



THE UNIVERSITY OF
SYDNEY

Submission to the Productivity Commission: Harnessing Data and Digital Technology & Building a Skilled and Adaptable Workforce

Centre for AI, Trust and Governance

29 September 2025

Contents

Executive Summary	2
Trust is multifaceted	3
Trust is multifaceted	3
AI regulation and the Interim Report's stance in favour of permissiveness	3
The innovation environment is very different from earlier waves of digital technology	3
There is more to good AI deployment than technical innovation or more apps	4
Skilled and adaptable workforces: the importance of investment in human creativity	4
Hallucination is a software and security risk too.....	5
Data and privacy	6
Challenges of outcomes-based regulation	6
The right of erasure	7
Sector-specific data access frameworks	7
Example: Data Sharing and Algorithmic Tenancies	7
References.....	9
Contact.....	10

Executive Summary

This submission responds to the Productivity Commission's Interim Report on **Harnessing Data and Digital Technology** and **Building a Skilled and Adaptable Workforce**. These domains are not merely technical or economic in nature; they are deeply embedded within broader social, cultural, and institutional transformations that demand nuanced and contextually grounded policy responses.

The submission highlights relevant University of Sydney research and makes several key points:

- Trust is multifaceted: a narrow focus on 'secure and responsible data handling' will not be sufficient to create societal trust in data use or AI development;
- The Interim Report's permissive approach to AI regulation takes insufficient account of the changed environment within which innovation now operates. A greater focus on how to ensure an openness of the new environment to *competition and interoperability* as well as small-scale innovation will be important for Australia going forward;
- Innovation doesn't just involve "more apps"; innovation can involve ground-up, careful, deployment that shapes AI use to ensure protection of a wide range of interests: as University of Sydney researchers have seen in their detailed work on AI deployment;
- On privacy: while 'outcomes based' regulation is attractive in theory it has considerable challenges in this context, and in Australia. Below we note critical challenges that need to be dealt with, and raise questions about the value of alternative pathways given that Australian legal frameworks need to be interoperate with global regulatory frameworks;
- We have concerns about proposing more data-sharing frameworks without taking into account power relations in the relevant contexts. Residential tenancy is an area fraught with power imbalances which will likely undermine consumers' ability to 'take advantage of their own data'. The case of tenancy data, for instance, illustrates how algorithmic decision-making can undermine individual agency and exacerbate informational asymmetries. International research underscores the need for governance models that foreground user control, algorithmic transparency, and ethical data stewardship.

A skilled and adaptable workforce must be equipped not only with technical competencies but also with critical capacities to navigate, interrogate, and shape the socio-technical systems that increasingly define our collective futures. This includes addressing the epistemic limitations of AI, such as hallucinations in software engineering, which introduce both productivity burdens and security vulnerabilities. Ultimately, the responsible growth of AI in Australia will depend on a regulatory architecture that is coherent, inclusive, and attuned to the complex realities of digital transformation.

Trust is multifaceted

Trust is multifaceted

The Report highlights that the benefits of increased data use and AI adoption “can only be realised if they are underpinned by trust that data is being handled safely and securely”. In other words, both individuals and businesses must feel confident that digital systems are safe, secure, and responsible before they will fully embrace them. While this emphasis on trust is important, the reasoning remains somewhat narrow.

Trust is treated largely as a matter of security and responsible handling, without sufficient attention to deeper issues such as fairness, transparency, accountability, and long-term governance. As demonstrated in ARC Laureate Professor Terry Flew’s extended research program on Mediated Trust, trust depends on more than just security: it depends also on taking into account a much wider set of interests and concerns in regulation. In our view, the Interim Report pays insufficient attention to the importance of considerations like fairness and broader distribution of benefits from new technologies. Failure to pay attention to these concerns is likely to undermine the drive for innovation or productivity, by undermining the societal trust on which technology adoption depends, and on which our system of government depends.

AI regulation and the Interim Report’s stance in favour of permissiveness

The innovation environment is very different from earlier waves of digital technology

The Interim Report suggests that more permissive AI regulation, and the opening up and sharing of Australian data will promote innovation in Australia. This claim needs to be assessed against a broader market background, and a recognition of the reality of highly concentrated markets, protected by significant data control and intellectual property rights that limit competition.

The development of AI is proceeding in a manner that is diametrically opposed to the foundational logic of the early internet. The internet was initially characterized by open protocols, decentralized development, and a non-commercial, academic ethos; its power lay in its openness. In stark contrast, contemporary AI development is defined by closed, proprietary systems, the centralized control of massive datasets, and intense commercial and geopolitical competition. The power of modern AI lies in its proprietary nature.

AI development is currently dominated by a small group of large companies that emerged victorious from the platform era. Their existing control over vast, proprietary data streams, the essential fuel for training sophisticated machine learning models, gives them an immense, and perhaps insurmountable, competitive advantage in relation to certain kinds of AI development and use. AI as currently developed, therefore, is poised not to disrupt the existing digital

monopolies but to entrench them further, creating even higher barriers to entry for potential competitors.

This inversion of the internet's foundational logic implies that the governance models that eventually and imperfectly emerged to deal with the open web – and arguments for broad data-sharing that imagine a world of open, data-driven innovation, may be completely inadequate for the age of AI. A technological system built on closed, proprietary data and opaque algorithms cannot be governed by the same principles of openness and transparency that were at least theoretically possible in an earlier era. Policies based on more open access to Australian and Australians' data are incapable of promoting and enabling Australian innovation unless issues such as competition and tools to promote interoperability are addressed.

AI must be understood as a "socially shaped technology." It is not an autonomous, magical force appearing out of nowhere, but is fundamentally an aggregation of data and images already circulating in the public domain, shaped at every stage by human values, commercial interests, existing legal rights such as the way we frame intellectual property rights, and institutional priorities.

There is more to good AI deployment than technical innovation or more apps

The key challenge, therefore, is not to surrender to technological determinism but to actively analyse **and shape** the social, economic, and political pathways along which AI develops. Australia can and should seek to enable local innovation and local research and development; invest in AI research and development including along alternative technological pathways and with open models; and seek to shape local deployment and use according to local priorities and norms.

In University of Sydney's own research projects AI deployment in public sector organisations, such as the work of Professor Kalervo Gulson and colleagues with education policymakers and educational institutions across Australia, and Dr Bello y Villarino's work within the NSW public sector, we see the importance of nuance, of careful co-design, consideration of data and data quality, and the importance of recognising a range of interests and goals, including but also beyond innovation, or 'more apps'. This work speaks to the need for care and not just speed; work at the local and small-scale level and not one-size-fits-all (or none) deployment to achieve genuine and sustainable benefits for the public sector in providing services important to Australians.

Skilled and adaptable workforces: the importance of investment in human creativity

A distinctive feature of current generative AI is that its impacts are particularly likely to be in knowledge-related work and occupations, including in particular creative industries and creative work. This contrasts with previous technological waves, which have sought to automate unskilled and semi-skilled work.

While artists have always drawn upon new technologies to enable original forms of creative

practice, there has been a renewed debate around the propensity of Generative AI to deskill creative labour and decentre minority voices and the creative industries of small-to-medium sized nations. There is, however, a gap between those who articulate critical or utopian accounts of the impact of AI on creative industries and processes, and the on-the-ground experience of particular firms, industries, practitioners and practices. The question of what AI can or cannot do is something of a moving frontier and claims about whether or not it can enhance creativity typically elide a clear definition of what constitutes human creativity.

Recent empirical studies (Anderson et al., 2024; Doshi & Hauser, 2024) suggest that while AI systems may outperform individuals in certain tasks, they tend to produce homogenised outputs when deployed at scale. This raises concerns about the erosion of diversity in creative expression, particularly in domains such as writing, design, and visual communication.

Miranda-Peña et al. (2024) highlights the risk that widespread adoption of AI tools may narrow the expressive range of creative outputs, thereby undermining the pluralism that is essential to cultural vitality. To mitigate this, Australia must invest in human-centred design, support for creative experimentation, and educational frameworks that prioritise critical thinking and aesthetic literacy.

A skilled and adaptable workforce is not merely one that can operate digital tools; it is one that can interrogate, repurpose, and reimagine them in ways that reflect and enrich our collective life. This requires a commitment to fostering creativity, diversity, and ethical reflection at all levels of education and professional development.

Hallucination is a software and security risk too

When considering the widespread use of AI, it is important to note that Large Language Models (LLMs) are known to produce hallucinations, outputs that are syntactically plausible but factually incorrect. Software engineering is not exempt from this phenomenon. Tanzil et al. (2024) demonstrate that identifying and correcting hallucinated code introduces additional burdens for developers, thereby diminishing the productivity gains that such tools ostensibly offer.

More critically, hallucinations in code generation may pose significant security risks. One emerging concern is “package hallucination,” wherein LLMs suggest the use of libraries that do not exist. As Spracklen et al. (2025) and Krishna et al. (2025) warn, malicious actors could exploit these hallucinations by registering packages that match the hallucinated names, thereby creating vectors for supply chain attacks. This underscores the need for rigorous validation protocols, secure development practices, and heightened awareness of the limitations of AI-generated code.

Data and privacy

Challenges of outcomes-based regulation

In theory, an outcomes-based privacy framework could offer greater flexibility and efficiency, particularly for AI developers working with large, evolving datasets. The Interim Report, however, may underestimate the challenges:

- This form of regulation can be demanding: firms must exercise judgment and demonstrate they are achieving the right results.
- Measuring outcomes is complex and often slow, meaning harm may be visible only after it has already occurred, making regulation too reactive.
- To avoid too slow a response, and provide the certainty often demanded by firms, regulators likely will seek to issue detailed guidance to define “fair outcomes,” effectively recreating prescriptive rulebooks and adding duplication.
- The approach relies heavily on regulatory capacity, to engage with firm definitions of privacy-respecting outcomes, and audit or oversee those outcomes. Assessing privacy “outcomes” can be complex, especially for AI systems like large language models, where data flows are opaque and impacts are diffuse.
- Data demands compound these challenges, since regulators must collect and interpret evidence on outcomes, a resource-intensive task. There is likely to be conflict between regulators’ need for information/transparency, and organisations’ desire to preserve commercial confidentiality over data sources, data use, and algorithms and their outcomes.
- Without clearly defined enforcement mechanisms, such a system could lead to under-regulation or be perceived as favouring industry over public interest, potentially eroding public trust further.

Claims that regulation will hinder innovation and national progress should be balanced against the risks arising from market concentration and monopoly power, poor consumer outcomes, and limited options for resolving issues without costly legal action.

Overall, while the outcomes based approach to regulation is presented as a modern, adaptable alternative to rigid privacy laws yet without strong oversight, meaningful accountability, and alignment with global frameworks, there is a risk that the approach will prove too vague to offer real protection, and too inconsistent to support the responsible growth of AI technologies in Australia the report seeks to achieve. Without robust oversight and alignment with international standards, Australia may find itself out of step with jurisdictions such as the EU and UK, inviting fragmentation and regulatory arbitrage.

To the extent that outcomes-based regulation is based on strong regulator capacity, we note that Australia’s Privacy Commissioner is radically under-resourced compared to regulators in domains cited by the Productivity Commission, in particular financial services.¹

¹ We note that, as detailed in the ADM+S Submission, it seems from Australia’s Budget Papers that the OAIC has less than 5% of the annual budget of the Australian Securities and Investments Commission (ASIC), and less than 10% of the budget of the ACCC.

The right of erasure

The Interim Report recommends against a ‘right to erasure’ similar to that found in the EU’s GDPR on the basis that it would be costly to enforce and of limited benefit.

In our view, this underestimates or underplays the benefits that can accrue to Australia from harmonising with/interoperating with widely-recognised international norms. The right of erasure is recognised elsewhere, and the **absence** of such a right in Australia – in conjunction with other ways Australian privacy law provides less rights to users – has arguably already had some impact on trust in Australia’s regulatory framework. It has been widely reported in Australia’s news media that Australians’ data has been used to train AI without their consent or even knowledge – with opt outs being offered in countries where privacy law provides individuals with more rights.² This divergence could undermine international trust in Australian-developed AI systems and complicate data portability in cross-border applications.

Sector-specific data access frameworks

The Interim Report proposes that Australia could develop sector-specific data access frameworks, outside the framework of the CDR, including via industry-led initiatives. While we can see the arguable benefits in terms of efficiencies within sectors, we note:

- This risks creating fragmented regulatory environments that hinder the ability of AI developers to build systems that rely on cross-sector data integration. We have seen this already: there has been considerable debate in Australia over when CDR data stops being CDR data.
- Developers may struggle with compliance complexity, especially for multi-domain applications (e.g., AI health chatbots that also manage financial advice);
- It risks creating data silos, impeding training of general-purpose AI models.

Example: Data Sharing and Algorithmic Tenancies

Emerging evidence from overseas and University of Sydney research raises questions about the proposal for more industry-led data sharing initiatives.

For example, researchers at the University of Sydney have a number of projects on rental vulnerability, private rental systems and data hunger in housing, concerning the limitations, risks, and asymmetries inherent in platform-mediated data ecosystems.³ Use Case 2 in the Report refers to real estate tenancy data (p.46). While the suggestion to standardise the

² See, for example, [‘Meta using Instagram and Facebook posts from as far back as 2007 to train artificial intelligence tools’](#), ABC News, 10 June 2024; [‘Meta admits Australians cannot opt out of ‘predatory’ AI data scrape’](#), Sydney Morning Herald, 11 September 2024; [“‘Next gold rush’: Australians’ data being used to train AI”](#), News.com.au, 27 September 2024; [‘Meta’s AI is scraping users’ photos and posts. Europeans can opt out, but Australians cannot.’](#) The Guardian, 11 September 2024. We note that [the position has become more complicated in 2025](#), with some training allowed by the Irish Data Protection Commissioner with safeguards. Notably, when Meta proposed large scale AI training on people’s personal data, more protective European law allowed for a pause, a public discussion, and safeguards. In Australia, it just happened, without debate, and with people offered zero choice. Australian parliamentarians simply told it had already happened.

³ For more information, see the work of [Dr Justine Humphry](#) and [Dr Sophia Maalsen](#).

format of rental ledgers and increase accessibility may have some benefit, the case study fails to consider the wider context of rental applications in a platform data driven era.

The experience of data-sharing is affected by who holds market and other power. In the case of Australia, and data relating to rentals and housing, people seeing housing are not in a position to negotiate what data sharing they will, or will not participate in.

Recent research on Open Banking and tenancy referencing in the UK highlights this point. Ciocanel et al (2024) show that UK tenants feel like they lack choice and control over what data they need to share, as they note “consumers have a more passive role as they only must allow data access and are not involved in the decisions, nor are they in control of the specific details gathered or shared” (1812).

This lack of control and insight into decision making is further highlighted in work by Wallace et al who note that the increasing use of fully automated systems that “generate algorithmically profiled outputs based on open banking data (Ciocănel et al 2024), including detailed income and expenditure for affordability assessments (Wallace et al, 2025 11). However, the use of open banking and APIs cannot predict future employment status and thus cannot always be an accurate marker of ability to pay rent (Wallace et al 2025, 13). While building on existing platform dynamics, generative AI introduces novel and existential threats to information integrity and social trust. Particular dangers are presented by generative AI’s lack of fidelity to facts, evidence, or truth; “deepfakes,” the potential for the mass-automated creation and dissemination of highly convincing misinformation, and the fundamental blurring of the lines between human and machine-generated content. This challenge moves beyond the problem of algorithmic amplification of existing human-created falsehoods to the industrial-scale production of new, synthetic, and potentially deceptive realities. This escalates the crisis of trust from a problem of curation to a problem of creation, threatening the very possibility of a shared understanding of the world.

References

Carlota Perez and Tamsin Murray Leach, Technological Revolutions: Which Ones, How Many and Why It Matters: A Neo-Schumpeterian View, Historical Background Paper nos. WP7-D7.1, Beyond 4.0 (UCL institute for innovation and Public Purpose, 2022).

Carl Benedikt Frey and Michael Osborne, “The Future of Employment: How Susceptible Are Jobs to Computerisation?” *Technological Forecasting and Social Change* 114, no. C (2017): 254–80; Martin Ford, *The Rise of the Robots* (OneWorld Publishers, 2015).

Terry Flew, Jonathon Hutchinson and Wenjia Tang, “Artificial Intelligence and the Creative Industries,” in *De Gruyter Handbook of Creative Industries* (DeGruyter, 2025).

Margaret Boden, “Creativity and Artificial Intelligence: A Contradiction in Terms?,” in *The Philosophy of Creativity: New Essays* (Oxford University Press, 2014).

Butler, T., Gozman, D., & Lyytinen, K. (2023). The regulation of and through information technology: Towards a conceptual ontology for IS research. *Journal of Information Technology*, 38(2), 86-107.

Currie, W. L., Gozman, D. P., & Seddon, J. J. (2018). Dialectic tensions in the financial markets: A longitudinal study of pre-and post-crisis regulatory technology. *Journal of Information Technology*, 33(4), 304-325.

Gozman, D., & Willcocks, L. (2019). The emerging Cloud Dilemma: Balancing innovation with cross-border privacy and outsourcing regulations. *Journal of Business Research*, 97, 235-256.

Vallis, C., Wilson, S., Gozman, D., & Buchanan, J. (2024). Student perceptions of AI-generated avatars in teaching business ethics: We might not be impressed. *Postdigital Science and Education*, 6(2), 537-555.

Ciocănel, A., Wallace, A., Beer, D., Cussens, J., & Burrows, R. (2024). Open Banking and data reassurance: the case of tenant referencing in the UK. *Information, Communication & Society*, 27(9), 1810–1825. <https://doi.org/10.1080/1369118X.2024.2310481>

Wallace, A., Beer, D., Burrows, R., Ciocănel, A., & Cussens, J. (2025). Algorithmic tenancies and the ordinal tenant: digital risk-profiling in England’s private rented sector. *Housing Studies*, 1–21. <https://doi.org/10.1080/02673037.2025.2453005>

Barrett R Anderson, Jash Hemant Shah, and Max Kreminski. 2024. Evaluating creativity support tools via homogenization analysis. In *Extended Abstracts of the 2024 CHI Conference on Human Factors in Computing Systems, CHI EA '24*, New York, NY, USA. Association for Computing Machinery.

Anil R. Doshi and Oliver P. Hauser. 2024. Generative ai enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*, 10(28):eadn5290.

Clarissa Miranda-Pena, Andrew Reeson, Cécile Paris, Josiah Poon, and Jonathan K.

Kummerfeld. 2024. Do LLMs Generate Creative and Visually Accessible Data visualisations. In Proceedings of the 22nd Annual Workshop of the Australasian Language Technology Association, pages 12–29, Canberra, Australia. Association for Computational Linguistics.

Minaoar Hossain Tanzil, Junaed Younus Khan, and Gias Uddin. ChatGPT incorrectness detection in software reviews. In Proceedings of the IEEE/ACM 46th International Conference on Software Engineering, ICSE '24, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400702174. doi: 10.1145/3597503.3639194. URL <https://doi.org/46010.1145/3597503.3639194>

Joseph Spracklen, Raveen Wijewickrama, A H M Nazmus Sakib, Anindya Maiti, BimalViswanath, and Murtuza Jadliwala. We have a package for you! a comprehensive analysis of package hallucinations by code generating llms. In SEC'25: Proceedings of the 34th USENIX Conference on Security Symposium. USENIX Association, 2025. URL <https://arxiv.org/abs/2406.10279.456>

Arjun Krishna, Erick Galinkin, Leon Derczynski, and Jeffrey Martin. Importing phantoms: Measuring llm package hallucination vulnerabilities, 2025. URL <https://arxiv.org/abs/2501.19012>

Contact

Centre for AI, Trust and Governance

aitg.centre@sydney.edu.au

sydney.edu.au

CRICOS 00026A