Filling the Gap: Expert Consultation Highlights AI Regulatory Needs

Submission to the Productivity Commission's 5 Pillars Interim Reports Consultation: Harnessing Data and Digital Technology

Submitted
11 September 2025

Authors
Emily Grundy, Greg Sadler, Luke Freeman

About Good Ancestors
Good Ancestors is an Australian charity dedicated to improving the long-term future of humanity by providing rigorous, evidence-based, and practical policy recommendations for Australia's biggest challenges. We have been deeply engaged in the AI policy conversation since our creation, working with experts around the world and helping to organise Australians for AI Safety.

Contact
If you would like to discuss this submission, please let us know at contact@goodancestors.org.au.

Good Ancestors

Table of Contents

Introduction

The Productivity Commission's interim report on Harnessing Data and Digital Technology called for a pause on mandating AI guardrails, gap analyses of existing rules, and AI regulation as a last resort. This Submission attaches an expert gap analysis completed by Good Ancestors to answer the Productivity Commission's request.

The expert gap analysis finds that, while existing regulators are well placed to address many AI threats, general-purpose AI models pose specific threats that are likely and consequential.[1] The gap analysis found that up to 93% of experts consider current Government measures inadequate for managing threats from general-purpose AI models. AI threats can transcend existing regulatory boundaries and require coordinated, upstream intervention that existing regulators are not well-placed to provide.

The Australian Government began consulting on AI safety in June 2023[2] and formally acknowledged the types of AI threats covered by our gap analysis when it signed The Bletchley Declaration in November 2023.[3] Over subsequent years of consultation, AI has continued to develop with AI models crossing chemical, biological, radiological, and nuclear (CBRN) risk thresholds in 2025,[4,5] with labs predicting Artificial General Intelligence as early as 2026.[6] AI development has outpaced Australia's ability to consult. Further consultation will lead to further delay, leaving Australia further behind.

The net benefit Australia gets from AI is the benefit minus the harm. The Productivity Commission's interim report takes the potential benefits of AI seriously, dedicating appendix B.4 to quantifying the economic potential. Yet, the interim report provides no comparable assessment of potential AI-related costs or harms. The interim report cites the International AI Safety Report, but largely ignores AI expert warnings and fails to incorporate quantitative harm assessments into its economic analysis.

Instead, the interim report explores how regulation could go wrong. The report extensively analyses how 'burdensome regulatory requirements can stifle innovation', but provides no parallel analysis of the economic costs of inadequate regulation or AI-related incidents. This leads to the report being overly optimistic, playing up benefits and playing down risks. We think the report should adopt a balanced approach in how it forecasts both benefits and harms. There is strong and credible evidence to support both the opportunities and risks of AI, and this evidence is substantial enough to justify efforts to mitigate the risks in the interests of achieving a net-benefit for Australians.

[1] Sadler, G et al. (2025, August 19). *Australian AI Legislation Stress Test: Expert Survey*. Good Ancestors.
[2] Department of Industry, Science and Resources. (2023, June 1). *Safe and responsible AI*. Australian Government.
[3] Department of Industry, Science and Resources. (2023, November 3). *Australia signs international declaration on AI safety*. Australian Government.
[4] Google DeepMind. (2025, August 1). *Gemini 2.5 Deep Think Model Card*. Google DeepMind.
[5] OpenAI. (2025, August 7). *GPT-5 System Card*. OpenAI.
[6] Anthropic. (2025, March 6). *Anthropic's recommendations: OSTP U.S. AI action plan*. Anthropic.

Evidence from the Australian AI Legislation Stress Test

Good Ancestors' Australian AI Legislation Stress Test directly addresses the Commission's call for gap analyses by stress testing current Australian laws against specific AI threats. We surveyed 64 experts with expertise spanning AI, public policy, cybersecurity, national security, and law, and conducted a detailed legal analysis of how current laws would respond to AI threat scenarios.

These stress tests help regulators and policymakers determine whether AI threats can be tolerated, whether existing laws are adequate, whether they can be strengthened, or whether new laws for general-purpose AI are needed.

Key findings include:

- Across all threats assessed, the vast majority of experts found current Government measures inadequate.
- Existing regulators are well placed to address many, but not all, AI risks. 'Upstream' regulation at the model development level would be more efficient and effective than 'downstream' regulation for addressing AI threats from general-purpose AI that transcend traditional regulatory boundaries.[7]

We have provided the report as a supporting document alongside our submission.

Transparency and validation

Good Ancestors has made the stress test, including the methodology, threat definitions, and stress test scenarios, widely available—please contact us for survey resources and the cleaned dataset. We have met to discuss the report with the lead department and used the Australian Government regulator stocktake to write to relevant regulators, inviting them to replicate the analysis with the benefit of their subject matter-specific knowledge.[8] Throughout these engagements, no regulator has expressed ownership of these risks or competency in mitigating them.

The Commission's final report can draw from our AI Legislation Stress Test and the broader literature to boost the analytical rigour of its discussion of AI risks to match the rigour dedicated to AI benefits.

[7] Mansur, E. T. (2011). *Upstream versus Downstream Implementation of Climate Policy*. The Design and Implementation of US Climate Policy, p 179-193, National Bureau of Economic Research, Inc.
[8] Department of Finance. (Accessed 2025, September 10). *Australian Government regulator stocktake*. Australian Government.

Limitations of the Commission's approach

The Commission's approach overlooks evolving risks that undermine existing regulatory approaches

The Commission's Draft Recommendation 1.1 calls for comprehensive gap analyses across all agencies. While we should avoid unnecessary and duplicative regulation, this approach has substantial flaws and would delay necessary action. Clear evidence of regulatory gaps already exists, and AI capabilities are advancing rapidly.

AI use cases and risks are constantly shifting

The Commission suggests mapping AI uses, assessing specific risks, and determining whether existing frameworks are sufficient or could be modified. However, this approach will only ever produce a point-in-time snapshot and doesn't align with the realities of AI development or adoption. AI systems are achieving such rapid progress that many benchmarks designed to test their limits are becoming 'saturated' and no longer helpful for measuring capabilities.[9] General-purpose AI systems can be used and misused in a near-infinite number of ways. Agentic AI has failure modes that are hard to anticipate. By the time regulators have conducted comprehensive gap analyses—if that is even possible—the risk and opportunity landscape will have shifted.

The best way to keep pace with a dynamic landscape of risks and opportunities is an **expert regulator** paired with an **independent technical safety body**–akin to the Civil Aviation Safety Authority (CASA) at the Australian Transport Safety Bureau (ATSB), respectively. There is already sufficient evidence of AI risks outside the remit of existing regulators to justify an AI Act creating an AI regulator and an AI Safety Institute. An AI regulator can modify its approach based on the evolving risk landscape, and an AI Safety Institute is best placed to build trusted relationships with industry and provide technical advice to government and the regulator.

General-purpose AI undermines existing regulatory approaches

General-purpose AI systems undermine existing regulatory approaches by operating across multiple regulatory domains and creating novel liability gaps.

Sector-specific regulation burdens regulators and industry

General-purpose AI systems operate across multiple regulatory domains simultaneously, including financial services, communications, and health. For example, a single general-purpose AI system could be used in financial services (involving ASIC, APRA, and RBA obligations), communications (triggering ACMA, DITRDCA, Home Affairs, and AGD obligations), and health applications (subject to AHPRA, TGA, and state-based obligations).[10]

[9] Stanford Institute for Human-Centered Artificial Intelligence (2025). *AI Index Report 2025: Chapter 2*. Stanford University.

[10] ASIC: Australian Securities and Investments Commission; APRA: Australian Prudential Regulation Authority; RBA: Reserve Bank of Australia; ACMA: Australian Communications and Media Authority; DITRDCA: Department of Infrastructure, Transport, Regional Development, Communications and the Arts; Home Affairs: Department of Home Affairs; AGD: Attorney General's Department; AHPRA: Australian Health Practitioner Regulation Agency; TGA: Therapeutic Goods Administration.

The current regulatory approach creates impractical compliance burdens for both regulators and industry. Without general laws defining standards for general-purpose AI systems and assigning liability amongst developers, deployers, and users:

a) Every regulator must set individual standards for their domain, and
b) General-purpose AI systems have to separately navigate the requirements of every regulator.

As a result, either general-purpose AI systems incur prohibitive compliance costs trying to meet every regulator's standards, or they're functionally prohibited in favour of narrow, sector-specific AI. This regulatory burden substantially hampers the potential economic benefit of general-purpose AI and invites non-compliance.

General-purpose AI agents create novel liability gaps

General-purpose AI agents also challenge fundamental legal concepts in responsibility and accountability. Our legal system is built on the principle that a wrongful act must typically be coupled with a culpable mental state (‘actus reus’ and ‘mens rea’ in criminal law). However, AI agents sever this connection. A user's intent is limited to their initial, often broad, prompt, and a system may go on to perform a harmful act far removed from that original instruction and entirely unknown to the user. This disconnect creates a legal gap that current civil and criminal law is not able to navigate. This gap is bridged elsewhere in law, such as in principal-agent responsibility, but it seems unlikely that an Australian court would hold an AI agent developer responsible where an agent exceeds its authority in the same way as a real estate agent or a lawyer acting on behalf of a client.

AI agents are being adopted across the economy, but currently operate without clear regulatory standards. Our stress test scenarios demonstrate how harm can occur through agents performing incompetent actions or by exceeding user authority. Expert consultation indicates that these risks are likely and consequential, and that current Government measures are inadequate for mitigating them.

We need cross-cutting regulation

These challenges require general laws that set standards for AI systems across sectors, define liability allocation between developers, deployers, and users, and provide regulatory certainty while reducing compliance burdens. We need coordinated, upstream intervention to balance the risks and opportunities of general-purpose AI in the overall public interest.

We need AI-specific regulation

Draft Recommendation 1.2 positions AI-specific regulation as a 'last resort' (p. 20). This is inconsistent with Australia’s approach to regulation across other technologies with specific risks, including therapeutic goods, aviation, communications, and automotive. Even laser pointers, bicycles, and button batteries are subject to specific regulations to address their specific risks. Cars have airbags, banks adopt policies limiting money laundering, and aircraft have escape slides. These common-sense risk mitigations emerge from specific rules. They don’t come from general laws and aren't 'last resort' measures.

AI systems capable of autonomous operation, self-modification, or providing dangerous capabilities warrant efficient and effective safety standards tailored to their specific risks. Holding malicious actors liable after they use AI to cause national-scale harm is not sufficient. While existing laws and regulators can manage many AI risks, some significant risks are not currently mitigated.

Not all AI threats can be managed via existing approaches

The Commission emphasises that 'few of AI's risks are wholly new issues', highlighting that problems like deepfakes, malicious actors, and copyright infringement existed before AI. However, this:

1. **Undervalues the AI risks that are novel.** Our expert consultation and stress test demonstrate that novel AI risks, like AI agents exceeding their authority, misuse of open-weight AI models, and loss of control scenarios, are likely and consequential, and the Government is currently ill-prepared to mitigate them. For instance, the stress test details how open-weight AI models present a novel blend of risks and opportunities, and that currently the Government has no tools to ensure that balance is struck in the overall public interest.

2. **Overlooks that AI can reduce the expertise required for malicious acts**, enabling actors with minimal skills to commit sophisticated attacks that previously required substantial technical knowledge. Evidence demonstrates how AI democratisation, combined with insufficient safety standards, can result in misuse. Anthropic's August 2025 Threat Intelligence Report detailed how Claude has been misused for large-scale extortion and AI-generated ransomware by actors with only basic coding skills. This includes 'vibe hacking', where attackers with no technical expertise completed sophisticated cyber attacks after jailbreaking large language models.[11]

3. **Underestimates how AI risks bypass the scope of current mitigations.** For instance, some risks are managed via professional standards (e.g., requirements on doctors or lawyers) or place-based requirements (e.g., regulation of labs or facilities). These regulations have been effective because those people and places were “chokepoints” for threats. As AI reduces and removes person-based and place-based chokepoints, the efficient and effective intervention point moves outside the ambit of current regulators. Expanding the scope of existing regulators is unlikely to be effective. For instance, the Medical Board of Australia and the Australian Health Practitioner Regulation Agency, which are responsible for medical professional standards, are not well placed to regulate AI developers or deployers, even as AI products begin standing in for medical professionals.

4. **Ignores that we routinely regulate skills or tools that significantly enhance known kinds of wrongdoing.** Most obviously we regulate weapons, rather than relying only on prohibitions on violent crime. We regulate banks to limit the ability of wrongdoers to launder money or finance terrorism. Regulating dangerous AI capabilities at their source is consistent with long-standing regulatory approaches.

General-purpose AI that does not meet appropriate safety standards amplifies risks across numerous fields. Targeted interventions at key points in the AI supply chain and lifecycle would be efficient—only needing a single intervention—and effective—intervening at pivotal points. If those risks are instead addressed on an ad-hoc basis, they would be inefficient—resulting in significant duplication across fields—and ineffective—because many subject matter regulators do not have the power to regulate across the AI supply chain and lifecycle.

AI-specific regulation can be technology-neutral

Draft recommendations 1.1, 1.2, and 1.3 highlight the value of technology-neutral regulations. This is widely accepted for two key reasons:

1. Technology-neutral laws prevent a good or service from being disadvantaged relative to an equivalent good or service based on an arbitrary technical characteristic.

[11] Anthropic. (2025, August). *Threat intelligence report*. Anthropic.

2. Technology-neutral laws need to be updated less frequently as technology evolves.

While these arguments are strong, the interim report misunderstands what "technology neutral" means in legal theory and how laws can be drafted to avoid these pitfalls while still addressing specific risks.[12]

The *Therapeutic Goods Act 1989* is an example of successful technology-neutral legislation. The Act defines 'therapeutic goods' by their function rather than the specific technological composition. The regulator asks whether a good is used in relation to preventing, diagnosing, curing, or alleviating a disease, ailment, defect, or injury. A good can be 'therapeutic' based on this technology-neutral approach, regardless of its specific chemistry or construction. This approach achieves both the desirable outcomes of removing arbitrary technical characteristics that could result in unfair disadvantage, and allowing new technology to be regulated without amending the Act.

This same technology-neutral approach applies in other cases, like the *Telecommunications Act 1997,* which regulates communications and communication services providers, or the *Civil Aviation Act 1988*, which regulates air navigation against performance metrics. The *Privacy Act 1988* also takes this technology-neutral approach, including tiers of rules that guide the application of high-level principles to specific technologies.

Australia can draft technology-neutral AI regulations by defining AI models and systems based on functions or risk tiers, not by reference to specific machine learning approaches.

For instance, the OECD[13] and The Safe and Responsible AI in Australia consultation[14] used the same technology-neutral definition of AI systems:

> *An 'AI system' is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.*

The EU AI Act's technology-neutral definition of AI is:[15]

> *'AI system' means a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments*

Canada's AI and Data Act's proposed technology-neutral definition of AI was: [16]

> *'AI system' means a technological system that, autonomously or partly autonomously, processes data related to human activities through the use of a genetic algorithm, a neural network, machine learning or another technique in order to generate content or make decisions, recommendations or predictions.*

We recommend that the final report be updated to clarify that AI regulations, and an Australian AI Act, can be technology-neutral, that examples of technology-neutral definitions of AI already exist, and that these yield all the benefits of technology-neutral regulation.

---

[12] Bennett Moses, L. (2007). *Recurring Dilemmas: The Law's Race to Keep Up With Technological Change*. University of New South Wales Law Research Series, 21.

[13] Organisation for Economic Co-operation and Development. (2024, March). *Explanatory memorandum on the updated OECD definition of an AI system*. OECD.

[14] Department of Industry, Science and Resources. (2024). *Proposals paper for introducing mandatory guardrails for AI in high-risk settings*. Australian Government.

[15] EU Artificial Intelligence Act. (2024). *Article 3: Definitions*. European Union.

[16] Parliament of Canada. (2022). *Digital Charter Implementation Act, 2022 (Bill C-27)*. Government of Canada.

Developers need appropriate accountability

Efficient and effective regulatory obligations match the risk to the actor best able to mitigate it. Without appropriate developer regulations, obligations will often fall only on Australian deployers and users–who have little practical ability to manage the risks of 'black box' AI systems. AI developers currently have no obligations to test their systems, verify those tests with third parties, or publish the results. Some developers do limited testing voluntarily, but independent evaluators have found these voluntary efforts inadequate.[17]

Absent an AI Act, Australian businesses are in an impossible situation. Not only are they required to mitigate risks that are beyond their technical expertise, but they are not necessarily even aware that the risks exist. If regulators or courts find imposing liability inappropriate in those circumstances, it would leave Australians exposed to harm without any practical recourse or access to justice.

The starting point has to be an obligation for AI developers to assess and disclose model risks and capabilities. The logical extension is an obligation to meet safety standards when risks are identified.

Lack of trust slows adoption and limits economic benefits

Australians' lack of trust in AI threatens its economic potential. Australians are among the least trusting of AI globally, with 96% holding concerns about generative AI[18] and only 36% trusting AI systems.[19] Trust drives adoption, and the Tech Council of Australia's analysis shows adoption pace dramatically impacts AI's value to the Australian economy. Its modelling indicates that slow adoption produces 61% less economic benefit (tens of billions of dollars) than fast-paced adoption.[20]

While the Commission acknowledges that sensible regulation "build[s] community trust and business confidence" (p. 10), the proposed delay undermines this goal. Government inaction, combined with mounting AI risks, could further erode trust instead of building confidence. Already, only 30% of Australians believe current Government measures are adequate for making AI safe, with 77% supporting AI regulation.[21] As capabilities continue to outpace regulation, this trust deficit will likely worsen, impacting AI adoption and reducing the net benefit of this technology.

Absent proposals for building trust, the Commission should adjust its economic benefit projections to reflect slow-adoption scenarios. The Commission's economic analysis assumes AI adoption rates of 70% (i.e., AI is used in 70% of tasks where it could be, p. 102). This figure is based purely on economic profitability, assuming that if it makes financial sense for an actor to use AI, then they will use it. This overlooks the myriad of non-economic factors that affect technology adoption, including trust and regulatory certainty. In stark contrast, the Tech Council of Australia estimates slow-pace adoption rates could be as low as 13%. We recommend that, if the final report does not include credible trust-building measures, the AI adoption rate estimate be substantially adjusted to reflect the impact of these non-economic factors.

[17] Future of Life Institute. (2025, July 17). *AI Safety Index: Summer 2025*. Future of Life Institute.
[18] Australian Competition and Consumer Commission. (2025, June 23). *Digital Platform Services Inquiry: Final Report*. Australian Government.
[19] KPMG and University of Melbourne. (2025, May). *Trust, Attitudes and Use of AI: Global Report*. KPMG.
[20] Microsoft and Tech Council of Australia. (2023, July). *Australia's Generative AI Opportunity.* Microsoft and Tech Council of Australia.
[21] KPMG and University of Melbourne. (2025, May). *Trust, Attitudes and Use of AI: Global Report*. KPMG.

Australia's Path Forward: Leadership and AI Regulation

Why Australia must act now

The pace of AI development makes regulatory delay dangerous

Regulatory delay becomes increasingly dangerous as AI capabilities advance. Draft recommendation 1.3 suggests pausing mandate guardrails for high-risk AI while we wait for every government agency to complete detailed gap analyses. The Commission doesn't specify what 'complete' entails, how many agencies would be involved, or provide timelines. Based on previous AI consultation timelines, such a comprehensive review of Australia's regulatory landscape could take years, during which AI capabilities will continue advancing. Already, since the Government began consulting on these issues, AI agents have become widespread, AI models have crossed new CBRN thresholds,[22,23] and AGI capabilities are forecasted as early as 1-2 years away.[24]

The Commission's Terms of Reference require analysis of "implementation feasibility and risks" (p. 5), yet the report fails to address risks of its proposed approach. While one section appears dedicated to these risks ("Are there risks of waiting for gap analyses before taking regulatory action?", p. 22), it actually focuses on the risks of mandating guardrails rather than the dangers of inaction. The Commission should fulfil its Terms of Reference by analysing how regulatory delay could leave Australia exposed to emerging and growing AI risks.

Australia should play an active role in global AI governance

The interim report states that "AI specific regulation is relatively rare globally" but frames this as reason to wait and be a "regulation taker" rather than pursue leadership (p. 22). This underestimates Australia's opportunity to shape global AI governance standards. The vast majority of Australians, 94%, believe Australia should play a leading role in the international governance and regulation of AI.[25] We've been regulatory leaders in the past: pioneering cigarette regulation and both inventing and mandating cockpit voice recorders in commercial aircraft. Australian innovation and regulatory foresight can become global standards.

The Commission also notes Australia should be "an active participant in international forums" and ensure "its interests are strongly represented" (p. 22), but meaningful participation requires sovereign technical capability, not just diplomatic engagement. Currently, Australia largely relies on foreign AI Safety Institutes and company self-assessments for risk evaluations. Australia needs domestic technical capability to meaningfully contribute to international networks, set strategic priorities reflecting Australian values, and independently verify external risk evaluations.

[22] Google DeepMind. (2025, August 1). *Gemini 2.5 Deep Think Model Card*. Google DeepMind.
[23] OpenAI. (2025, August 7). *GPT-5 System Card*. OpenAI.
[24] Anthropic. *Anthropic's recommendations: OSTP U.S. AI action plan*. Anthropic.
[25] Saeri, A., Noetel, M., & Graham, J. (2024, March 8). *Survey Assessing Risks from Artificial Intelligence: Technical Report*. Ready Research and The University of Queensland.

Our recommendations

The overall benefit Australia achieves from AI will depend on how effectively we manage risks *and* harness benefits. The Commission's final report should acknowledge the evidence that justifies immediate regulatory action rather than recommending further delay. Without action, net benefits decrease.

Recommendation 1: Acknowledge regulatory gaps

Evidence, including our stress test, demonstrates that certain AI risks are likely, consequential, and unable to be effectively addressed by existing laws or regulators. The final report should be updated to acknowledge that:

- Some AI risks are novel, cross cutting, or so consequential that they require new regulation.
- The pace of AI development consistently exceeds regulatory review processes and government consultations.
- Sufficient evidence of regulatory gaps exists to justify building governance structures now.

Recommendation 2: Create appropriate governance structures

The appropriate governance structure is analogous to aviation–an overarching act creating a statutory expert regulator and an independent technical body.

Australia needs these three complementary elements:

- **An AI Act** providing a technology-neutral legislative framework for AI regulation. See below box for an explanation of how an AI Act could work.
- **An AI regulator** to enforce standards, coordinate across sectors, and adapt to evolving risks.
- **An Australian AI Safety Institute** to conduct independent technical evaluations, accelerate safety research, and have a seat at the global table.

Recommendation 3: Include credible trust-building measures or reduce assumptions about the economic benefit of AI

Australians do not trust AI, and that lack of trust is slowing AI adoption. Continued government inaction and mounting AI risks, will likely reduce trust further. The PC should either recommend credible trust-building measures that would remove trust as a barrier, or revise down its forecasted AI adoption rates to reflect these non-economic factors. Our preference is credible trust-building measures.

How could an AI Act work?

The interim report expresses concern about the potentially broad scope of AI regulation and the lack of international consensus on how to regulate AI (p. 11). Well-supported approaches adopted from established regulatory regimes can address both concerns. Specifically, an AI Act can require the regulator or minister to adopt internationally recognised standards or best practices. This means Australia can avoid falling behind or getting ahead of the global regulatory consensus. It also positions Australia as a global consensus builder, ready to work with partners to shape global norms. Further, an AI Act can specifically preserve the operation of existing regulators and avoid duplication. An AI Act can leave CASA to deal with AI in aviation, or the TGA to deal with AI in medicine, with an AI regulator addressing only the gaps and coordinating with existing regulators.

The regulation of AI does not have to be complicated or risk overreach. An AI Act could have four key features:

1. **Create a technology-neutral definition of AI.**
    - A technology-neutral AI definition would define AI based on its functions, not the specific technology, and differentiate between low-risk AI or narrow AI and high-risk AI, to ensure that inherently safe products are not regulated unnecessarily. Examples are provided above.
    - An AI Act could adopt the model in Section 7 of the *Therapeutic Goods Act 1989*, which allows the Secretary to deem when goods are or are not therapeutic to allow the resolution of scope disputes as AI evolves.

2. **Establish the primacy of existing regulators.**
    - Existing regulators are experts in their domains and should retain primacy within their spheres of responsibility. This could be supported by a mechanism for regulators to "refer power" to the AI regulator where a coordinated or centralised approach is warranted.

3. **Establish an expert AI regulator.**
    - This could be modelled on existing expert regulators, like CASA or the TGA.
    - This should include regulatory and compliance powers applying across the AI supply chain.

4. **Allow the adoption of internationally recognised standards or best practices** as they emerge and apply them to specific parts of the AI supply chain and lifecycle.
    - This could be modelled on Section 98 of the C*ivil Aviation Act 1988*, or a range of other regulations that allow for delegated legislation or disallowable instruments after having regard to standards and best practice.
    - This could be supported by appropriate definitions of the AI supply chain (e.g., AI developers, AI deployers, AI infrastructure operators, etc.) and the AI lifecycle. Taxonomies of AI supply chains and AI lifecycles already exist, including from the OECD.[26]

Optionally, an AI Act could:

1. **Ban unacceptable uses of AI**.
    a. There is growing agreement about specific unacceptable AI capabilities, like "nudify" or social scoring.[27] An AI Act would be a logical place to enact such a ban, including because it would have the framework to target specific areas of the AI supply chain and AI lifecycle and be backed by a regulator.

2. **Create the statutory basis for an Australian AI Safety Institute**.
    a. This could be modelled on Part 2 Division 1 of the *Transport Safety Investigation Act 2003*, which creates the ATSB as an independent technical agency. A non-regulatory technical institute is a natural partner for a technical regulator and can help build a positive safety culture focused on good outcomes, not apportioning blame.

[26] Organisation for Economic Co-operation and Development. (2022, February). *OECD Framework for the Classification of AI systems*. OECD.
[27] EU Artificial Intelligence Act. (2024). *Article 5: Prohibited AI Practices*. European Union.

Conclusion

Australia can support AI innovation while managing catastrophic risks. The Productivity Commission's recommendations risk delay while AI capabilities advance unchecked. Expert assessment demonstrates that existing measures are inadequate for managing national-scale AI risks. Without appropriate laws for general-purpose AI, we will have gaps, parties being held responsible for risks they cannot mitigate, and access to justice challenges for people experiencing AI harms. A well-designed AI Act, drawing on established approaches, can target the risks and increase the net benefit that Australians achieve from AI. Australia should act now, implementing safeguards that complement existing frameworks.