




**Human Technology
Institute**



Submission to the Productivity Commission: Harnessing data and digital technology interim report

15 September 2025

About the Human Technology Institute

The Human Technology Institute (HTI) is building a future that applies human values to new technology. HTI embodies the strategic vision of the University of Technology Sydney (UTS) to be a leading public university of technology, recognised for its global impact specifically in the responsible development, use and regulation of technology. HTI is an authoritative voice in Australia and internationally on human-centred technology. HTI works with communities and organisations to develop skills, tools and policy that ensure new and emerging technologies are safe, fair and inclusive and do not replicate and entrench existing inequalities.

The work of HTI is informed by a multi-disciplinary approach with expertise in data science, law and governance, policy and human rights.

For more information, contact us at hti@uts.edu.au

Acknowledgement of Country

UTS acknowledges the Gadigal people of the Eora Nation, the Boorooberongal people of the Dharug Nation, the Bidiagal people and the Gamaygal people upon whose ancestral lands our university stands. We would also like to pay respect to the Elders both past and present, acknowledging them as the traditional custodians of knowledge for these lands.

Authors: Professor Nicholas Davis, Sarah Sacher and Professor Edward Santow

Contents

About the Human Technology Institute	2
Executive summary	4
Recommendations	6
Australia's response to AI must balance potential productivity effects with broader impacts of adoption	9
Factoring in costs of AI risks into productivity projections	9
Addressing AI-related harms should be central to Australia's regulatory and policy response	10
Further interrogation of productivity projections in light of uncertainties and contingencies	14
Realistic assessment of the distribution of productivity gains across the population	16
An alternative productivity framing: public policy should promote equitable distribution of gains in Australia's economy and society	16
The Productivity Commission's Blueprint for AI Regulation	17
Interim Report recommendations	17
A 'third way' for AI regulation in Australia	19
Proactive and proportionate regulation for AI supports productivity goals	21
Australia would not be served by being merely a 'regulation taker'	26
Privacy Act reform proposals	28

Executive summary

The Human Technology Institute (HTI) welcomes the opportunity to make a submission to the Productivity Commission's regarding its Interim Report, *Harnessing data and digital technology* (Interim Report). In this submission, HTI addresses the Productivity Commission's broader productivity framing; its recommended blueprint for artificial intelligence (AI) regulation; and its proposals for privacy law reform.

Balancing productivity effects with broader impacts of AI adoption

The economic and broader opportunities associated with AI, as well as the risks of harm, are well understood. Good public policy requires a high degree of rigour in assessing both these opportunities and risks.

The Interim Report forcefully sets out the upside economic potential of increased AI adoption in Australia, including where this potential is speculative or is contingent on a range of unproven assumptions. By contrast, the Interim Report appears to adopt a more sceptical attitude towards the risk of harm, even where some of these harms have already materialised and are being experienced by people in Australia.

In preparing its final report, we urge that the Productivity Commission adopt a balanced approach that gives appropriate weight to opportunities and risks. Good public policy is predicated on a realistic assessment of this combination of reward and risk, and where risks engage fundamental human rights that those be addressed robustly. In the context of AI adoption, any gains in productivity must be assessed alongside likely externalities, especially risks of harm to the community and to the market. These include the economic costs associated with AI risks.

The Commission also should take into account how projected productivity gains associated with increased AI adoption are likely to be distributed across the Australian economy and society. Without intervention, there is a high risk that productivity gains will be enjoyed disproportionately by a small number of large companies, many of which order their operations and tax affairs in jurisdictions outside of Australia, whereas the risks will be borne disproportionately by smaller Australian enterprises, individuals and governments. In formulating law and policy in this area, the Australian Government has an important role in promoting an equitable distribution of benefits and risks. HTI urges the Commission to consider the policy and regulatory settings needed to promote an equitable distribution of the productivity dividends associated with increased AI adoption in Australia.

An Australian approach to regulation for AI

The Productivity Commission has stated that its preferred approach to regulation for AI is one that is proportionate, risk-and outcomes-based. HTI supports this approach. It is also an effective way to apply the vision set out by the Treasurer, the Hon Jim Chalmers MP, who has argued for Australia to adopt a 'middle course' on AI, stating that 'we cannot let AI rip, nor can we pretend it's not happening.'¹

Adopting this middle course means creating a path that is neither a deliberately permissive approach, nor a highly-prescriptive, technology-specific legislative approach. Instead, a good middle course for Australia would apply the following elements:

¹ Jim Chalmers, 'Australia and the AI revolution – turning algorithms into opportunities' *The Guardian* (Online) 3 August 2025 <<https://www.theguardian.com/business/2025/aug/03/australia-shouldnt-fear-the-ai-revolution-new-skills-can-create-more-and-better-jobs>>.

1. A reform strategy based on an expedited analysis of gaps in existing law, which have the effect of leaving certain AI-related risks unaddressed. How this analysis is undertaken is important: given that stakeholders and authoritative bodies differ on what such a gap analysis would entail. It would be a mistake to see this task as requiring a comprehensive analysis of the entire corpus of Australian law, with a view to finding deficiencies. That task is too big, and in any event unnecessary to generate the requisite policy insights. Instead, HTI urges that a simpler, more expedited gap analysis would focus on the areas of greatest risk associated with AI, and how current law addresses those risks. Where there is a deficiency, this would suggest a need for reform of law, policy or some combination of both.
2. Where law reform is warranted, the Government's default position should be that targeted, technology-neutral legislation is preferable. Only where this is unlikely to achieve the Government's objective is technology-specific reform warranted.
3. The Government should prioritise implementing reform recommendations that have already been carefully tested in areas relevant to data, digital technology and AI. The Government's recent reviews, especially the Attorney-General's Department review of the *Privacy Act 1988* (Cth) and the Robodebt Royal Commission's recommended framework for automated decision-making by government, are important examples of such reform recommendations.
4. The Government should adopt flexible, economy-wide legislation to address high-risk AI, in the form of framework legislation. Framework legislation is necessary to achieve clarity and consistency on a range of AI-related concepts throughout the Australian economy – an outcome that cannot be achieved by relying solely on updates to existing laws. Framework legislation would provide a default set of legal definitions and rules, which could be departed from where warranted in other legislation. It would set out a nuanced process to assess AI risk and take an outcomes-based approach, so that it remains applicable as AI technology develops.
5. It is a truism to observe that good law and policy are only as effective as their application and enforcement, and so the Government should prioritise support for the broader regulatory ecosystem. This involves ensuring that regulators have the right mix of capability, resources, functions and powers to fulfil their oversight functions in the context of AI.

The more detailed recommendations, summarised at the end of the executive summary, seek to give life to this approach.

AI-driven productivity improvement hinges on several key preconditions. These include regulatory certainty, public trust in AI products and services, and responsible AI adoption practices, including meaningful engagement with workers. By failing in one or more of these areas, Australia risks the worst of all worlds: simultaneously failing to meet the productivity potential of AI while experiencing costly AI harms that lead to poorly-designed, reactive regulation.

Lastly, Australia should adopt a regulatory approach that fits its own legal system, politics, and values, while still aligning as much as possible with international standards. It would be a mistake for Australia to simply follow others' rules, especially given shifting geopolitical dynamics. With major jurisdictions like the US and EU taking different paths, Australians cannot assume they will benefit from regulatory protections imposed in other jurisdictions.

Privacy Act reform proposals

The Productivity Commission was right to identify the problems with consent as a mechanism to protect Australians' personal data in the 21st century. However, the Interim Report's proposed solution – namely, an 'alternative compliance' pathway under the Privacy Act – is speculative and untested. Moreover, it gives short shrift to the deep work that had been done to develop a different solution through many years of research, consultation and consensus building by government, industry and civil society.

The Attorney-General Department's 2022 *Privacy Act Review Report* was the product of precisely this sort of deep work. The Government has committed to implement the majority of the reforms, which are actively in train. The Report's recommendations also align with community expectations that favour strengthened privacy protections for individuals – noting that the Privacy Act exists to protect the human right to privacy, in line with Australia's international law obligations.

That this reform process has been slow, and remains unfinished, and that its recommendations have not achieved *universal* agreement among stakeholders, are not reasons to abandon the Attorney-General's Department's central recommendations – especially in circumstances where the alternative recommendation is highly unlikely to achieve this level of agreement. That conclusion is fortified by the fact that the Interim Report's proposal is likely to reduce the protection to the human right to privacy. There is little, if any, evidence to support the view that privacy laws such as the General Data Protection Regulation (GDPR) 'stifle' innovation. Additionally, the Interim Report's analysis does not take into account the financial and human costs associated with *not* adequately protecting privacy, nor the many harms that the GDPR has likely *prevented*.

HTI does not address other law reform proposals discussed in the Interim Report – including those related to copyright law, and data sharing reforms. However, it notes a general concern regarding proposals that would undermine existing protections, without full and balanced consideration of potential impacts on human and legal rights, including privacy.

Recommendations

Recommendation 1

The Productivity Commission's final recommendations should be made by reference to a thorough articulation and considered assessment of AI-related risks of harms. Risks to be considered include:

- implications for individual human rights and legal rights (for example, intellectual property rights, employment rights)
- implications for individuals' physical and mental safety
- adverse impacts to groups of individuals or collective rights of cultural groups
- costs to the broader Australian economy
- impacts on Australian workforces
- risk to Australia's liberal-democratic system
- risks to the environment
- upside risks or likelihood to bring relevant public benefit.

Recommendation 2

To assist the Productivity Commission to implement Recommendation 1, and to inform law and policy reform related to AI, the Australian Government should commission independent modelling of the nature and extent of harms that are occurring through the adoption and deployment of AI systems. This could be achieved through tasking regulators, individual departments or an external expert body to collate evidence of current and emerging harms across various dimensions.

Recommendation 3

Productivity projections should be rigorously tested by reference to a larger body of research, and in light of the likely distribution of productivity gains to individuals associated with AI adoption. The Productivity Commission should actively consider the policy and regulatory settings needed to ensure that productivity dividends are shared equitably in the Australian economy and society, including to individuals.

Recommendation 4

The Productivity Commission should consider an approach to regulation for AI in Australia that would involve:

- coordinated reform to priority areas of existing law – guided by an expedited, high-level gap analysis, and conducted in parallel to the mandatory guardrails process
- the development of flexible and proportionate framework legislation as the chosen model for mandatory guardrails for high-risk AI. Framework legislation would provide a default set of legal definitions and rules, and could be departed from in other legislation
- dedicated support for the broader regulatory ecosystem, to ensure that regulators have the right mix of capabilities, resources and powers to effectively fulfil their oversight functions in the context of AI.

Recommendation 5

Priority reforms, which have been the subject of extensive multi-sector consultation, and are the subject of existing government reform commitments, should be urgently implemented. These include tranche 2 reforms to the *Privacy Act 1988* (Cth), and the development of a framework for use of automated decision making by government.

Recommendation 6

The Productivity Commission should conduct a thorough cost-benefit analysis of three options for AI regulation (a deliberately permissive approach; a prescriptive approach; and HTI's proposed 'third way'). This assessment should include a focus on:

- costs associated with AI risks
- contingencies that influence the success of productivity measures – including for example, regulatory certainty, consumer trust, incentives or disincentives for good AI governance, support for workers to transition, likely rates of successful AI adoption across particular sectors or types of roles, infrastructure and investment
- distribution of productivity benefits to individuals (including workers and disadvantaged groups) across areas such as income, quality of life and leisure time, as distinct from productivity benefits for business and the national GDP.

Recommendation 7

In light of the Australian Government's acceptance of existing recommendations to reform the Privacy Act, and the reform process currently in train, the Productivity Commission should reconsider its proposals to amend the Act.

Australia's response to AI must balance potential productivity effects with broader impacts of adoption

The issues at the heart of the Productivity Commission inquiry are multifaceted. Any gains in productivity from adopting AI must be assessed alongside potential externalities, especially risks of harm to the community and to the market. Prioritising adoption without adequate safeguards will tend to magnify these risks. Any potential increase in productivity should be weighed against the impact of these externalities. Broadly speaking, the Productivity Commission should advocate an approach that pursues increased productivity in ways that properly account for the risks posed by such externalities.

Factoring in costs of AI risks into productivity projections

The Interim Report acknowledges potential harms attributable to AI systems but does not adequately quantify or balance them against the projected benefits. The Interim Report estimates the cumulative GDP impact of potential labour productivity growth linked to the adoption of AI systems – a total additional contribution of \$116bn over the next ten years. But the Commission does not provide any comparable estimates of the economic costs associated with AI-related risks.

Nevertheless, it is possible to quantify many forms of AI-related harms in economic terms. One source of harms referred to by the Productivity Commission is when AI is used maliciously, including to disseminate fake content, manipulate public opinion, and for use in criminal offences. A quantified example in this category is the costs of scams for Australians, which reached \$2.7 billion in 2023, and were over \$2 billion in 2024.² In 2025, the use of AI has been attributed as driving a more than 25% increase in scam losses in the first four months of 2025 compared to the previous period.³ Hence, there is clearly significant *additional downside risk* associated with increased use of AI, in areas such as scam costs to Australians. Based on these figures, that could be in the order of \$500 million per year.

Another meaningful source of costs to the economy relates to harms that result when automated systems fail to operate effectively. While Robodebt was not an AI system (in the sense that debt notices were created by reference to a relatively crude algorithm), it illustrates the economic and social costs that can arise when automated decision-making tools are deployed without adequate governance. This scheme was intended as an efficiency measure projected to save the government \$1.7 billion over five years. Instead it resulted in \$2.4 billion in payouts to those affected, inclusive of the largest class-action settlement in Australian history.⁴

The number of AI systems used to make or influence decisions that directly affect Australians is rising rapidly, but systematic, fit-for-purpose governance to prevent the use of AI leading to other, Robodebt-like failures is lagging. A 2024 research report

² 'Scam Losses Decline, but More Work to do as Australians Lose \$2.7 Billion', *Australian Competition & Consumer Commission* (Web Page, 28 April 2024) <<https://www.accc.gov.au/media-release/scam-losses-decline-but-more-work-to-do-as-australians-lose-27-billion>>; 'Australians Better Protected as Reported Scam Losses Fell by Almost 26 Per Cent', *National Anti-Scam Centre* (Web Page, 11 March 2025) <<https://www.nasc.gov.au/news/australians-better-protected-as-reported-scam-losses-fell-by-almost-26-per-cent>>.

³ Shannon Williams, 'AI-fuelled Scams Surge as Australian Losses Jump 28 Percent', *IT Brief Australia* (Web Page, 23 August 2025) <<https://itbrief.com.au/story/ai-fuelled-scams-surge-as-australian-losses-jump-28-percent>>.

⁴ Dan Jervis-Bardy, 'Robodebt Victims Win Record \$548.5m Settlement from Government, Taking Total Payout to \$2.4bn', *The Guardian* (Web Page, 4 September 2025) <<https://www.theguardian.com/australia-news/2025/sep/04/robodebt-victims-win-record-settlement-centrelink-government-compensation>>.

commissioned by the NSW Ombudsman revealed the use (and proposed use) of AI in automated decision making across a wide range of contexts, including across every NSW state government portfolio.⁵ A 2025 report by the Australian National Audit Office found that a single agency – the Australian Tax Office – had 43 AI models in production, yet 74% of these did not have a completed data ethics assessment, and the ATO as a whole lacked clearly-defined assurance and approval arrangements for testing.⁶

Examples from around the world demonstrate that when algorithmic decision-making systems fail, the costs are significant. False accusations of childcare fraud generated by a Dutch automated system that relied on machine learning affected 35,000 families, resulted in the resignation of the Rutte government in 2021, and is estimated to cost at least \$1.8 billion, with some estimates topping \$25 billion.⁷ In July 2025, the UK Government revealed that it had paid more than \$2.2 billion to 7,900 postmasters across four compensation schemes as part of the Post Office Horizon IT scandal.⁸ Failures in the Canadian Government's automated payroll software resulted in a 2020 compensation agreement to pay public sector workers \$384 million,⁹ with total public funds spend on the system estimated to be more than \$4 billion.¹⁰ Of course, direct compensation costs and other charges related to fixing failing AI or algorithmic systems are likely to represent just a fraction of the real cost to humans.

The additional risks – and therefore negative productivity impact – affecting Australians, which is attributable to AI, can be estimated by considering three factors: the adoption rate of AI systems in critical decision-making systems across government and industry; the average rate of failure, misuse or over-use of AI systems that results in economic and other harms to Australians; and the potential quantum of this harm.

Addressing AI-related harms should be central to Australia's regulatory and policy response

The importance of protecting against AI harms should be self-evident. Any policy or regulatory response to AI should centre Australian people and communities, as underlined by the Australian Government.¹¹ There is a substantial evidence base pointing to a range of AI harms. These human costs are significant even where not directly measurable. Certain harms, and human rights impacts, must be considered

⁵ Kimberlee Weatherall et al, 'Automated Decision-Making in NSW: Mapping and Analysis of the Use of ADM systems by State and Local Governments' (Research Report, ARC Centre of Excellence for Automated Decision-Making and Society, March 2024)

<<https://cmsassets.ombo.nsw.gov.au/assets/Reports/Executive-Report-ADMS.pdf>>.

⁶ Australian National Audit Office, *Governance of Artificial Intelligence at the Australian Taxation Office* (Report No 26 2024-25) 43-6 <https://www.anao.gov.au/sites/default/files/2025-02/Auditor-General_Report_2024-25_26.pdf>.

⁷ 'Childcare Scandal Compensation May Cost €14 billion: NOS', *Dutch News* (Web Page, 14 May 2025) <<https://www.dutchnews.nl/2024/05/childcare-scandal-compensation-may-cost-e14-billion-nos/>>.

⁸ 'What are the Different Post Office Compensation Schemes?', *British Broadcasting Corporation* (Web Page, 8 July 2025) <<https://www.bbc.com/news/articles/cev4mw43w13o>>.

⁹ Andrew Duffy, 'PSAC Strikes Phoenix Compensation Deal with Federal Government', *Ottawa Citizen* (Web Page, 10 July 2020) <<https://ottawacitizen.com/news/local-news/psac-strikes-phoenix-compensation-deal-with-federal-government>>.

¹⁰ 'Compensation from Phoenix Payroll Disaster Rises to \$711M', *Western Standard News* (Web Page, 11 February 2025) <https://www.westernstandard.news/news/compensation-from-phoenix-payroll-disaster-rises-to-711m/62042?utm_source=website&utm_medium=related-stories>.

¹¹ See, eg, Department of Industry, Science and Resources (Cth), *Safe and Responsible AI in Australia Consultation: Australian Government's Interim Response* (Report, January 2024) 19 <https://storage.googleapis.com/converlens-au-industry/industry/p/prj2452c8e24d7a400c72429/public_assets/safe-and-responsible-ai-in-australia-governments-interim-response.pdf>.

and addressed through a law or policy response to AI, and cannot simply be 'offset' by projected productivity gains.

Some examples of AI-related harms that have directly affected Australians

- Online Safety and children's rights:
 - eSafety Commissioner data from June 2025 shows that explicit deepfake images of underage Australians have doubled since 2023.¹²
 - A journalistic investigation uncovered instances of young people in Australia being encouraged to take their own life by AI chatbots.¹³
- Consumer rights:
 - 83% of Australians have experienced negative consequences from digital design features designed to influence their behaviour.¹⁴
 - Airbnb adopted an AI system that assessed customer "trustworthiness" based on publicly available data. This led to Australian customers being banned from the platform without explanation, with sex workers appearing to have been targeted for exclusion.¹⁵
- Privacy and worker rights:
 - 76% of Australians have experienced harm because of a data breach.¹⁶
 - An Australian company turned laptops into covert recording devices to monitor employees working from home.¹⁷
- Democracy and national security: A largely automated pro-Russian news website was used to spread propaganda on Australian topics in the lead-up to the 2025 federal election.¹⁸

There is significant variability in how risks associated with AI are developing over time. Some of these risks have already manifested in such a way that some, or many, people have suffered significant harm. Other risks are still emerging in the sense that, without effective mitigations, some forms of increased AI adoption are likely to lead to increased harm. It would be short sighted to dismiss those harms that have been documented internationally – particularly since they are often connected with the same global companies that Australian people and businesses rely on, in economies that are further along in the AI adoption process than Australia.

¹² Farid Farid, 'Real Punishment for "Terrifying Tech" Deepfake Makers', *The Canberra Times* (Web Page, 7 August 2025) <<https://www.canberratimes.com.au/story/9035465/real-punishment-for-terrifying-tech-deepfake-makers/>>.

¹³ April McLennan, 'AI Chatbots Accused of Encouraging Teen Suicide as Experts Sound Alarm', *Australian Broadcasting Corporation* (Web Page, 12 August 2025) <<https://www.abc.net.au/news/2025-08-12/how-young-australians-being-impacted-by-ai/105630108>>.

¹⁴ Consumer Policy Research Centre, *Duped by Design: Manipulative Online Design* (Report, June 2022) 6 <<https://cprc.org.au/wp-content/uploads/2022/06/CPRC-Duped-by-Design-Final-Report-June-2022.pdf>>.

¹⁵ Jarni Blakkarly, 'Is Airbnb Using an Algorithm to Ban Users from the Platform?', *CHOICE* (Web Page, 21 March 2022) <<https://www.choice.com.au/consumers-and-data/data-collection-and-use/how-your-data-is-used/articles/airbnb-banning-users>>.

¹⁶ 'Australian Community Attitudes to Privacy Survey 2023 Infographic', *Office of the Australian Information Commissioner* (Web Page) <<https://www.oaic.gov.au/engage-with-us/research-and-training-resources/research/australian-community-attitudes-to-privacy-survey/australian-community-attitudes-to-privacy-survey-2023-infographic>>.

¹⁷ David Marin-Guzman, 'Company Turned Laptops into Covert Recording Devices to Monitor WFH', *The Australian Financial Review* (Web Page, 24 August 2025) <<https://www.afr.com/work-and-careers/workplace/company-turned-laptops-into-covert-recording-devices-to-monitor-wfh-20250822-p5mp0z>>.

¹⁸ Ange Lavoipierre and Michael Workman, 'Pro-Russian Influence Operation Targeting Australia in Lead-up to Election with Attempt to "poison" AI Chatbots', *Australian Broadcasting Corporation* (Web Page, 3 May 2025) <<https://www.abc.net.au/news/2025-05-03/pro-russian-push-to-poison-ai-chatbots-in-australia/105239644>>.

While we agree in principle with the Interim Report's assertion that the regulatory response should focus on mitigating the additional risk of harms occurring due to AI, the adoption of AI systems can both *amplify* the impact of relatively, mature, existing risks (for example, cybersecurity vulnerabilities and attacks), and *trigger* existing risks in new ways and contexts (for example, widespread discrimination in hiring practices). In some cases, AI systems may also create entirely new risks (for example, widespread deskilling thanks to over-reliance on LLMs).

We caution against treating the desire to focus on additional risks of harm as a reason to overlook long-standing risks from data and digital systems that have already been identified as problematic elsewhere. Effective regulation must recognise that AI can and will exacerbate harms such as bias, manipulation, or concentration of market power, both by amplifying their effect and triggering them in novel ways.

Accordingly, both the adaptation of existing, technology-neutral rules and the targeted introduction of new, technology-specific measures may be required to address the scale and speed of AI-related risk amplification. Furthermore, given the pace of AI development and deployment, there is a strong case for surfacing and taking swift action on AI-specific risks, rather than waiting for problems to crystallise before regulatory responses are implemented. If AI technology continues to develop and be adopted without a concurrent investment in fit-for-purpose governance driven by clear regulatory incentives, the potential for AI-related harms will increase. While the examples above focused on how AI systems can harm individuals directly, AI systems also pose systemic risks. For example, the OECD and the Financial Stability Board have warned that reliance on AI systems in financial settings, without effective technical and governance protections, pose systemic risks to the stability of financial systems.¹⁹ A decision to remain passive in light of these emergent and increasing harms would simply be poor policy.

The harms associated with AI are not dispersed or experienced equally by all. AI misuses can, and does, bring disproportionate harms to certain groups, including children, people with disability, LGBTQIA+ people and Aboriginal and Torres Strait Islander people. Additionally, widespread AI adoption has implications for Australian workforces. The Interim Report observes that 'there could be a sizeable number of workers that are displaced', and points to a need for retraining and the availability of 'the social safety net of unemployment assistance'.²⁰ More work is needed here. A deeper analysis might involve assessing how AI adoption is likely to affect workers across various roles and parts of the economy; identifying specific policy interventions to assist workers to upskill and successfully transition; and determining whether law reform is needed to shore up employment legislation, including to protect workers rights, and encourage engagement with workers during AI adoption processes.

Placing greater weight to the claimed (and often speculative) productivity benefits, without giving due weight to possible or likely risks, also negates important values-based discussion and debates. Regarding AI's impact on copyright and the creative industries, Arts Minister the Hon Tony Burke MP has stated that it is 'not useful to

¹⁹ Organisation for Economic Cooperation and Development, *Regulatory Approaches to Artificial Intelligence in Finance* (Paper No 24, September 2024) 18 <https://www.oecd.org/en/publications/regulatory-approaches-to-artificial-intelligence-in-finance_f1498c02-en.html>; Financial Stability Board, *The Financial Stability Implications of Artificial Intelligence* (Report, 14 November 2024) 12-6 <<https://www.fsb.org/2024/11/the-financial-stability-implications-of-artificial-intelligence/>>.

²⁰ Commissioner Stephen King, Commissioner Julie Abramson, Productivity Commission, *Harnessing Data and Digital Technology: Interim Report* (Report, August 2025) 13-4 <<https://www.pc.gov.au/inquiries/current/data-digital/interim/data-digital-interim.pdf>>.

frame conversations about the arts around economic activity'.²¹ The OECD has similarly observed that there is 'an existential question' to be considered, regarding whether AI development should be permitted 'at the cost of whole sections of the economy, and at the expense of enormous groups of workers and their valuable intellectual property'.²²

We must be careful to avoid a trade-off between projected benefits to the Australian economy, and known harms to the public. Instead, broader impacts and risks of AI adoption should be considered as part of a robust cost-benefit analysis, to help rigorously assess AI policy proposals and ensure that the net benefits of any proposed option outweigh the likely harms to Australians, and that harms are addressed appropriately.

HTI also recommends that the Australian Government undertake a parallel process to form a more accurate picture of the nature and extent of harms that are occurring through the operation of AI. This could be achieved through tasking regulators, individual departments or an external expert body to collate evidence of current and emerging harms. This should in turn feed into law reform processes across various sectors, as well as any economy-wide law.

Recommendation 1

The Productivity Commission's final recommendations should be made by reference to a thorough articulation and considered assessment of AI-related risks of harms. Risks to be considered include:

- implications for individual human rights and legal rights (for example, intellectual property rights, employment rights)
- implications for individuals' physical and mental safety
- adverse impacts to groups of individuals or collective rights of cultural groups
- costs to the broader Australian economy
- impacts on Australian workforces
- risk to Australia's liberal-democratic system
- risks to the environment
- upside risks or likelihood to bring relevant public benefit.

Recommendation 2

To assist the Productivity Commission to implement Recommendation 1, and to inform law and policy reform related to AI, the Australian Government should commission independent modelling of the nature and extent of harms that are occurring through the adoption and deployment of AI systems. This could be achieved through tasking regulators, individual departments or an external expert body to collate evidence of current and emerging harms across various dimensions.

²¹ Josh Bornstein and Sophie Cunningham, 'Stealing from Artists Won't Make Australia Richer. Here's Why', *The Australian Financial Review* (Web Page, 15 August 2025) <<https://www.afr.com/life-and-luxury/arts-and-culture/stealing-from-artists-won-t-make-australia-richer-here-s-why-20250815-p5mn84>>.

²² Organisation for Economic Cooperation and Development, *The Impact of Artificial Intelligence on Productivity, Distribution and Growth: Key Mechanisms, Initial Evidence and Policy Challenges* (Paper No 15, April 2024) <https://www.oecd.org/en/publications/the-impact-of-artificial-intelligence-on-productivity-distribution-and-growth_8d900037-en.html>.

Further interrogation of productivity projections in light of uncertainties and contingencies

There is a body of independent academic research that provides a nuanced view of the productivity potential of AI. HTI draws on a literature review produced by Per Capita and the Centre of the Public Square to highlight an alternative evidence base,²³ alongside other sources.

First are studies that point to high rates of failure for AI-related projects. Research by Douglas Gray and Evan Shellshear found that failure rates of 80% are common for data science projects.²⁴ Similarly, a recent MIT study found that 95% of organisations adopting GenAI found zero return on investment despite enterprise investment of \$30 billion to \$40 billion into GenAI.²⁵ Productivity gains are realised only when AI systems move successfully from experimentation into sustained, value-creating use. If a high percentage of projects fail to reach that point, adoption and productivity assumptions must be revised down. In other words, the drag created by high failure rates acts as a hidden externality in productivity calculations: it depresses realised gains relative to potential. Without adequate mechanisms of learning and governance, such costs may compound across the economy as firms replicate costly mistakes.

Second, there are a series of reports indicating that the productivity gains associated with AI adoption are significantly uncertain and ‘murky’.²⁶ They are also highly context-dependent. Whether AI adoption translates into productivity gains depends on policy and regulatory decisions that determine the course of AI’s development; the business context in which AI is adopted; and the organisational settings within companies, amongst other factors.

- An OECD report entitled *The impact of AI on Productivity, Distribution and Growth* describes factors which make predicting AI’s impact on productivity unclear and multifaceted. The report found initial micro evidence of productivity gains; however, the impact on longer-term and aggregate productivity remains uncertain. Additionally, the hyper-concentration of AI development and AI technology owners within a handful of tech companies may limit the broader diffusion of AI benefits and create further risks.²⁷
- A paper by the US Federal Reserve found that the timeline for AI productivity will be ‘inherently slow’. GenAI’s contribution to productivity growth will depend on the speed with which adoption is attained and, historically, the process for integrating revolutionary technologies into the economy is a protracted one. Significant

²³ ‘Checking for Mistakes, the Hidden Negative Externalities in the AI Productivity Debate’, *Per capita* (Web Page, 19 August 2025) <https://percapita.org.au/our_work/checking-for-mistakes-in-ai/>.

²⁴ Douglas Gray and Evan Shellshear, *Why Data Science Projects Fail: The Harsh Realities of Implementing AI and Analytics, without the Hype* (CRC Press, 1st ed, 2025) <<https://www.routledge.com/Why-Data-Science-Projects-Fail-The-Harsh-Realities-of-Implementing-AI-and-Analytics-without-the-Hype/Gray-Shellshear/p/book/9781032660301?srsId=AfmBOorlbFrWf0hjZ6Ph5yC7OisrIFTC9by7jgUjaylX09P8bhMLIKwb>>.

²⁵ Madison Mills, ‘MIT Study on AI Profits Rattles Tech Investors’, *Axios* (Web Page, 21 August 2025) <<https://www.axios.com/2025/08/21/ai-wall-street-big-tech>>.

²⁶ Jon Whittle, ‘Does AI Actually Boost Productivity? The Evidence is Murky’, *The Conversation* (Web Page, 11 July 2025) <<https://theconversation.com/does-ai-actually-boost-productivity-the-evidence-is-murky-260690>>.

²⁷ Organisation for Economic Cooperation and Development, *The Impact of Artificial Intelligence on Productivity, Distribution and Growth: Key Mechanisms, Initial Evidence and Policy Challenges* (Paper No 15, April 2024) <https://www.oecd.org/en/publications/the-impact-of-artificial-intelligence-on-productivity-distribution-and-growth_8d900037-en.html>; ‘Checking for Mistakes, the Hidden Negative Externalities in the AI Productivity Debate’, *Per capita* (Web Page, 19 August 2025) <https://percapita.org.au/our_work/checking-for-mistakes-in-ai/>.

investment and infrastructure are needed for success, but investment to deploy new technologies is ‘fraught with risk’ and building to meet anticipated demand can lead to ‘disastrous consequences’ for the economy.²⁸

- A report by the Social Policy Group found Australia is ill-suited to capitalise on the wealth generation promised by AI when compared to complex economies. It found that in an economy like Australia’s – dominated by raw material exports, limited value-added production, and consumption-driven services – AI-driven productivity gains are capped, translating primarily into layoffs rather than growth.²⁹

Third, there are several studies and surveys that have queried the extent to which AI can reliably boost individual worker productivity in real-world circumstances:

- Upwork surveyed 2,500 workers across the US, UK, Canada, Australia. It found that the majority of high expectations around AI boosting productivity come from C-Suite leaders and are detached from the reality of worker experience, with 77% of employees reporting that AI has increased their workloads.³⁰
- Stack Overflow conducts an annual survey which polls over 65,000 professional developers. While adoption of AI tools for developers were high with around 76% of developers using them, 66% have reported that their productivity has slowed due to additional work needed to validate and correct AI’s outputs.³¹
- A longitudinal economic study from the University of Chicago and the University of Copenhagen looked at data from 25,000 employees across 7,000 workplaces in Denmark over two years. Researchers concluded that AI chatbots have had no significant impact on earnings or recorded hours in any occupation.³²

These points are raised not to undermine the strong case for AI adoption as an enhancer of productivity, but to ensure that the analysis is conducted with sufficient relevant data. The projections in the Interim Report must be rigorously tested. On the whole, the Productivity Commission appears to accept the upside potential of AI at face value, notwithstanding some highly speculative and contingent assumptions. At the same time, the Productivity Commission is sceptical of the existence and momentum of AI harms, despite significant and undeniable evidence of harms that are already occurring, let alone those projected by leading experts to pose a serious and likely future risk.

²⁸ ‘Checking for MistAlkes, the Hidden Negative Externalities in the AI Productivity Debate’, *Percapita* (Web Page, 19 August 2025) <https://percapita.org.au/our_work/checking-for-mistakes-in-ai/>; Martin Neil Baily et al, ‘Generative AI at the Crossroads: Light Bulb, Dynamo or Microscope?’ (Working Paper No 53/2025, Board of Governors of the Federal Reserve System, 2025) 49 <<https://www.federalreserve.gov/econres/feds/files/2025053pap.pdf>>.

²⁹ The Social Policy Group and Carla Wilshire, *Artificial Intelligence and the Great Entrenchment* (Report, December 2024) 10 <https://socialpolicy.org.au/wp-content/uploads/2024/12/SPG_AI_and_the_Great_Retrenchment.pdf>;

³⁰ ‘Checking for MistAlkes, the Hidden Negative Externalities in the AI Productivity Debate’, *Percapita* (Web Page, 19 August 2025) <https://percapita.org.au/our_work/checking-for-mistakes-in-ai/>; Kelly Monahan and Gabby Burlacu, ‘From Burnout to Balance: AI-Enhanced Work Models’, *Upwork* (Web Page, 23 July 2024) <<https://www.upwork.com/research/ai-enhanced-work-models>>.

³¹ ‘Checking for MistAlkes, the Hidden Negative Externalities in the AI Productivity Debate’, *Percapita* (Web Page, 19 August 2025) <https://percapita.org.au/our_work/checking-for-mistakes-in-ai/>; ‘2025 Developer Survey’, *Stack Overflow* (Web Page, 2025) AI section <<https://survey.stackoverflow.co/2025/ai>>.

³² ‘Checking for MistAlkes, the Hidden Negative Externalities in the AI Productivity Debate’, *Percapita* (Web Page, 19 August 2025) <https://percapita.org.au/our_work/checking-for-mistakes-in-ai/>; Anders Humlum and Emilie Vestergaard, ‘Large Language Models, Small Labor Market Effects’, (Working Paper No 33777, Massachusetts Institute of Technology, May 2025) <<https://www.nber.org/papers/w33777>>.

Realistic assessment of the distribution of productivity gains across the population

The Interim Report has stated that the productivity benefit will translate to \$4,300 per capita in 10 years' time due to AI adoption.

HTI has been unable to identify in history even one productivity gain, derived through the development and use of a new technology, which has benefited every person equally, and it is well understood that AI harms disproportionately impact vulnerable and marginalised groups. Additionally, while productivity gains to the national GDP usually correlate with higher wages, this isn't automatic, nor is it guaranteed to be an equal distribution across the population. A study published in *Nature* on the nuanced relationship between technological innovation and income inequality found that technological innovation inadvertently exacerbates income disparities, with a pronounced effect in developed economies.³³

Inevitably, many Australian workers will be negatively impacted by AI adoption, especially if it proceeds with minimal regulation or policy intervention. It is stating the obvious to note that a person who loses their job due to the implementation of AI systems in their company, does not stand to gain as much as a CEO or shareholder that may see millions in profit through that same process. Certain cohorts and communities are also projected to face significant inequities through AI transition. The Interim Report observes, for example, that administrative roles are more likely to be displaced by AI – and this in itself is likely to disproportionately impact women in the workforce.³⁴

Indeed, Jobs and Skills Australia's recent report on the implications of GenAI for the workforce found that women, older workers, First Nations peoples, and people with disability may face disproportionate risks due to occupational concentration and digital access gaps.³⁵ These impacts of the AI transition must be mitigated through proactive regulation and policy intervention – a necessary step to prevent entrenchment of inequality.

An alternative productivity framing: public policy should promote equitable distribution of gains in Australia's economy and society

The Interim Report highlights how Australians may benefit from an AI-driven productivity boom. For example, automation can increase 'leisure time and living standards'.³⁶

It is important to assess these claims against the real-world economic and other implications of minimally-regulated AI adoption for individual Australians, as distinct from industry and the national GDP.

³³ Anran Xiao et al, 'Bridging the Digital Divide: The Impact of Technological Innovation on Income Inequality and Human Interactions' (2024) 11(809) *Humanities and Social Sciences Communications* 1 <<https://www.nature.com/articles/s41599-024-03307-8>>.

³⁴ Kweilin Ellingrud et al, 'Generative AI and the Future of Work in America', *McKinsey* (Web Page, 26 July 2023) <<https://www.mckinsey.com/mgi/our-research/generative-ai-and-the-future-of-work-in-america>>; Commissioner Stephen King, Commissioner Julie Abramson, Productivity Commission, *Harnessing Data and Digital Technology: Interim Report* (Report, August 2025) 12-3 <<https://www.pc.gov.au/inquiries/current/data-digital/interim/data-digital-interim.pdf>>.

³⁵ Job and Skills Australia, *Our Gen AI Transition: Implications for Work and Skills* (Report, 14 August 2025) <<https://www.jobsandskills.gov.au/publications/generative-ai-capacity-study-report>>.

³⁶ Commissioner Stephen King, Commissioner Julie Abramson, Productivity Commission, *Harnessing Data and Digital Technology: Interim Report* (Report, August 2025) 11 <<https://www.pc.gov.au/inquiries/current/data-digital/interim/data-digital-interim.pdf>>.

HTI's view is that, without targeted intervention, AI productivity gains may concentrate among large enterprises and metropolitan centres, while small businesses, regional communities, and lower-skilled workers face displacement. In other words, even if the Interim Report's estimates of a \$116bn increase in economic activity are correct, structural barriers and market dynamics related to data, digital and AI products and services suggest these benefits may flow disproportionately to those already advantaged.

For example, it is relevant to consider questions such as the following. Will increased AI-driven productivity necessarily improve quality of life for Australian workers in the absence of proactive regulation? The notion that AI adoption will lead to more leisure time for everyday people is a trope well-explored in utopian science fiction novels, but does not translate to current trends of intrusive AI-driven productivity tools pressing workers to work faster and harder for no additional reward. Australian Amazon warehouse workers, who may be penalised if they fail to meet high-pressure performance metrics while being constantly monitored,³⁷ may be creating productivity gains for a major corporation—but at great cost to the individual.

Will AI adoption by major companies necessarily lead to better quality products and services that Australian consumers rely on in the absence of proactive regulation? Perhaps not in circumstances where consumers are required to dedicate significant personal time towards navigating a series of AI chat-bots and automated phone lines to deal with complex consumer issues, or where they have no choice but to engage with opaque and inaccurate dynamic pricing systems when using a rideshare app or booking a holiday.³⁸

It should be a policy objective for the productivity benefits of AI adoption to flow to the individual. An alternative productivity framing could focus on exploring how to ensure that people have a share of the productivity dividends of AI in their day to day lives, and develop AI policy and regulatory settings needed to manifest this outcome.

Recommendation 3

Productivity projections should be rigorously tested by reference to a larger body of research, and in light of the likely distribution of productivity gains to individuals associated with AI adoption. The Productivity Commission should actively consider the policy and regulatory settings needed to ensure that productivity dividends are shared equitably in the Australian economy and society, including to individuals.

The Productivity Commission's Blueprint for AI Regulation

Interim Report recommendations

The Interim Report argues that 'poorly designed regulation could stifle AI investment without improving outcomes'. It recommends a regulatory approach that is

³⁷ Select Committee on the Impact of Technological and Other Change on the Future of Work and Workers in New South Wales, Parliament of New South Wales, *Final Report: Workplace Surveillance and Automation* (Report No 2, November 2022) 15-6

<<https://www.parliament.nsw.gov.au/lcdocs/inquiries/2591/Report%20No.%20%20-%20Future%20of%20work%20and%20workers%20in%20New%20South%20Wales.pdf>>; Daniel Ziffler, 'With a "CCTV Camera in Every Laptop" as Workplace Surveillance Extends to the Home, Governments are Trying to Determine a Limit', *Australian Broadcasting Corporation* (Web Page, 11 November 2024) <<https://www.abc.net.au/news/2024-11-11/a-cctv-camera-in-every-laptop-as-workplace-surveillance-extends-/104504744>>.

³⁸ *Trivago N.V. v Australian Competition and Consumer Commission* [2020] FCAFC 185; *Australian Competition and Consumer Commission v Trivago N.V. (No 2)* [2022] FCA 417.

‘proportionate, risk-based, outcomes-based and technology-neutral where possible’.³⁹ HTI agrees with this statement, as an articulation of policy objectives.

As the Productivity Commission finalises its recommendations, it should prioritise specific reform that adopts the Commission's preferred proportionate, risk-and outcomes-based approach to regulation in this area.

It would be absurd to suggest that the Australian Government must choose one of two extreme approaches to policy and regulation for AI: either an approach that focuses only on risk, or one that ignores risk. Instead, the Productivity Commission rightly recommends a balanced approach that takes advantage of productivity and other economic opportunities associated with AI, while addressing the risks. To make that balanced approach real, the Commission should be as clear on the practical steps to address risks as it is on steps needed to take advantage of the opportunities. More specifically, the Productivity Commission should give due weight to the several, well-developed reform proposals that have been the subject of extensive consultation, such as the Attorney-General Department's proposed reforms to federal privacy law.

In simple terms, there are three broad categories of reform for AI that Australia can choose from:

- A deliberately permissive approach to regulation for AI that primarily relies on existing laws, takes a ‘wait and see’ approach to harms, and in some cases, undoes existing protections to enable faster AI adoption. This approach is most commonly associated with the United States federal government approach under President Trump, although it should be acknowledged that individual US states do not uniformly adopt that approach.
- Detailed, prescriptive legislation that applies to all or a significant portion of the development, use and deployment of AI. The EU's AI Act is an example of this approach.
- A flexible approach that favours technology neutrality, takes a balanced approach to AI opportunities and harms, and can incorporate both coordinated reforms to existing laws, and outcomes-based framework legislation. This is HTI's preferred approach.

Whether or not this was the Productivity Commission's intent, the Interim Report has been widely interpreted as supporting a deliberately permissive approach to regulation for AI. Implementation of these recommendations would delay proactive steps towards addressing AI harms through a consistent, cross-cutting process, by pausing the mandatory guardrails, in order to pursue measures that assist businesses at the forefront of AI development and diffusion. While the Interim Report supports updating technology-neutral laws to ensure they are fit for purpose to address AI harms (an important step), it also makes sweeping proposals regarding privacy and copyright law, which would erode key protections for individuals and content creators.

The Interim Report recommendations are underpinned by a number of assumptions:

- minimal regulation, or at least a slower regulatory response, will lead to higher productivity;
- mandatory guardrails would necessarily be technology specific and prescriptive (i.e., not ‘outcomes based’);

³⁹ Commissioner Stephen King, Commissioner Julie Abramson, Productivity Commission, *Harnessing Data and Digital Technology: Interim Report* (Report, August 2025) 9, 19
<<https://www.pc.gov.au/inquiries/current/data-digital/interim/data-digital-interim.pdf>>.

- ‘untested’ mandatory guardrails would ‘raise uncertainty’ as compared to incremental reform via gap analyses;
- a ‘comprehensive’ gap analysis is needed prior to overarching mandatory guardrails if those reforms are to be effective.

In this section, HTI unpacks these assumptions and sets out its view of a ‘third way’.

A ‘third way’ for AI regulation in Australia

Good regulation for AI can facilitate innovation and guide it in a positive direction. A balanced approach to AI regulation would give proportionate attention both to AI opportunities and harms.

In its submission to the Department of Industry, Science and Resource’s (DISR) mandatory guardrails consultation, HTI set out an approach to AI regulation that would coordinate law reform behind a clearly defined regulatory objective: to protect people from harm, and to support innovation for economic benefit and societal wellbeing.⁴⁰

HTI’s proposed approach would involve taking proactive and coordinated measures to address AI harms – including via targeted reforms to technology neutral legislation; and the adoption of flexible, economy-wide legislation to address high-risk AI. The latter could take the form of a ‘framework legislation’ model, as contemplated through the mandatory guardrails consultation process.

HTI’s model would adopt a nuanced approach to assessing AI risk. This involves the articulation of a definition of risk, and a principles-based process for balancing AI risks and opportunities to enable consistent assessments across a range of use contexts.

Assessing risk

HTI supports a human-rights based approach to risk that would support the development of a risk matrix, and would take into account a range of risks and mitigating factors, including:

- upside risks (i.e. positive opportunities) and downside risks (ie, threats)
- a broad range of risks that government and the private sector typically consider would be relevant—including economic, commercial, social, political, environmental and safety risks
- the context in which the relevant risk arises. For example, AI used in high stakes decision-making contexts, such as law enforcement, would involve a higher level of risk than, say, AI in a computer game
- any risk mitigants, including human oversight and governance safeguards, as well as the existence of effective regulation in respect of a particular area or activity
- a proportionality-based mechanism for balancing risks, rights and interests that come into conflict with each other.

In addition to coordinated law reform, this option would prioritise dedicated support for the broader regulatory ecosystem. This involves ensuring that regulators have the right mix of capabilities, resources and powers to effectively fulfil their oversight functions in

⁴⁰ Human Technology Institute, Submission to Department of Industry, Science and Resources, *Introducing Mandatory Guardrails for AI in High-risk Settings: Proposals Paper* (4 October 2024) 34 <<https://utsd8.prod.acquia-sites.com/sites/default/files/2024-10/HTI-submission-DISR-mandatory-guardrails-proposals-paper.pdf>>.

the context of AI. Support for regulators is needed to ensure the fair and consistent enforcement of existing laws in the AI context; and to reduce uncertainty for business through regulator-issued guidance, standards and stress-testing.

Flexible framework legislation

As recognised by the Productivity Commission, technology-neutral law already applies to AI, including high-risk AI. There is, however, a need to achieve cross-sectoral clarity and consistency on a range of AI-related concepts.

Relying *solely* on reforms to existing laws will not be sufficient to address shared, intersecting harms associated with high-risk AI, and it is unlikely that a sector-specific approach to reform would achieve a coherent and certain regulatory response. There is also a risk of inconsistencies across legal regimes, and the potential for unnecessary complexity, fragmentation and confusing overlap. This in turn would increase compliance burdens for business.

The option of framework legislation would provide a default set of legal definitions and rules, which could be departed from in other legislation. In this way, framework legislation would operate alongside domain-specific legal definitions and rules, rather than duplicating them. Unlike an EU-style AI Act, framework legislation would not attempt to comprehensively regulate AI, anticipate every high-risk use case, or adopt rigid, technology-specific provisions. Instead it would favour a flexible, outcomes-based approach, designed to remain applicable as technology develops.

Framework legislation would provide clarity by setting out:

- a clear regulatory objective
- a normative foundation for understanding risk; a clear definition of risk; and a process for assessing risk
- requirements for AI developers and deployers to take reasonable steps to mitigate risks (i.e. to comply with process-based mandatory guardrails)
- a rebuttable presumption that where a person is responsible for making a decision using AI, that person is legally liable for the impact of that decision
- enforcement through appropriate mechanisms, including oversight by relevant regulators.

Reforms to existing technology-neutral laws

HTI supports the Interim Report's conclusion that regulation for AI should include a process for updating existing laws.

There are a number of possible options for a regulatory gap analysis. A 'comprehensive' gap analysis by reference to the entire corpus of current law is a Herculean task, and wholly unnecessary for achieving the objective of effective AI reform.

The alternative is a much simpler gap analysis process that would assess the suitability of existing law against some of the key harms associated with AI. This could be conducted relatively quickly, especially since it is already largely clear where priority reforms are needed. In HTI's view, key priorities for technology-neutral law reform are as follows:

- Implementation of outstanding law reform recommendations that relate to AI, such as the Privacy Act Review and the framework for automated decision-making by

government, as recommended by the Robodebt Royal Commission. These can be actioned immediately.

- Clarification and strengthening of laws in priority areas – including competition and consumer law, discrimination law, online safety, copyright and broader intellectual property law, and employment law. Reviews or reform initiatives in some of these areas are already being progressed by Government. This will involve more detailed consideration of regulatory gaps with respect to the identified priority areas, taking into account existing enforcement mechanisms and standards.

Importantly, this process could occur alongside, and in conversation with, the development of mandatory guardrails reform.

On the other hand, delaying AI reform to conduct comprehensive gap analysis would simply prolong existing regulatory uncertainty and increase exposure to AI-related harms. A passive approach could also result in AI reforms effectively being set by external events as government is forced to react to emergencies, most of which are exacerbated by the absence of clear legal guardrails. We have seen this most acutely in the context of cybersecurity,⁴¹ online safety,⁴² and social media crises.⁴³ Legislating in haste, in direct response to a crisis, has been shown time and again to result in sub-optimal legislation.

Proactive and proportionate regulation for AI supports productivity goals

The goals of minimising AI risks through regulation, and realising economic benefit, depend on interconnected and symbiotic factors. Good regulation and regulatory enforcement can prevent harms, while facilitating productivity and fostering positive innovation as part of a balanced AI ecosystem.

Productivity through AI adoption relies on several preconditions – these include regulatory certainty, public trust, and responsible AI adoption practices. All of these hinge on proactive and proportionate AI regulation.

Regulatory certainty as a driver of productivity

Without clarity about what is and isn't permissible (or what will be decided in the future), companies can be inhibited from investing in AI. Failure to address this uncertainty creates a barrier to innovation in AI development and is a drag on AI's potential productivity growth.

HTI and the e61 Institute conducted forthcoming research on the relationship between AI, productivity and regulatory certainty.⁴⁴ This research found that without credible, durable policy signals focused on responsible and productive AI use, firms have weaker incentives to invest in the data, infrastructure, applications, complementary capabilities and organisational change that productivity requires. Regulatory ambiguity

⁴¹ Jim Chalmers and Michelle Rowland, 'Changes to Protect Consumers Following Optus Data Breach', *The Treasury* (Joint Media Release, 6 October 2022) <<https://ministers.treasury.gov.au/ministers/jim-chalmers-2022/media-releases/changes-protect-consumers-following-optus-data-breach>>.

⁴² Tom Lowrey, 'Federal Government Moves to Criminalise 'Doxxing' after Publication of Jewish Australians' WhatsApp Messages', Australian Broadcasting Corporation (Web Page, 13 February 2024) <<https://www.abc.net.au/news/2024-02-13/federal-government-to-criminalisedoxxing/103458052>>; 'Public Consultation on Doxxing and Privacy Reforms', *Attorney-General's Department* (Web Page, 6 September 2024) <<https://consultations.ag.gov.au/integrity/doxxing-and-privacy-reforms/>>.

⁴³ Paul Karp, 'Australia Passes Social Media Law Penalising Platforms for Violent Content', *The Guardian* (Web Page, 4 April 2019) <<https://www.theguardian.com/media/2019/apr/04/australia-passes-social-media-law-penalising-platforms-for-violent-content>>.

⁴⁴ Forthcoming e61 and HTI research.

delays adoption, may distort investment, and risks driving unsustainable or low-value reorganisation. This suggests a clear productivity case for the Australian government to reduce regulatory uncertainty.

Key findings on the connection between regulatory uncertainty and productivity include:

- AI requires large, complementary and partly irreversible investments with long payback periods. Even small rule changes can materially change expected returns and design choices.
- Without a clear national regulatory approach, firms face legal ambiguity on data use, liability, intellectual property, employment law and professional obligations.
- Investment is sensitive to both the quality of rules and their predictability over time, particularly in regulated sectors where compliance confidence is essential.
- Uncertainty raises perceived risk on long-horizon, cross-functional AI investments. It delays and distorts capital spending and encourages short-termism.
- Uncertainty increases the risk of coordination failure: firms may wait for others to act, mitigate risks inefficiently, or design systems to align with foreign regulatory regimes. The result is less experimentation, slower diffusion, and missed productivity gains.
- Unlocking AI's productivity dividend requires not just innovation in technology, but coherence and credibility in the institutional frameworks that govern its use.

Not all regulatory approaches support productivity equally. e61 and HTI assessed three possible approaches to AI regulation, and ultimately found that a pragmatic approach that favours technology neutrality (akin to HTI's proposed 'third way') is most conducive to productivity.

Productivity dividends associated with three possible regulatory approaches	
Deliberately permissive approach	<p>A deliberately permissive approach could take one of two main forms. It could involve consciously adopting laissez faire approach, declining to create any new regulation or update existing law to address AI-related risks. Alternatively, it could take the more radical stance of exempting certain activities associated with the development and use of AI from complying with otherwise-applicable law.</p> <p>Either form could minimise the short-term regulatory burden for firms and maximise the initial uptake of AI systems. If credible and consistent, this could promote productivity-enhancing reorganisation across the economy. But this approach would face significant challenges in maintaining regulatory certainty in practice.</p> <p>One such challenge from this approach is state governments stepping into what they perceive as a regulatory vacuum, creating a patchwork of sub-national regulation. A recent example of this is in AI-enabled workplace surveillance, where a number of states, including NSW and Victoria, are in the process of reforming their own laws.</p> <p>Another risk is that this approach fails to guard against a range of consequential emerging harms, resulting in long-term instability in government policy. For example, without clear liability regimes,</p>

	<p>developers and users are unlikely to mitigate second-order societal impacts and even some direct harms flowing from AI, increasing both the risk of harm and the likelihood of costly regulatory reversals.</p> <p>Over time, policy reversals or reactive interventions to emerging, AI-related risks and harms could undermine both business and consumer confidence and weaken AI adoption and investment. It could thus fail to build the institutional foundations to foster significant long-term productivity gains.</p>
Prescriptive, AI-specific approach	<p>A prescriptive, AI-specific regime aims to provide clarity for AI developers, deployers and users, particularly in high-risk contexts. By creating risk-based rules and setting compliance standards specific to AI systems, it can deliver short-term certainty about obligations, supporting adoption where both use cases are related risks are clear. Because productivity gains depend on predictable rules that let firms invest and deploy at scale, this promise of certainty is the main channel to higher productivity.</p> <p>However, this approach risks being both inflexible and increasingly burdensome in fast-changing contexts. Detailed rules can struggle to remain durable as technology evolves. As a result, the effort of interpreting rules and implementing compliance such as human oversight may overwhelm both firms and regulators, potentially creating a disproportionate burden on smaller firms. When rules date quickly or diverge, planning becomes less predictable, undermining the very certainty needed for investment and diffusion so while such a framework may enhance trust in narrow settings, it is unlikely to sustain long-term regulatory certainty.</p> <p>A further concern is that this approach would create a new class of technology-specific regulation, applying across existing regulatory frameworks in product and labour markets. Being overly prescriptive about how technologies should be designed and implemented, regardless of the actual risk could become a drag on innovation and could limit the productivity growth from new business models and reallocation of resources across the economy.</p>
Pragmatic reform that favours technology neutrality	<p>A key strength of this approach is that it embeds a degree of discipline around the identification of harms. It encourages policy makers and regulators to rely first on current regimes that seek to address risks that arise regardless of the technology in question, and depart from these only where there is a case for doing so, such as where those laws cannot be effectively applied or where a genuinely new issue arises that merits a novel regulatory response.</p> <p>Compared with a permissive approach that proves unstable in practice, or a prescriptive approach that risks rigidity, a pragmatic, technology-neutral model offers greater regulatory certainty. It avoids speculative rule-making aimed at anticipating every possible high-risk use case, and it does not signal a doctrine of non-interference. Instead, by orienting firms toward compliance with outcome-focused laws already in force, it establishes a stable baseline of expectations: businesses know which obligations apply, regulators can act</p>

	<p>consistently, and both parties can adapt incrementally as new risks emerge.</p> <p>However, its success depends on consistent interpretation, proactive guidance, and adequate regulatory resourcing. Without these, firms would still face ambiguity in areas like privacy, intellectual property law, algorithmic accountability, and the application of employment and consumer law in novel contexts.⁴⁵</p>
--	---

HTI and e61's report concluded that while Australia's current legal environment is not lawless, its lack of direction and credible commitments has created real uncertainty for firms and workers. Reducing regulatory ambiguity can be seen as a strategic lever for long-term growth: at both the firm and economy levels, governance should not be positioned as a constraint on AI, but the foundation for investment in productive transformation.⁴⁶

Building public trust and realising community expectations

The Interim Report underscores the importance of community trust several times, recognising that 'consistent and reliable regulation can help promote trust in AI technology'.⁴⁷

The Australian public has markedly low trust in AI. KPMG and University of Melbourne research showed that only 36% of Australians are willing to trust AI, with 78% concerned about negative outcomes.⁴⁸ This lack of trust is not misplaced. Australians are increasingly exposed to deepfakes, scams, mis- and dis-information, surveillance in retail environments, and automated processes to determine important matters, such as social security payments. The community expects that governments will act, through regulation and policy, to ensure stronger protections.

Importantly, Australians recognise the value of a balanced approach to AI regulation. Polling research, commissioned by the Minderoo Foundation, found that 61% of Australians favour a balanced but firm approach to AI regulation – one that protects people while fostering innovation. But when a balanced approach was removed as an option in the question-set, support for strict regulation jumped to 64% – even if that meant sacrificing innovation and productivity. Additionally, 81% said AI must benefit all Australians, not just corporations.⁴⁹

It is clear from these results that Australians recognise the productivity and innovation potential of AI, but do not want this to come at the expense of their rights and safety. Failing to build trust with the public through proactive and sensible regulation that responds to the public's concerns could sabotage AI's potential, by chilling AI adoption.

Building trust is particularly important for public sector uses of AI. There is enormous opportunity for AI to foster productivity across government, and in turn, for government to demonstrate responsible AI practice and drive broader trust in AI, in line with its role

⁴⁵ Forthcoming e61 and HTI research.

⁴⁶ Forthcoming e61 and HTI research.

⁴⁷ Commissioner Stephen King, Commissioner Julie Abramson, Productivity Commission, *Harnessing Data and Digital Technology: Interim Report* (Report, August 2025) 17 <<https://www.pc.gov.au/inquiries/current/data-digital/interim/data-digital-interim.pdf>>.

⁴⁸ The University of Melbourne and KPMG International, *Trust, Attitudes and Use of Artificial Intelligence: A Global Study 2025* (Report, 2025) <<https://assets.kpmg.com/content/dam/kpmgsites/xx/pdf/2025/05/trust-attitudes-and-use-of-ai-global-report.pdf>>.

⁴⁹ Talbot Mills Research, *Minderoo Foundation: AI Combined Report* (Report, August 2025) 17 <<https://cdn.minderoo.org/assets/documents/AI-attitudes-report-August-2025.pdf>>.

as an 'exemplar'.⁵⁰ On the other hand, poor AI adoption can lead to loss of confidence over the long-term, as the ongoing fallout from the Robodebt scheme illustrates. The Australian public rightly holds government to a higher standard, and regulation is needed to ensure that this standard is met. Progressing law reform such as the development of a framework for automated decision making by government, would go a long way to improving trust, minimising harms, and ensuring considered and successful AI adoption by government.

Incentivising responsible AI adoption as a driver of productivity

It is well understood that strong AI governance and early consideration of risk, are key to successful AI adoption. Sensible AI regulation can encourage this kind of good practice.

To illustrate: HTI's and Essential Research's *Invisible Bystanders* and *From Invisible to Involved* reports found that worker engagement in AI development improves the uptake of AI to promote productivity.⁵¹ Through close consultation with workers, this research found that workers have a nuanced understanding of the benefits of AI. When they are excluded from technology decisions, risks of failure and mistrust grow. But when they're included, AI solutions are more likely to be adopted and succeed, being better aligned with real-world needs. Here, consideration of risks to workers and workers' rights during AI adoption is both a protective measure, and an enhancer of productivity.

Research on benefits of worker engagement in the context of AI adoption

Organisations that engage their workers on AI experience:

- 5.9 times more financial benefit.⁵²
- 2.1 times greater return on investment.⁵³
- 4.7 times more AI fluency among workers.⁵⁴
- 4.6 times higher top worker performance.⁵⁵

⁵⁰ Commissioner Stephen King, Commissioner Julie Abramson, Productivity Commission, *Harnessing Data and Digital Technology: Interim Report* (Report, August 2025) 10 <<https://www.pc.gov.au/inquiries/current/data-digital/interim/data-digital-interim.pdf>>.

⁵¹ Human Technology Institute, *From Invisible to Involved: A Guide to Worker Engagement on AI* (Report, June 2025) <<https://www.uts.edu.au/globalassets/sites/default/files/2025-06/25.06.04-hti-guide-to-worker-engagement-on-ai.pdf>>; Essential Research and Human Technology Institute, *'Invisible Bystanders': How Australian Workers Experience the Uptake of AI and Automation* (Report, May 2024) <https://www.uts.edu.au/globalassets/sharedmedia/documents/CSJI/essentialresearchuts_invisible_bystanders_0524.pdf>; Llewellyn Spink and Nicholas Davis, 'Companies are Betting on AI to Help Lift Productivity. Workers Need to be Part of the Process', *The Conversation* (Web Page, 19 June 2025) <<https://theconversation.com/companies-are-betting-on-ai-to-help-lift-productivity-workers-need-to-be-part-of-the-process-258396>>.

⁵² Worker engagement empowers workers to personally derive value from AI. Such organisations are 5.9x as likely to get significant financial benefits from AI compared with organisations where employees do not get value from AI: Sam Ransbotham et al, 'Achieving Individual and Organizational Value with AI', MIT Sloan Management Review (Web Page, 31 October 2022) .

⁵³ High-performing organisations focus on deep and detailed AI implementations, rather than solutions that only touch the surface. With an average of 3.5 use cases, these organisations anticipate generating 2.1x greater ROI on their AI initiatives than counterparts. From Potential to Profit: Closing the AI Impact Gap, BCG AI Radar (Web Page, 15 January 2025).

⁵⁴ In a survey among N=1,500 professionals, Salesforce found that employees who are heard are 4.6x more likely to feel empowered to perform their best work: 'How Engaged Employees Are the Path to Success', Salesforce (Blog Post, 21 August 2019).

⁵⁵ Gallup's quarterly workforce study reveals that when leaders have communicated a plan for AI implementation, workers feel 4.7x more comfortable using AI in their role: 'Strategy Will Fail Without Culture That Supports', Gallup (Web Page, 1 November 2024).

On the other hand, AI adoption can be undermined by poor implementation processes that do not take into account implications for workers or their place within broader organisational context. This is evidenced by a number of backflips made by major companies after making workers redundant in preference for AI – and then having to hire them back.⁵⁶ These incidents cause harm to the affected workers, reputational damage for businesses, and no doubt result in a productivity deficit.

ADM-S Centre research draws similar conclusions in respect of data governance. It found that the benefits of AI are compromised if organisations do not maintain high-quality and well-structured data or adequately manage privacy and cybersecurity risks arising from the transfer of data to multinational tech companies.⁵⁷

These findings are two among many that indicate that incentivising good AI governance and practice can help to unlock the productivity potential of AI.⁵⁸

Meanwhile, a policy response that promotes productivity without sufficient guardrails incentivising good behaviour can embolden careless or unscrupulous actors to innovate in a manner that threatens the wellbeing of Australians, and undercuts government's broader economic aims. As noted by the US Consumer Financial Protection Bureau, positive innovation and competition depend on 'companies competing on the merits of their products or services', rather than by exploiting legal uncertainties to engage in unethical behaviour.⁵⁹

Industry cannot be relied upon to adopt AI responsibly, in the absence of mandatory requirements. This point has been recognised by prominent industry figures. Sam Altman, for example, has stated that companies need regulation to keep them on track,⁶⁰ while Daniel Petre has argued that companies cannot be trusted to build AI without laws 'forcing them to do the right thing'.⁶¹

Australia would not be served by being merely a 'regulation taker'

The Interim Report argues that Australia should be a 'regulation taker', since international regulatory approaches to AI are likely to greatly influence the regulatory

⁵⁶ Commonwealth Bank of Australia, Swedish fintech company, Klarna, and Duolingo: see, eg, 'CBA Backflips on AI Cuts, but the Threat Remains', *Finance Sector Union* (Media Release, 21 August 2025); Chris Morris, 'Going "AI First" Appears to be Backfiring on Klarna and Duolingo', *Fast Company* (Web Page, 12 May 2025) <<https://www.fastcompany.com/91332763/going-ai-first-appears-to-be-backfiring-on-klarna-and-duolingo>>.

⁵⁷ Fan Yang and Jake Goldenfein, 'Does AI Really Boost Productivity at Work? Research Shows Gains Don't Come Cheap or Easy', *ARC Centre of Excellence for Automated Decision-Making and Society* (Web Page, 15 August 2025) <<https://www.admscentre.org.au/does-ai-boost-productivity/>>.

⁵⁸ See, eg, Douglas Gray and Evan Shellshear, *Why Data Science Projects Fail: The Harsh Realities of Implementing AI and Analytics, without the Hype* (CRC Press, 1st ed, 2025) <<https://www.routledge.com/Why-Data-Science-Projects-Fail-The-Harsh-Realities-of-Implementing-AI-and-Analytics-without-the-Hype/GrayShellshear/p/book/9781032660301?srsId=AfmBOorlbFrWf0hjZ6Ph5yC7OisrIFTC9by7jgUjaylX09P8bhMLIKwb>>.

⁵⁹ Seth Frotman and Erie Meyer, 'CFPB Comment on Request for Information on Uses, Opportunities, and Risks of Artificial Intelligence in the Financial Services Sector', Consumer Financial Protection Bureau (Web Page, 12 August 2024) <<https://www.consumerfinance.gov/about-us/newsroom/cfpb-comment-on-request-for-information-on-uses-opportunities-and-risks-of-artificial-intelligence-in-the-financial-services-sector/>>.

⁶⁰ James Clayton, 'Sam Altman: CEO of OpenAI Calls for US to Regulate Artificial Intelligence', *British Broadcasting Corporation* (Web Page, 17 May 2023) <<https://www.bbc.com/news/world-us-canada-65616866>>.

⁶¹ Paul Smith, 'Tech Crowd Can't be Trusted to Regulate Themselves on AI: Daniel Petre', *The Australian Financial Review* (Web Page, 11 August 2025) <<https://www.afr.com/technology/tech-crowd-can-t-be-trusted-to-regulate-themselves-on-ai-daniel-petre-20250811-p5mm3c>>.

landscape in Australia.⁶² While it would be hubristic to suggest that Australia should adopt an entirely unilateral approach to regulation in this area, in light of shifting geopolitical trends, there are strong arguments in favour of Australia adopting a path that aligns with its values.

Since major jurisdictions such as the US and the EU have diverged in their regulatory approaches, there is no guarantee that Australians will feel the residual protections of international regulatory regimes. Global technology companies are already taking advantage of weaker legal protections in certain jurisdictions, including Australia. Previously, the so-called 'Brussels Effect' enabled Australians to benefit from stronger protections as global companies uniformly altered their operations to comply with EU laws.⁶³ This is no longer the case. Meta recently confirmed to the Australian Parliament, for example, that it has enabled its European customers to opt out of Meta using their personal data to train their generative AI models but has not made this change in Australia, given that there is no legal requirement to do so.⁶⁴

Relying on international regulation would also neglect to address the many local Australian companies, including SMEs, that do not have a global presence – and are therefore untethered to any international regulatory regime.

Australia should align with international regulatory approaches as far as possible, and learn lessons from international regimes. However, it is also important for Australia to adopt a regulatory approach that is attuned to Australia's own legal and political structure and values – *especially* since geopolitical trends are currently in flux. Amongst other things, this is essential for ensuring that AI harms are appropriately mitigated for Australian people and communities – for example, to address implications of AI for Aboriginal and Torres Strait Islander communities, including with respect to representation, non-discrimination, and intellectual property rights.

Taking a passive approach by relying solely on international regulation also undermines regulatory certainty for Australian companies. Because of the variance between major jurisdictions, international experience does not provide a blueprint for Australian workers and firms.⁶⁵ Investors, boards, executives and workers are left without a clear sense of what this means domestically.

Recommendation 4

The Productivity Commission should consider an approach to regulation for AI in Australia that would involve:

- coordinated reform to priority areas of existing law – guided by an expedited, high-level gap analysis, and conducted in parallel to the mandatory guardrails process
- the development of flexible and proportionate framework legislation as the chosen model for mandatory guardrails for high-risk AI. Framework legislation

⁶² Commissioner Stephen King, Commissioner Julie Abramson, Productivity Commission, *Harnessing Data and Digital Technology: Interim Report* (Report, August 2025) 22 <<https://www.pc.gov.au/inquiries/current/data-digital/interim/data-digital-interim.pdf>>.

⁶³ Anu Bradford, *The Brussels Effect: How the European Union Rules the World* (Oxford University Press, 2020).

⁶⁴ David Swan, 'Facebook, Instagram are Using Your Data: And You Can't Opt Out', *The Sydney Morning Herald* (Web Page, 16 June 2024) <<https://www.smh.com.au/technology/facebook-instagram-are-using-your-data-and-you-can-t-opt-out-20240613-p5jlhc.html>>.

⁶⁵ HTI and e61 research

would provide a default set of legal definitions and rules, and could be departed from in other legislation

- dedicated support for the broader regulatory ecosystem, to ensure that regulators have the right mix of capabilities, resources and powers to effectively fulfil their oversight functions in the context of AI.

Recommendation 5

Priority reforms, which have been the subject of extensive multi-sector consultation, and are the subject of existing government reform commitments, should be urgently implemented. These include tranche 2 reforms to the *Privacy Act 1988* (Cth), and the development of a framework for use of automated decision making by government.

Recommendation 6

The Productivity Commission should conduct a thorough cost-benefit analysis of three options for AI regulation (a deliberately permissive approach; a prescriptive approach; and HTI's proposed 'third way'). This assessment should include a focus on:

- costs associated with AI risks
- contingencies that influence the success of productivity measures – including for example, regulatory certainty, consumer trust, incentives or disincentives for good AI governance, support for workers to transition, likely rates of successful AI adoption across particular sectors or types of roles, infrastructure and investment
- distribution of productivity benefits to individuals (including workers and disadvantaged groups) across areas such as income, quality of life and leisure time, as distinct from productivity benefits for business and the national GDP.

Privacy Act reform proposals

The Interim Report makes two recommendations regarding reforms to the Privacy Act:

- The Australian Government should amend the Privacy Act to provide an alternative compliance pathway that enables regulated entities to fulfil their privacy obligations by meeting criteria that are targeted at outcomes, rather than controls-based rules.
- The Australian Government should not amend the Privacy Act to introduce a 'right to erasure', as this would impose a high compliance burden on regulated entities, with uncertain privacy benefits for individuals.

The Privacy Act has undergone a significant review process through the Attorney-General Department's 2022 Privacy Act Review Report,⁶⁶ and in September 2023 the Government committed to implement the majority of the reforms proposed by the

⁶⁶ Attorney-General's Department (Cth), *Privacy Act Review: Report 2022* (Report, 2022) <https://www.ag.gov.au/sites/default/files/2023-02/privacy-act-review-report_0.pdf>.

review.⁶⁷ The reform recommendations developed through this process are well-researched and have been the subject of extensive stakeholder and broader community consultation – some going back almost two decades. Several of these reforms have already been implemented through the ‘first tranche’ of reforms to the Privacy Act, and the Attorney-General has stated that the Australian Government will introduce further reforms as part of a second tranche.⁶⁸

In light of the extensive public and industry consultation, over multiple processes and several years, and the careful scrutiny that the Government has already given to these reforms to the Privacy Act, these reforms should not be abandoned or replaced with less effective privacy protections, without the Productivity Commission articulating much clearer and more detailed justification, and extensive community, civil society and industry support. The primary recommendation for an ‘alternative compliance’ pathway under the Privacy Act is a sweeping and untested solution that would involve an overhaul of privacy law. It would also involve significant uncertainty and costs to business to adjust to a wholly new model, as an alternative to updating the existing model that businesses already comply with, and understand.

In its analysis of privacy law, the Productivity Commission rightly diagnoses problems with the Privacy Act’s over-reliance on consent as a model for protecting privacy in the era of AI, proposing its own ‘outcomes-based’ approach as an alternative. However, addressing problems with the consent model is at the centre of the existing proposals for reforms to the Privacy Act made via the Privacy Act Review. Notably, the introduction of a fair and reasonable test is designed to address issues with ineffectual consent processes. This test would require those collecting and using personal information to act fairly and reasonably regardless of whether the individual, whose personal information they are dealing with, has provided their consent.

Additionally, it is not accurate to classify the Privacy Act as broadly ‘prescriptive’ and ‘controls-based’ rather than ‘flexible’ and ‘outcomes based’ – the Privacy Act is built around privacy principles that are intended to apply flexibly to a range of contexts, and the reforms proposed via the Privacy Act Review would take it further down an outcomes-based path, while ensuring the law remains applicable to rapidly developing AI technology.

The Privacy Act exists to protect the human right to privacy for the Australian people, in line with Australia’s international law obligations. It seeks to instantiate this protection in ways that facilitate other legitimate activities, such as trade and commerce, though those other activities must never be used to justify the violation of this fundamental right. To put this another way, the right to privacy is intended to operate in tension with industry’s legitimate commercial activities, but that tension should not be resolved through unjustified restrictions on the right to privacy.

Relatedly, data rights, including the right to erasure, are important measures that provide individuals with agency over their own personal data and reputation. While there may be some complexities for business associated with implementing such reforms, this is not in itself a sufficient reason to discard them.

Proposals that would undermine the strength of privacy protections are wildly out of step with community expectations. Research conducted by the OAIC in 2023 found that 85% of Australians want more control and choice over the collection and use of

⁶⁷ Australian Government (Cth), *Government Response: Privacy Act Review Report* (Report, September 2023) <<https://www.ag.gov.au/sites/default/files/2023-09/government-response-privacy-act-review-report.PDF>>.

⁶⁸ K&L Gates, ‘Australian Privacy Law Reform Tranche 2: The Time for Conversation is Over’, *Jdsupra* (Web Page, 5 September 2025) <<https://www.jdsupra.com/legalnews/australian-privacy-law-reform-tranche-2-3541970/>>.

their personal information; 87% believe the sale of personal information or trading in personal information is unfair and unreasonable; and 61% are uncomfortable with businesses using AI to make decisions about them that use their personal information. Ultimately, the Australian government has a responsibility to respond to the concerns of the Australian people.

The Interim Report surveys the costs of protecting the right to privacy, including with respect to legal and administrative requirements and compliance costs for business, and claims that the GDPR has ‘stifled firm productivity and innovation’ in the EU. However, this analysis does not take into account the financial and human costs associated with *not* adequately protecting privacy, nor the many harms that the GDPR has likely *prevented*. For example, according to a 2023 IBM report, the average cost to organisations of a data breach in Australia has grown 32% in the last 5 years, reaching AUD \$4.03 million. Globally, the study also found that 95% of studied organisations, including Australian organisations, have experienced more than one data breach.⁶⁹

The argument that privacy regulation necessarily hinders innovation in the AI context is also debatable. California is perhaps the world’s most powerful incubator of new technology products and services despite having some of the strongest data protection rules, such as the California Consumer Privacy Act. Other studies have found benefits of regulation such as the discovery of ‘overlooked efficiencies’ (for example, requirements to clean up data increases overall data quality, which renders data more useful for a range of AI applications).⁷⁰ Australia’s Privacy Commissioner has observed that privacy is not a barrier to the adoption and training of AI models, pointing to OAIC issued-guidance on GenAI that are designed to enable businesses to do just that, in a responsible manner.⁷¹

Recommendation 7

In light of the Australian Government’s acceptance of existing recommendations to reform the Privacy Act, and the reform process currently in train, the Productivity Commission should reconsider its proposals to amend the Act.

⁶⁹ IBM Security, *Cost of a Data Breach Report: 2023* (Report, 26 August 2023)

<<https://mysecuritymarketplace.com/reports/cost-of-data-breach-report-2023/>>.

⁷⁰ Nicholas Martin et al, ‘How Data Protection Regulation Affects Startup Innovation’ (2019) 21 *Information Systems Frontiers* 1307, 1309 <<https://link.springer.com/article/10.1007/s10796-019-09974-2>>; Daniel Bachlechner, Marc Van Lieshout and Tjerk Timan, ‘Privacy as Enabler of Innovation’ (Conference Paper, IFIP International Summer School on Privacy and Identity Management (Privacy and Identity), August 2019) 2 <<https://inria.hal.science/hal-03378973/document>>.

⁷¹ Carly Kind, ‘Upend Privacy Laws for AI at the GP? This Will Not Benefit Patients’, *The Australian Financial Review* (Web Page, 14 August 2025) <<https://www.afr.com/politics/federal/upend-privacy-laws-for-ai-at-the-gp-this-will-not-benefit-patients-20250813-p5mmht>>; ‘Guidance on Privacy and Developing and Training Generative AI Models’, *Office of the Australian Information Commissioner* (Web Page, 23 October 2024) <<https://www.oaic.gov.au/privacy/privacy-guidance-for-organisations-and-government-agencies/guidance-on-privacy-and-developing-and-training-generative-ai-models>>.