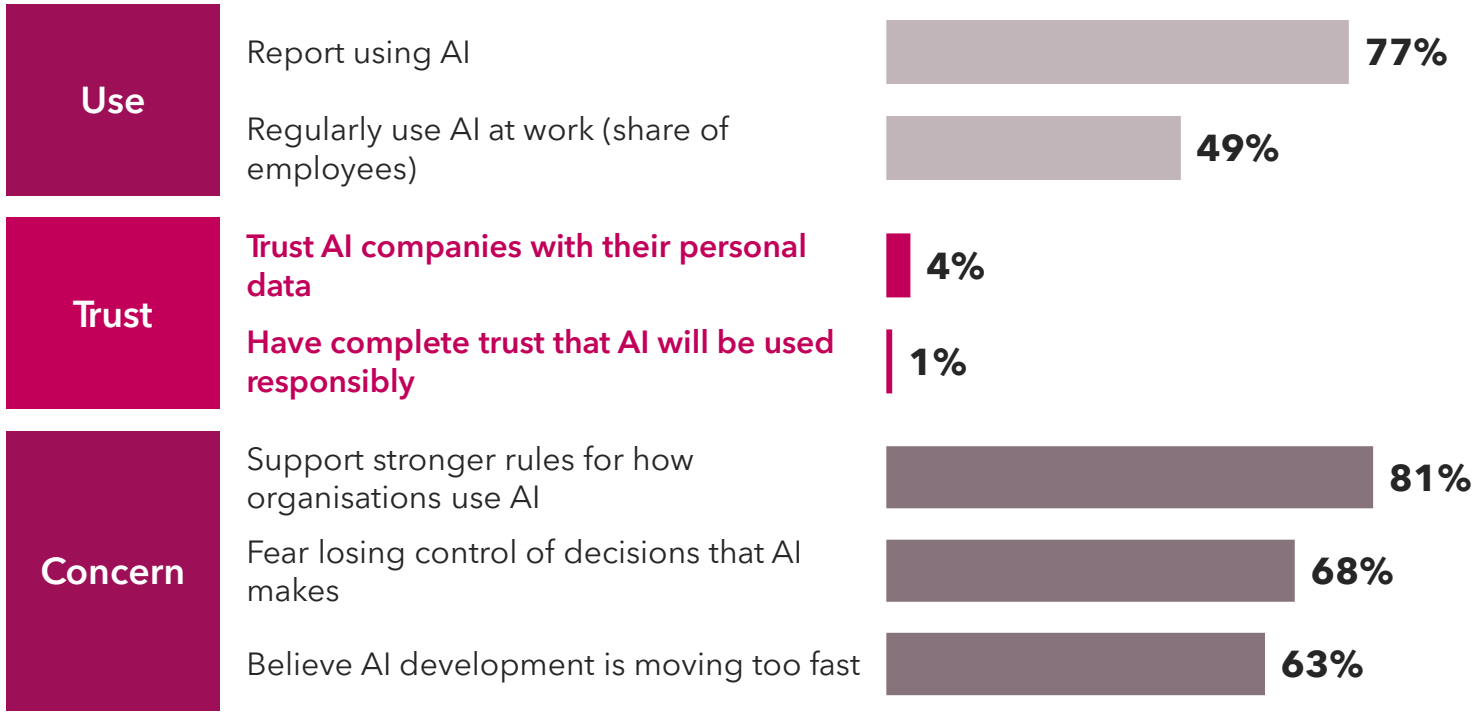

AI sentiment index, out of 100

52 Australia, equal lowest of 23 countries

66 Average across all 23 countries

Paradox of use and trust

High adoption does not automatically produce trust. Australia combines widespread AI use with equal-lowest sentiment globally, amid public concern over errors, deepfakes, privacy, and corporate behaviour

"I don't trust it. On the other hand, I use it a lot."
Sally McManus, ACTU Secretary, June 2026

Staff and customers will not give AI their trust by default. Leaders need to show how people check, explain and challenge each AI-assisted decision

Sources: EY Global AI Sentiment Study (2026); OAIC Australian Community Attitudes to Privacy Survey (2026); Tech Policy Design Institute (2026); Minderoo Foundation (2026); University of Melbourne and KPMG (2025); ABC News (2026)

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

Australian suspicion of AI has moved from opinion polls into courts, workplaces, and suburbs

PUBLIC PUSHBACK AGAINST AI, SEPTEMBER 2025 TO MAY 2026

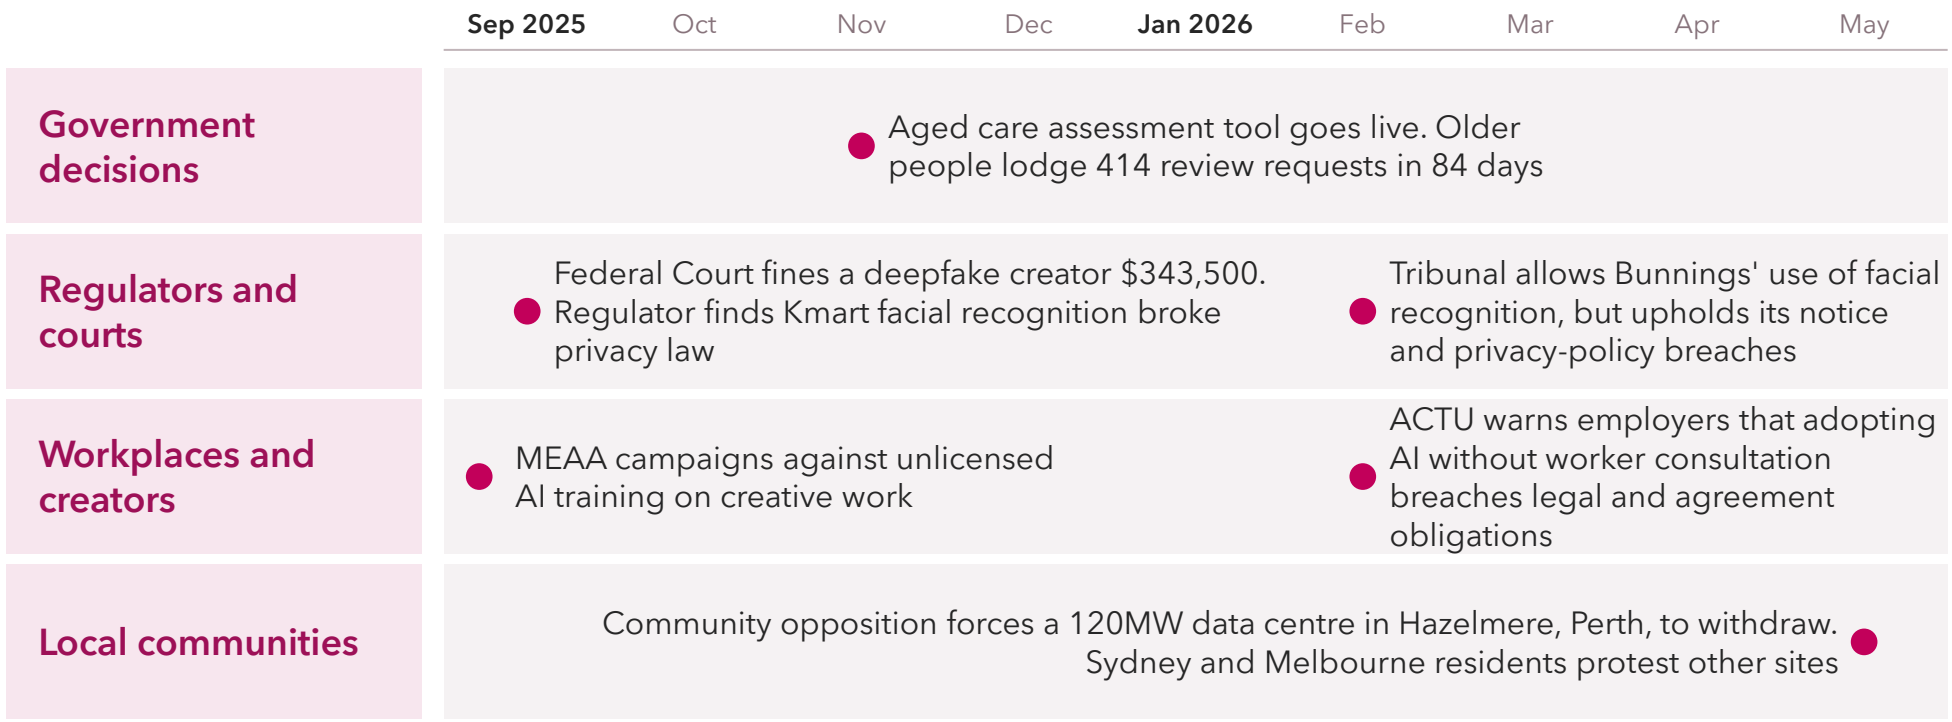

June 2026 warning

Assistant Minister Andrew Charlton told business leaders that Australia faces a US-style backlash against AI without community protections

Low trust puts up to $115 billion a year in generative AI benefits by 2030 at risk (Microsoft and Tech Council of Australia)

Sources: ABC News (2026); eSafety Commissioner (2025); OAIC (2025); MEAA (2025); ACTU (2026); The Guardian Australia (2026); Australian Ageing Agenda (2026); Bunnings [2026] ARTA 130; Charlton (2026); Microsoft and Tech Council of Australia (2023)

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

Some challenges with AI deployment

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

In an Anthropic trial, experts adopted an AI agent's randomly assigned views 86% of the time, then presented them as their own

86%

of analytical choices matched the agent's randomly assigned suggestion



The agent stated its view first on 77% of decisions. Experts contested the agent on 23%

Years of experience, expertise and trust in AI did not change the result

Anthropic ran a randomised trial. 50 data experts used an AI agent to analyse farm data and advise a board. A coin flip set which of two sound options the agent favoured on each decision

Measure	What experts reported	What the session records showed
Ownership of the decision	**52%** rated the final advice as more their own than the agent's	**82%** of final advice was predictable from the agent's leans alone, against 61% by guesswork
Reasons for each choice	**15%** mentioned the agent when asked to justify their choices	**77%** repeated the agent's talking points, at times almost word for word
Confidence and awareness	**Equal confidence** for experts who followed every lean and those who did not	**0%** suspected that the agent had been set to favour one side

Source: Herdağdelen, Joren and Tamkin (2026), Who's aligned to whom? How AI agents shape experts' research judgment calls, Anthropic

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

Speed in preparation is lost if every decision still waits for the same people

Parts of each decision must move to AI, with controls set by risk and reversibility

Stages of the decision cycle

Define the choice → **Prepare evidence and analysis** → Challenge and compare → Decide and authorise the commitment → Review the results

Measured gains from AI in preparation and analysis

40% less time to complete professional writing tasks
(Noy and Zhang, 2023)

18% improved writing quality
(Noy and Zhang, 2023)

30% more output per hour for less experienced employees
(Brynjolfsson et al., 2025)

25% faster work on consulting tasks suited to AI capability
(Dell'Acqua et al., 2026)

The most experienced workers gained little from the same tool (Brynjolfsson et al., 2025)

Common human responsibilities

- People test each option against the goals, risks, and trade-offs
- A named leader commits the resources and is accountable for the result
- People compare the results with the expected outcome and change course when needed

Research on decisions made with AI

A systematic review and meta-analysis of more than 100 experiments found that human and AI pairs did worse than the better of the two alone on decision tasks. **Two failure points include**:

1. **Automation Bias / Over-reliance**: Humans accept incorrect AI outputs
2. **Algorithm Aversion / Under-reliance**: Humans reject correct AI recommendations

The same pairs did better on content creation tasks (Vaccaro et al., 2024)

Faster work and more decision papers do not establish that a better choice is made

Sources: Noy and Zhang (2023); Brynjolfsson, Li and Raymond (2025); Dell'Acqua et al. (2026); Vaccaro, Almaatouq and Malone (2024)

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

Fluent AI output narrows the options people consider and hides errors, so people check less

	Problem	Outcome	Research indicates…
1	**Fewer options considered**	People anchor on AI suggestions, so a group's options converge	In a trial of about 300 writers, stories became more alike (Doshi and Hauser, 2024)
2	**Less precision**	Fluent output hides errors, and people check less	On a problem outside AI's capability, employees scored lower using AI (60-71%), compared to those not using AI (84%) (Dell'Acqua et al., 2026)
3	**Workslop**	Plausible but unchecked content moves into decision papers	Deloitte's A$440,000 review for DEWR included a fabricated court quote and false references. Deloitte agreed to repay the final instalment (Associated Press, 2025)
4	**Amplified bias**	An assistant that agrees with its user reinforces that user's errors	11 AI models affirmed users' actions 49% more often than people did, and users trusted the agreeable answers more (Cheng et al., 2026)

Each problem is a weakness in how people decide, and AI can make it faster and harder to see

Sources: Doshi and Hauser (2024); Dell'Acqua et al. (2026); Associated Press (2025); Cheng et al. (2026)

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

What the law already requires

Human accountability, explainability, and transparency

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

While AI plays a role in decision-making, most regulatory environments insist that people stay accountable for the outcome



Sources: Corporations Act 2001; FAR Act 2023; APRA (2023, 2026a, 2026b); ASIC (2024); DISR (2025); DTA (2025); AICD and HTI (2026); EU (2024); Royal Commission into the Robodebt Scheme (2023)

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

A duty cannot be passed to a tool, so the decision-maker must be able to explain and defend the choice

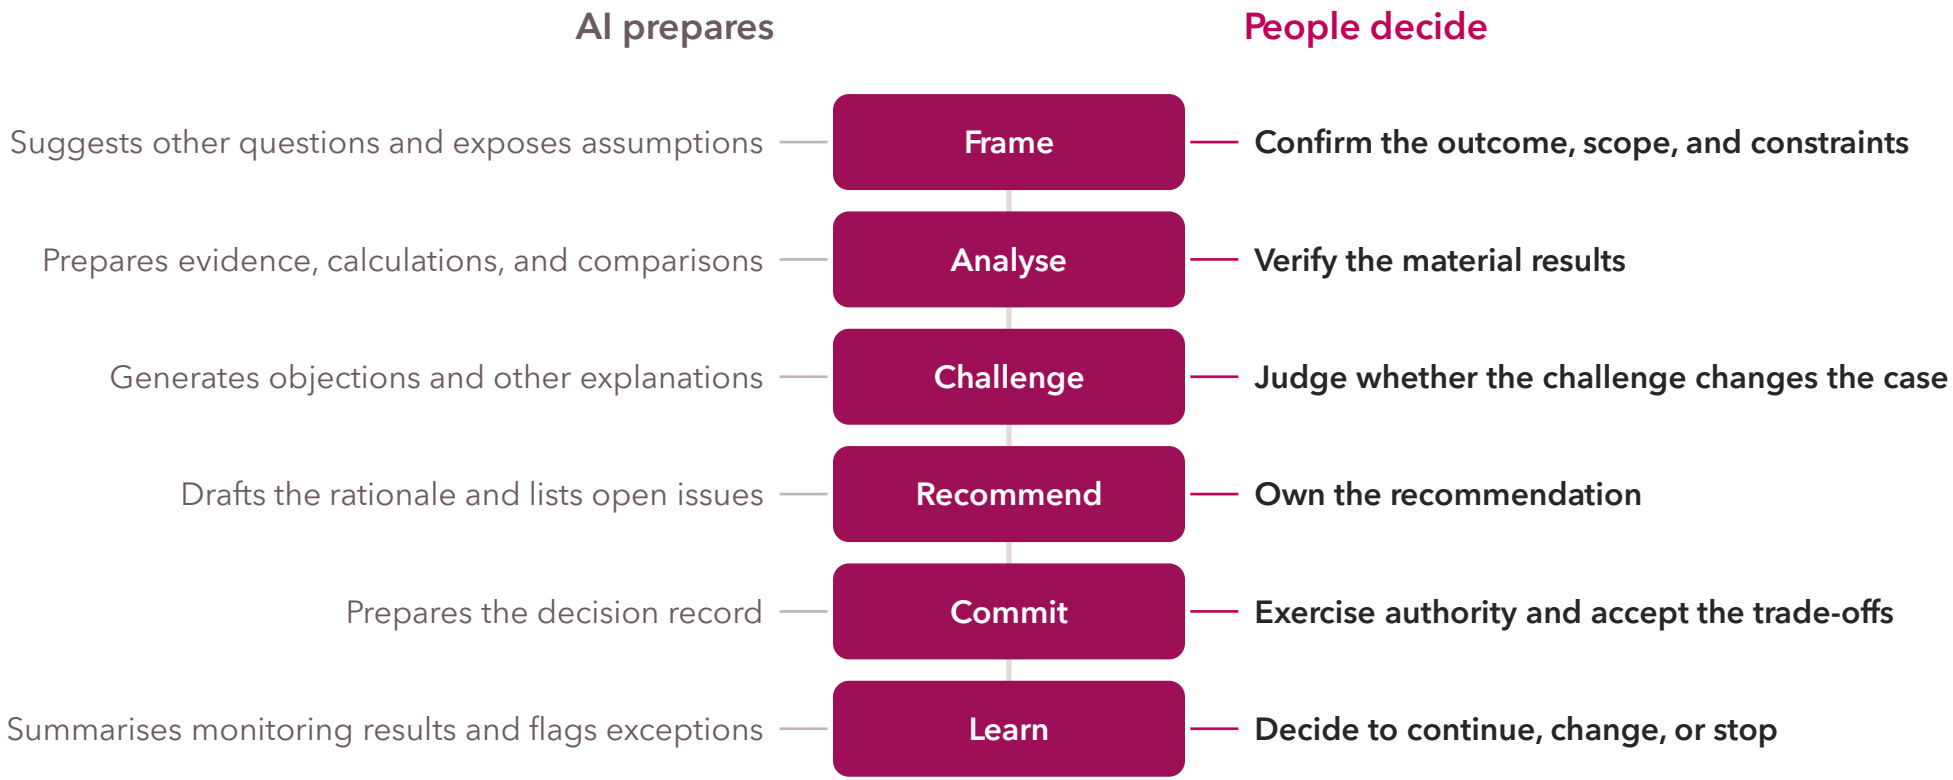

Public example: Air Canada, 2024

A tribunal held the airline liable for wrong fare advice from its chatbot, and rejected the claim that the bot was a separate entity

Sources: Adapted from Sebastian et al. (2026); Moffatt v Air Canada (2024)

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

From 10 December 2026, privacy policies must disclose program-assisted decisions, even with human sign-off

Scope test under APP 1.7, with all three conditions met

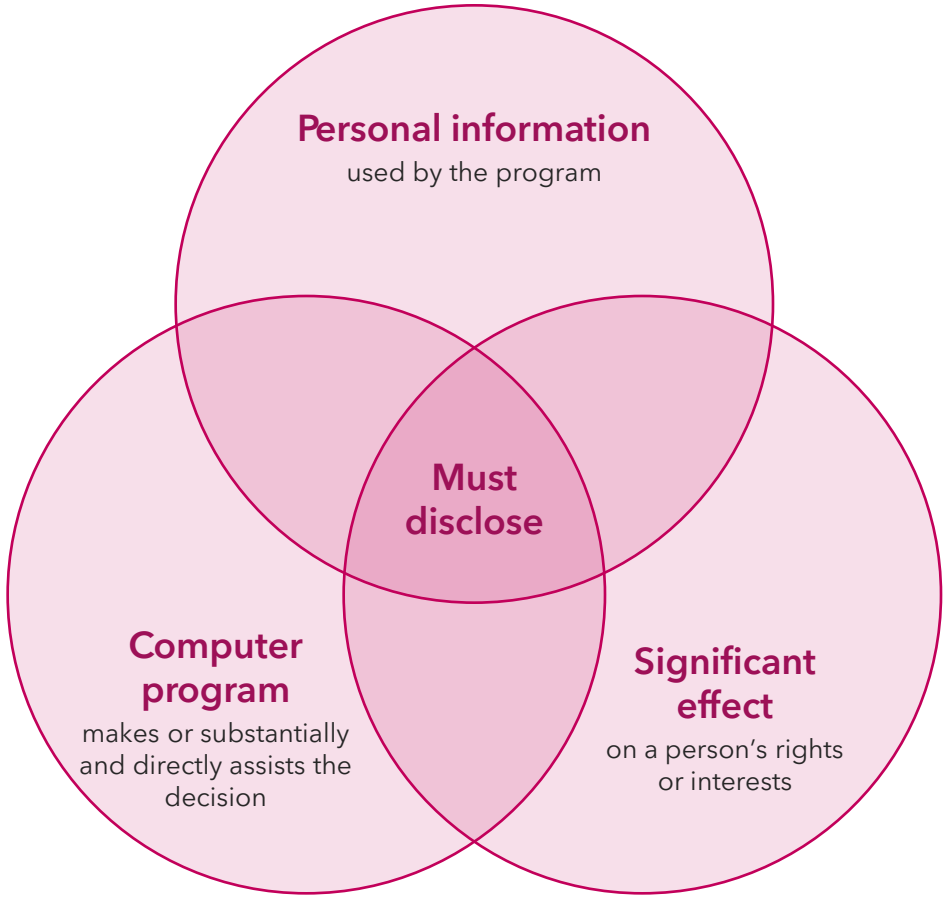

Privacy policy must state

a) The kinds of personal information that programs use
b) The kinds of decisions that programs make alone
c) The kinds of decisions where programs substantially and directly assist

Common assumption	Correct reading
A person signs off, so the tool is out of scope	Sign-off does not remove the duty when a program substantially assists
Recruitment tools are exempt as employee records	Applicants are not employees, so screening tools are in scope
The vendor of the tool holds the duty	The organisation that uses the tool holds the duty
Policy wording carries little legal risk	A non-compliant policy can attract penalties of up to $364,000

Sources: Privacy Act 1988 (Cth) Schedule 1, APP 1.7 to 1.9; OAIC, APP Guidelines Chapter 1, Version 2.0 (2026); Allens (2026); Maddocks (2026); Linklaters (2026)

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

Decision Controls to move faster

Effective controls for AI-assisted decision making

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

AI's greatest risks cluster in a few places – mainly autonomy, security, and trust

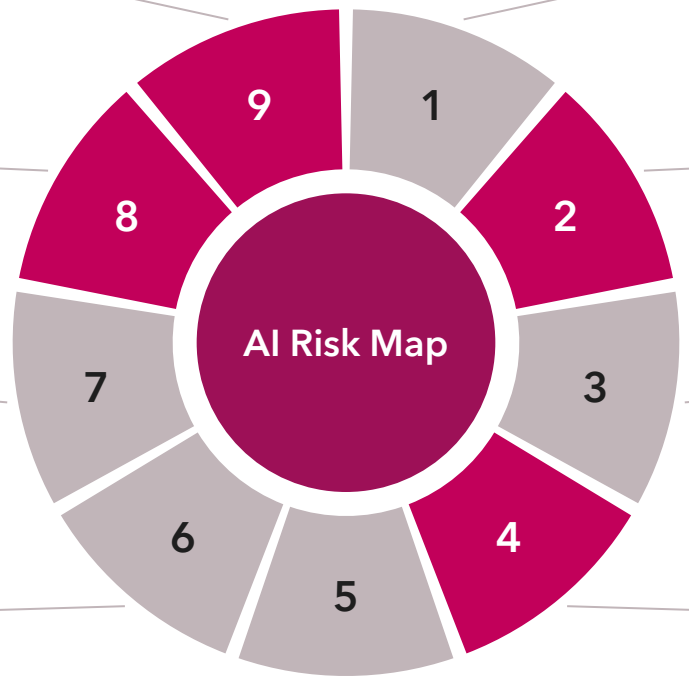

9 Trust and reputation

- **Customer trust**
- Regulator and stakeholder confidence
- Employee trust

8 Security and misuse

- **Prompt injection and manipulation**
- Data leakage
- Malicious use

7 Data and third-party dependency

- Data quality and provenance
- Vendor and model concentration
- Data residency and sovereignty

6 Workforce and capability

- Deskilling and judgement loss
- Role displacement
- Adoption failure

1 Strategic intent and objectives

- Objective misspecification
- Value case erosion
- Strategic dependency

2 Decision integrity

- Output error and hallucination
- Bias and unfair outcomes
- **Automation bias**

3 Oversight and control

- Ineffective human-in-the-loop
- Drift and monitoring gaps
- Shadow AI

4 Autonomy and delegation

- **Agent permission creep**
- Cascading agent actions
- **Unclear escalation thresholds**

5 Accountability and legal

- Accountability gaps
- Regulatory non-compliance
- Contractual and IP exposure

Domain with at least one high risk category

Domain with medium or low categories only

Source: Strategic Design by Argon & Co, AI Risk Map

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

The role AI plays should depend on how ambiguous the decision is and how costly an error would be

Low-ambiguity, low-cost decisions can go to AI quickly; high-stake, ambiguous decisions keep people in charge

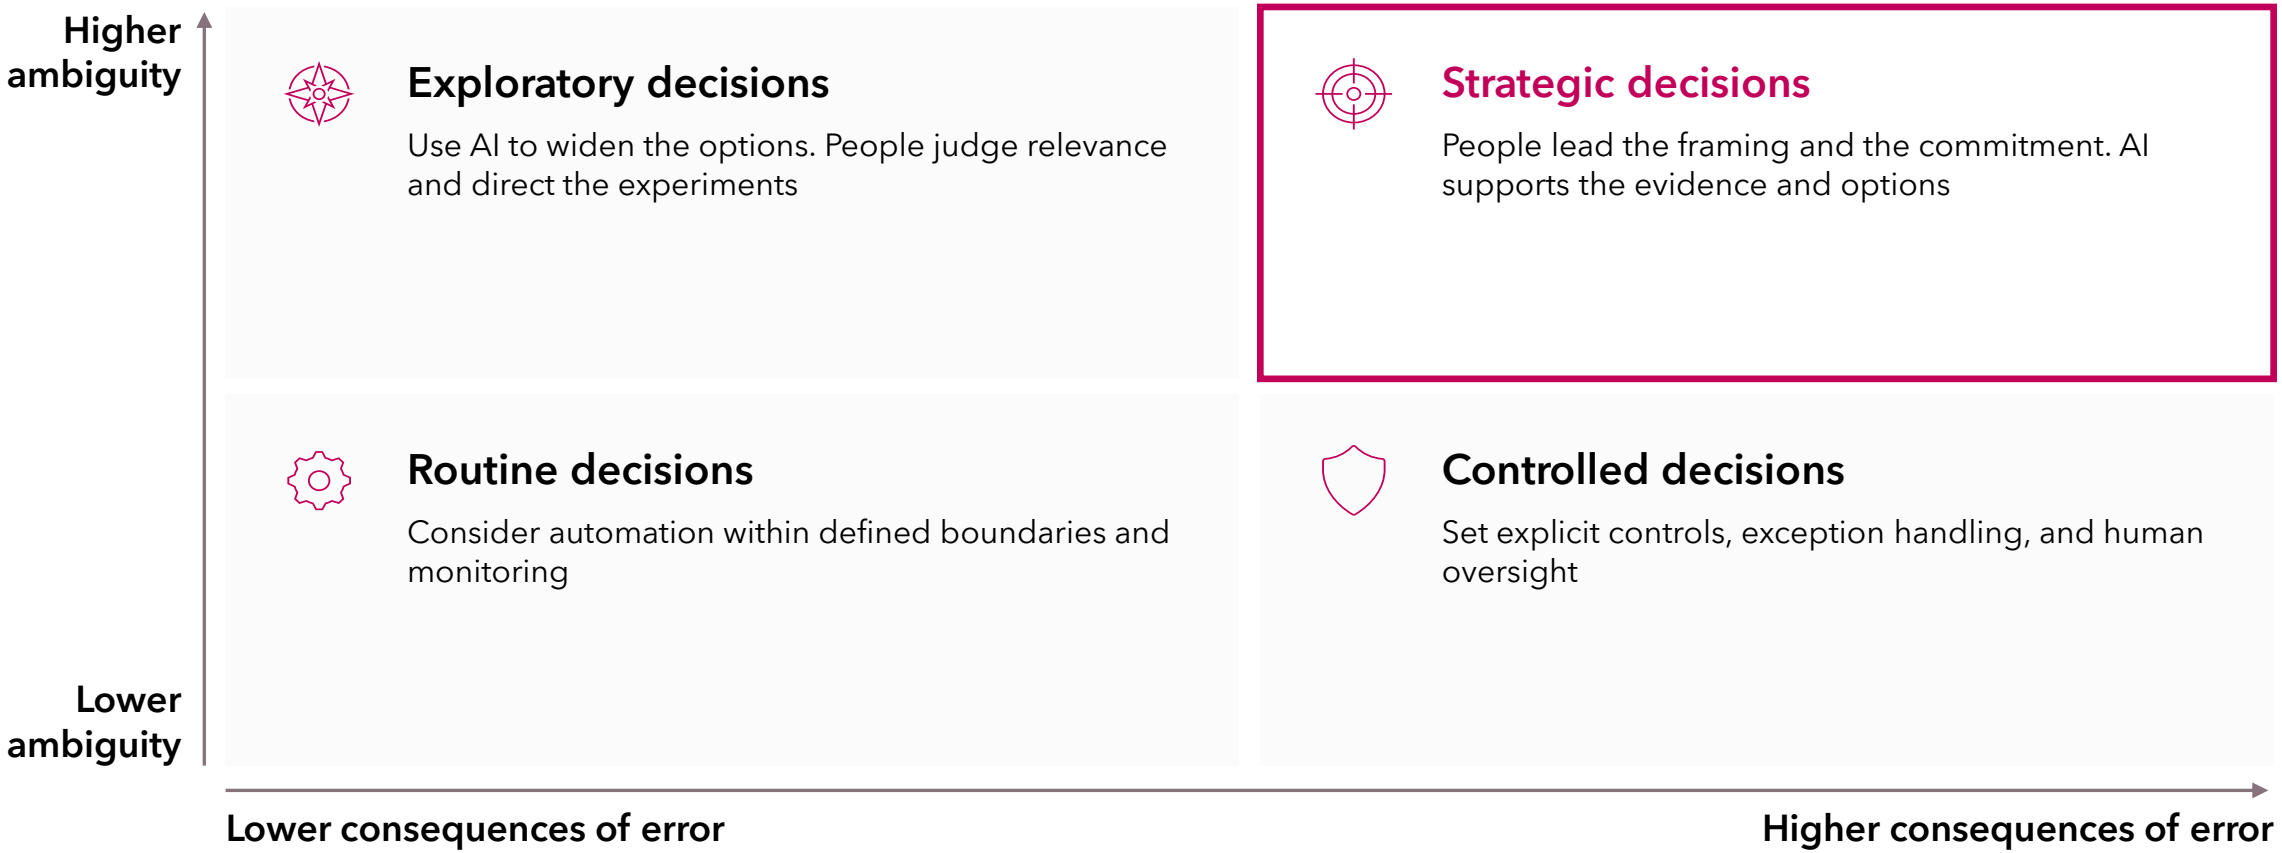

Source: Adapted from Sebastian et al. (2026)

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

People should set the question and the method before AI argues for an option

	01 Test the question	02 Set the method	03 Challenge the options
AI does Widens the options and the evidence	AI tests the request for other explanations, hidden assumptions, and missing information	AI checks the decision specification for unclear criteria, conflicting goals, and missing evidence	AI generates credible alternatives and builds the strongest case against the preferred option
People decide Narrow the choice and own the outcome	People confirm the decision to make now, its scope, and its owner	People approve the specification, set the trade-offs, and keep the decision authority	People confirm that each option is feasible, distinct, and within the mandatory limits
	The result is an agreed decision question	**The result is a clear analysis brief for AI**	**The result is a tested set of options**

Source: Adapted from Sebastian et al. (2026)

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

How easily a decision can be reversed should set how freely an AI agent may make it – most decisions tend to be 'two-way door decisions'

Hard to reverse → **Easy to reverse**

If the decision proves wrong, can it be undone at low cost?

Type 1

One-way door decisions

Consequential and irreversible, or nearly irreversible

Who decides	Senior leaders, after wide consultation
Pace	Methodical, careful, and slow
Method	Deliberate fully and debate before committing

Type 2

Two-way door decisions

Changeable and reversible, which covers most decisions

Who decides	High-judgement individuals, small groups close to the work or AI with frameworks and business rules
Pace	Quick
Method	Decide, learn from the result and correct course

Source: Adapted from Amazon, 2015 Letter to Shareholders (2016)

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

Enterprise and agentic AI controls

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

For enterprise AI, controls must layer as a system's reach widens and its logic adapts

	Static AI system Works within fixed rules, for example a credit-scoring model	Adaptive AI system Learns from live data, for example a fraud-detection model
Narrow reach	**Rules-based controls** Baseline safeguards for every AI system. Model cards, validation tests, bias tests, and sample reviews make each decision rule explicit **Example:** A bank sends each credit model card to an independent model-risk unit for review before go-live	**Post-output alignment controls** Test outcomes against policy, because the decision logic is hard to trace. Audits and user training help people challenge outputs **Example:** Microsoft assesses high-stakes models against libraries of expert-defined policies on what is fair
Broad reach	**Propagation-risk controls** Stop a local error from spreading to connected systems. Map the links, monitor shared infrastructure, and plan oversight in advance **Example:** The European Central Bank stress-tests how a cyber disruption travels across payment systems	**Integrated AI controls** Combine all three layers for systems that learn and act across functions, partners, and markets. Agentic AI sits here **Example:** The Bank of England warns that AI trading agents can amplify external shocks

AI solutions with broader reach change the nature of the risk

- In connected systems, small errors interact with other failures and grow into second-order effects
- A minor logistics API error can combine with a cyber incident at a payment gateway and add to a wider breakdown
- Agentic AI adds speed and reach, often with limited human oversight
- No single organisation can see the whole system. Control needs shared rules and escalation rights with partners and regulators

Controls must match the reach of a decision, and evolve as the decision-maker learns over time

Source: G. Lanzolla, M. Pagani and C.L. Tucci, 'Scaling AI With Adaptive Governance', MIT Sloan Management Review, Summer 2026.

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

Leading organisations already run specific controls for each of the four layers

Control layer	Specific controls	From research	Risk map domains
Rules-based controls Static, narrow reach. The baseline for all	• Model card: data checks, decision thresholds, and bias tests • Independent model-risk review before go-live • Sample tests against expert judgement, before and after launch	A bank ran 100 client files through a new lending model and compared its calls with relationship managers'	2 Decision integrity 5 Accountability and legal
Post-output alignment controls Adaptive, narrow reach	• Test outcomes against expert-defined policies • Audit failure scenarios and the impact on affected groups • Train users to treat outputs as signals, not decisions	Nasdaq analysts record why they close each alert and feed it back to recalibrate the model	2 Decision integrity 3 Oversight and control 9 Trust and reputation
Propagation-risk controls Static, broad reach	• Map dependencies across systems and partners • Monitor shared infrastructure for weak points • Agree escalation rights with partners and suppliers	European Central Bank stress tests simulate how a cyber disruption cascades across payment systems	7 Data and third-party 8 Security and misuse
Integrated controls Adaptive, broad reach. Includes agentic AI	• All three layers above, run together • A named owner for risks that cross unit and company lines	The Bank of England warns that AI trading agents could amplify shocks beyond human detection	4 Autonomy and delegation 1 Strategic intent

Every system starts with rules-based controls. Reach and adaptiveness decide what is layered on top

Source: G. Lanzolla, M. Pagani and C.L. Tucci, 'Scaling AI With Adaptive Governance', MIT Sloan Management Review, Summer 2026.

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

A model that passes its launch tests can still fail in use, so controls must continue after go-live

Most AI risks emerge at two moments, and each needs a different governance response

Before go-live: the model is in development

Failure modes

- The training data is biased, incomplete or out of date
- The model does not fit what the task requires
- Validation tests are too weak to find errors before launch

Example from the research

At one bank, an early credit-limit model showed that small changes to its inputs could cause unexpected shifts in its decisions

Governance response

Prevent errors before the model goes live. Test the data, decision logic, and bias, and get an independent review

After go-live: the model runs in the real world

Failure modes

- Model quality falls as the input data changes over time (data drift)
- The model gives outputs that look plausible but are false
- Users trust the outputs too much because they cannot detect errors

Example from the research

Nasdaq's surveillance systems create hundreds of alerts per second. Some illicit trades are designed to look compliant, so the model can fail to flag them

Governance response

Detect and correct errors while the system runs. Monitor drift, review outcomes, and send the findings back to the model team

Governance and risk management must cover the entire life cycle

Source: G. Lanzolla, M. Pagani and C.L. Tucci, 'Scaling AI With Adaptive Governance', MIT Sloan Management Review, Summer 2026.

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

An AI agent should gain authority only after it passes each evidence gate

Use-case record
Set before the first gate

Risk tier
Set by customer impact, regulatory impact, data sensitivity and model type

Mandatory controls
The checks each tier must pass before and after launch

Decision rights
The people who approve changes, and on which terms

Less authority → More authority

1 Build and test
The model runs on test data only. No customer sees a result

2 Recommend
The AI recommends on real cases. People make every decision

3 Act within limits
The AI decides inside the tier limits. Exceptions need business and risk sign-off

4 Act more widely
The AI learns in use or acts across teams and partners. Extra controls apply

Gate 1	Gate 2	Gate 3
Model card review	**Expert sample test**	**In-use review**
Data and bias checks pass. An independent team signs off	Experts review sample cases. All gaps are explained first	Drift and incidents stay in limits. System links are mapped

The evidence that lets an agent move faster is the same evidence that shows it is safe

Source: Adapted from G. Lanzolla, M. Pagani and C.L. Tucci, 'Scaling AI With Adaptive Governance', MIT Sloan Management Review, Summer 2026

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

Controls only hold when they are built into how the organisation works and learns

After leaders choose the right controls, they must build them into processes and systems and keep them up to date

1 Embed risk controls in daily operations

Response

Build the risk tier, the mandatory controls, and the decision rights into core systems. Oversight then happens in normal planning, audits, and reviews

In practice

A global bank coded its AI controls into credit approval. Staff cannot skip model validation, and each exception needs sign-off from business and risk team

Inaction risk

Controls sit in a separate compliance layer, and people can go around them

2 Reach clear decisions across different expert teams

Response

Agree shared risk thresholds, evidence standards and metrics. Joint model reviews then let teams test the same assumptions, even when they disagree

In practice

Structured reviews combine incident and near-miss analysis, drift monitoring, and a comparison of intended use with actual outcomes

Inaction risk

Teams stop a model too early from caution, or stall in debate about methods

3 Run governance as a learning system

Response

Give named owners the task of reviewing incidents and near misses. Turn each lesson into updated policies, thresholds, and controls

In practice

Each incident and near miss is logged with what went wrong. An owner turns it into a revised policy, threshold, or control, so the rules change as the AI does, with accountable human oversight

Inaction risk

Controls stay fixed while the AI systems and their risks change

Advantage comes not from adopting AI faster, but from governing it better

Source: G. Lanzolla, M. Pagani and C.L. Tucci, 'Scaling AI With Adaptive Governance', MIT Sloan Management Review, Summer 2026.

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

Summary

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

In summary, there are safe ways to deploy AI

1	**Pressure to move faster**	Boards want AI at speed, but Australians trust it less than almost anyone, and human review of AI output fails quietly
2	**Where speed and risk sit**	AI speeds up preparation, yet decisions still wait on people, and the biggest risks cluster in autonomy, security, and trust
3	**What the law already requires**	People stay accountable for AI-assisted decisions and must be able to explain them, with new disclosure rules coming soon
4	**Decision controls**	Match AI's role to how ambiguous, costly, and reversible a decision is, and give agents fixed rules with clear room for discretion
5	**Enterprise and agentic AI controls**	Add control layers as systems adapt and reach wider, grant authority gate by gate, and build controls into daily work

There's no need to reinvent the wheel, there are already ways to deploy AI safely

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

See real change.

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

Contact details and disclaimer

Strategic Design by Argon & Co (formerly Bevington Group) is a specialist consultancy with six core practices:



Strategy Design (including M&A services)



Operating Model Design & Restructuring



Process Reengineering & Automation



Enterprise Digital – Data – AI



Change Management & Delivery Support



Risk Intelligence

Questions or clarification of the content of this report can be directed to:

Report authors		
Alex Spork	+61 423 966 684	alex.spork@argonandco.com
James Tan	+61 403 458 099	james.tan@argonandco.com

Other enquiries can be directed to		
Roger Perry	+61 3 9663 5522	ceo@bevingtongroup.com
Strategic Design office	+61 3 9663 5522	admin@bevingtongroup.com

All content in this report is subject to copyright and cannot be reproduced without the express permission of Bevington Consulting Pty Ltd.

Disclaimer:

This document has been prepared in good faith based on information provided by client staff and other sources of information available at the date of publication without any independent verification. It is the responsibility of the client management to ensure that their staff provide accurate information, and that management are not aware of any issues that may impact the usefulness of the data for its intended purpose. Readers are responsible for their interpretation and actions based on the content of this publication. Bevington Consulting Pty Ltd will not be liable for any loss, damage, cost or expense incurred or arising by reason of any person or organisation using or relying on information in this report without taking, and acting on, further formal advice from Bevington Consulting Pty Ltd (trading as Strategic Design by Argon & Co).

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

References (1 of 4)

- Allens (2026) 'Automated decision-making transparency: what APP entities need to know about the APP 1 amendments', *Allens Insights*, 10 June. Available at: https://www.allens.com.au/insights-news/insights/2026/06/automated-decision-making-transparency-what-app-entities-need-to-know-about-the-app-1-amendments/.
- Amazon.com, Inc. (2016) *2015 letter to shareholders*. Seattle, WA: Amazon.com, Inc.
- Associated Press (2025) 'Deloitte to partially refund Australian government for report with apparent AI-generated errors', *Associated Press*, 7 October. Available at: https://www.usnews.com/news/business/articles/2025-10-07/deloitte-to-partially-refund-australian-government-for-report-with-apparent-ai-generated-errors.
- Australian Ageing Agenda (2026) 'Ruston questions accuracy of aged care assessment tool', *Australian Ageing Agenda*, February. Available at: https://www.australianageingagenda.com.au/executive/ruston-questions-accuracy-of-aged-care-assessment-tool/.
- Australian Broadcasting Corporation (2026) 'What do the pope and most Australians have in common? Neither trust AI companies', *ABC News*, 3 June. Available at: https://www.abc.net.au/news/2026-06-04/artificial-intelligence-australians-trust-survey-pope-leo-ai/106745098.
- Australian Council of Trade Unions (2026) 'AI: unions put corporate Australia on notice', Media release, 27 February. Available at: https://www.actu.org.au/media-release/ai-unions-put-corporate-australia-on-notice/.
- Australian Institute of Company Directors and Human Technology Institute (2026) *A director's guide to AI governance*. Sydney: AICD.
- Australian Prudential Regulation Authority (2023) *Prudential Standard CPS 230 Operational Risk Management*. Sydney: APRA.
- Australian Prudential Regulation Authority (2026a) *Letter to industry on artificial intelligence (AI)*, 30 April. Available at: https://www.apra.gov.au/news-and-publications/apra-letter-industry-artificial-intelligence-ai.
- Australian Prudential Regulation Authority (2026b) 'APRA calls for a step-change in AI-related risk management and governance', Media release, 30 April.
- Australian Securities and Investments Commission (2024) *Report 798: Beware the gap: governance arrangements in the face of AI innovation*. Sydney: ASIC.
- Brynjolfsson, E., Li, D. and Raymond, L.R. (2025) 'Generative AI at work', *The Quarterly Journal of Economics*, 140(2), pp. 889–942. doi:10.1093/qje/qjae044.
- Charlton, A. (2026) 'Speech to AFR AI Summit', Assistant Minister for Science, Technology and the Digital Economy, 2 June. Available at: https://www.minister.industry.gov.au/charlton/media/speech-afr-ai-summit.
- Cheng, M., Lee, C., Khadpe, P., Yu, S., Han, D. and Jurafsky, D. (2026) 'Sycophantic AI decreases prosocial intentions and promotes dependence', *Science*, 391(6792), eaec8352. doi:10.1126/science.aec8352.

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

References (2 of 4)

- Dell'Acqua, F., McFowland III, E., Mollick, E., Lifshitz, H., Kellogg, K.C., Rajendran, S., Krayer, L., Candelon, F. and Lakhani, K.R. (2026) 'Navigating the jagged technological frontier: field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality', *Organization Science*, 37(2), pp. 403–423. doi:10.1287/orsc.2025.21838.
- Department of Industry, Science and Resources (2025) *Guidance for AI adoption*. Canberra: Commonwealth of Australia.
- Digital Transformation Agency (2025) *Policy for the responsible use of AI in government, version 2.0*. Canberra: Commonwealth of Australia.
- Doshi, A.R. and Hauser, O.P. (2024) 'Generative AI enhances individual creativity but reduces the collective diversity of novel content', *Science Advances*, 10(28), eadn5290. doi:10.1126/sciadv.adn5290.
- eSafety Commissioner (2025) 'Court orders $343,500 penalty for posting deepfakes of Australian women', Media release, 26 September. Available at: https://www.esafety.gov.au/newsroom/media-releases/court-orders-343500-penalty-for-posting-deepfakes-of-australian-women.
- EY (2026) 'Australia ranks equal lowest on global AI sentiment as use rises, EY study finds', *EY Global AI Sentiment Study*, 11 May. Available at: https://www.ey.com/en_au/newsroom/2026/05/australia-ranks-equal-lowest-on-global-ai-sentiment-as-use-rises.
- Fortune (2025) 'Deloitte was caught using AI in $290,000 report to help the Australian government crack down on welfare after a researcher flagged hallucinations', *Fortune*, 7 October. Available at: https://fortune.com/2025/10/07/deloitte-ai-australia-government-report-hallucinations-technology-290000-refund.
- Glenn, R.W. (ed.) (2020) *Trust and leadership: the Australian Army approach to mission command*. Dahlonega, GA: University of North Georgia Press.
- The Guardian Australia (2026) 'Developer withdraws plans for Perth datacentre after fierce community opposition', *The Guardian*, 15 May. https://www.theguardian.com/technology/2026/may/15/developer-withdraws-plans-for-perth-datacentre-after-fierce-community-opposition
- Herdağdelen, A., Joren, H. and Tamkin, A. (2026) 'Who's aligned to whom? How AI agents shape experts' research judgment calls', *Anthropic,* 27 October. Available at: https://www.anthropic.com/research/agents-judgment-calls.

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

References (3 of 4)

- Lanzolla, G., Pagani, M. and Tucci, C.L. (2026) 'Scaling AI with adaptive governance', *MIT Sloan Management Review*, 67(4), Summer. Available at: https://sloanreview.mit.edu/article/scaling-ai-with-adaptive-governance/.
- Linklaters (2026) 'Data Protected: Australia', *Linklaters Insights*. Available at: https://www.linklaters.com/en/insights/data-protected/data-protected---australia (contributed by Allens, updated May 2026).
- Maddocks (2026) 'Ten weeks to go: OAIC releases guidance on new automated decision-making transparency obligations under APP 1', *Maddocks Insights,* 7 October. Available at: https://www.maddocks.com.au/insights/oaic-releases-guidance-on-new-automated-decision-making-transparency-obligations.
- Media, Entertainment & Arts Alliance (2025) 'Big Tech agreement to 'Stop Creative Theft' a win for media and creative workers', Media release, 22 August. Available at: https://www.meaa.org/mediaroom/big-tech-agreement-to-stop-creative-theft-a-win-for-media-and-creative-workers/.
- Microsoft and Tech Council of Australia (2023) *Australia's generative AI opportunity*. Sydney: Tech Council of Australia, July.
- Minderoo Foundation (2026) 'Australians support AI innovation, but safety must come first', February. Available at: https://www.minderoo.org/stories/ai-research-feb-2026/.
- Noy, S. and Zhang, W. (2023) 'Experimental evidence on the productivity effects of generative artificial intelligence', *Science*, 381(6654), pp. 187–192. doi:10.1126/science.adh2586.
- Office of the Australian Information Commissioner (2025) *Commissioner initiated investigation into Kmart Australia Limited (Privacy)* [2025] AICmr 155, 26 August. Sydney: OAIC.
- Office of the Australian Information Commissioner (2026a) *Australian Community Attitudes to Privacy Survey 2026*. Sydney: OAIC, May. Available at: https://www.oaic.gov.au/engage-with-us/research-and-training-resources/research/australian-community-attitudes-to-privacy-survey/australian-community-attitudes-to-privacy-survey2026/overview.
- Office of the Australian Information Commissioner (2026b) *Australian Privacy Principles guidelines: Chapter 1, APP 1 open and transparent management of personal information*, version 2.0, 30 September. Sydney: OAIC.
- Royal Commission into the Robodebt Scheme (2023) *Report*. Canberra: Commonwealth of Australia, 7 July.

© Copyright Bevington Consulting Pty Ltd. All rights reserved.

References (4 of 4)

- Sebastian, I.M., Weill, P., Haskamp, T. and vom Brocke, J. (2026) 'Designing decision rights for AI', *MIT CISR Research Briefing*, XXVI(6), 18 June. Available at: https://cisr.mit.edu/publication/2026_0601_AIDecisionMatrix_SebastianWeillHaskampVomBrocke.
- Tech Policy Design Institute (2026) *Earning trust: unlocking AI adoption for Australians*. Canberra: TPDi, May. Available at: https://techpolicy.au/earning-trust/.
- University of Melbourne and KPMG (2025) *Trust, attitudes and use of artificial intelligence: a global study 2025*. Melbourne: University of Melbourne. doi:10.26188/28822919.
- Vaccaro, M., Almaatouq, A. and Malone, T. (2024) 'When combinations of humans and AI are useful: a systematic review and meta-analysis', *Nature Human Behaviour*, 8, pp. 2293–2303. Doi:10.1038/s41562-024-02024-1.

Legislation and case law

- *Bunnings Group Limited and Privacy Commissioner (Guidance and Appeals Panel)*, [2026] ARTA 130 (Administrative Review Tribunal, 4 February 2026).
- *Corporations Act 2001* (Cth).
- *Financial Accountability Regime Act 2023* (Cth).
- *Moffatt v Air Canada*, 2024 BCCRT 149 (Civil Resolution Tribunal of British Columbia).
- *Privacy Act 1988* (Cth), Schedule 1 (Australian Privacy Principles), as amended by the *Privacy and Other Legislation Amendment Act 2024* (Cth).
- Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (*Artificial Intelligence Act*).

© Copyright Bevington Consulting Pty Ltd. All rights reserved.