



13 May 2024

Select Committee on Adopting Artificial Intelligence (AI)
PO Box 6100
Parliament House
Canberra ACT 2600

By email: aicommittee.sen@aph.gov.au

Dear Sir/Madam

Select Committee on Adopting Artificial Intelligence (AI)

AFMA welcomes the opportunity to provide comment to the Select Committee on Adopting Artificial Intelligence (AI).

AFMA represents the interests of over 125 participants in Australia's wholesale banking and financial markets. Our members include Australian and foreign-owned banks, securities companies, treasury corporations, and traders across a wide range of markets and industry service providers. Our members are the major providers of services to Australian businesses, institutional investors, and superannuation funds.

AI presents an exceptional opportunity for Australia to leverage its technological literacy to drive efficiencies in knowledge work, and ensure our economy remains competitive in the emerging AI supported global economy.

AFMA has supported ¹ Australia having AI standards that are interoperable with international standards, use globally consistent terminology, leverage international regulatory outcomes, and are careful to preserve the potential efficiency gains that AI is enabling.

AFMA continues to support careful, targeted, and proportionate regulation to minimise harms from AI in areas where these risks exist.

In our view, the risks associated with AI are mainly in non-wholesale contexts, and in wholesale financial markets and banking contexts can be targeted by firms with sensible risk management practices.

¹ [R62-20-Australia-s-AI-Action-Plan-AFMA-Submission.pdf](#), see also [R43-23 DISR - Safe and Responsible AI - Consultation - AFMA Submission.73b0f530b31d3.pdf \(storage.googleapis.com\)](#)

Market participants and regulators are alert to the potential for risks that are relevant to wholesale. And while many of these are already covered by existing regulation, for example those relating to frontrunning, spoofing and insider trading, should new theoretical risks emerge as practical matters the industry stands ready to work closely with the Government to anticipate and manage those risks in the most efficient way possible.

We note also that given the international context of AI development and delivery, a key part of enabling Australian AI success lies with supporting firms utilising third party deployment of systems.

We trust our comments below are of assistance to the Committee.

Yours sincerely

A handwritten signature in black ink that reads "Damian Jeffree". The signature is written in a cursive, slightly slanted style.

Damian Jeffree

Head of Financial Markets, Exchanges and Digital

AFMA's Response to the Select Committee's Terms of Reference

- a) recent trends and opportunities in the development and adoption of AI technologies in Australia and overseas, in particular regarding generative AI;

AI has recently had a defining breakthrough with the exponential improvements in Large Language Models (LLMs) and other types of generative AI that followed the application of large increases in training data and compute coupled with hardware improvements including the development of large purpose-built GPU-style AI chips.

These developments open many more areas of application for AI, such as automated coding, more sophisticated language parsing for compliance and legal purposes, and analysis of unstructured data for trading, research and other purposes.

Fully and properly deployed, AI will make very substantial contributions to productivity growth. Goldman Sachs estimates that Generative AI could raise global GDP by 7% and lift productivity growth by 1.5 percentage points over a 10-year period.²

It is incumbent upon national governments to support their economies maximizing these benefits in a timely way, as they are so significant, and will be a key area of international competition over the coming years.

- b) risks and harms arising from the adoption of AI technologies, including bias, discrimination and error;

AFMA supports the recognition that there are some areas where AI could be used by bad actors, and that there are facets of generative AI that require sensible management frameworks, particularly where they interact with natural persons.

In deploying these sensible management frameworks, firms will drive outcomes that:

- ensure alignment with existing regulatory requirements;
- are appropriately incentivised;
- are appropriately tested in line with existing requirements to ensure they are robust to noisy data and other input discrepancies;
- have appropriate monitoring by other AI or human agents (or have a similar separately incentivised self-monitoring function built-in); and
- meet existing requirements for other risk management measures including kill switches where appropriate.

Retail/Wholesale

This connection of many of the risks to natural persons is a key point. Many of the risks associated with Generative AI technologies are not as strongly connected with wholesale banking and financial markets. These areas of the economy are generally concerned with

² [Generative AI Could Raise Global GDP by 7% \(goldmansachs.com\)](https://www.goldmansachs.com/insights/articles/generative-ai-could-raise-global-gdp-by-7%-(goldmansachs.com))

business-to-business interactions and are areas where the participants are typically well-resourced and well-placed to manage the risks themselves. Allowing firms to manage these risks typically results in more efficient outcomes that benefit the engine room of the economy.

For example, wholesale market investors would not benefit from regulatory 'protection' from the judicious use of AI in developing financial market research. The market is fully capable of finding the optimal level of AI usage, which may well vary from time to time. Regulatory intervention is likely to create costs and barriers to innovation in such an area.

Following from this, the first step in targeting AI regulation should be to recognise and respect the Retail/Wholesale distinction which already underpins much of the financial services law.

This important distinction is between the consumer protection space (retail) and the wholesale (business) space. While the sum of regulatory actions over recent years has arguably diluted the clarity of this division, we are encouraged by the work of the ALRC which should reinforce the distinction.

Bias and Discrimination Challenges with Generative AI (and other Deep Learning Artificial Neural Network-based systems)

There are, as the Terms of Reference suggest, a number of potential risks and harms arising from the adoption of AI technologies.

An important distinction is that these tend to be most associated with Generative AI.

Forms of AI such as carefully hand-coded logistic regression classifier systems (e.g. IBM Watson) that might use fuzzy logic to generate outputs, or non-generative ANN-based machine learning systems that might learn to recognise handwriting, obstacles on the road and the like, typically do not have the same potential issues with bias and discrimination that might, in some circumstances, be associated with Deep Learning-Artificial Neural Network (DL-ANN)-based systems such as Generative AI.

Generative AI systems tend to be less predictable and have lower 'explainability' than the types of AI systems with more explicit logic designs that had dominated deployment in the prior period.

Non-DL-ANN systems have been sorting hand-written mail, reading numbers on cheques, playing chess, recognising faces, driving cars, flying aeroplanes and many other applications for decades without the need for much in the way of AI-specific regulatory intervention.

Regulatory responses driven by concerns around bias and the potential for bias and discrimination should be limited to systems where these are potential issues, leaving development in non-DL ANN systems unencumbered by what could be irrelevant risk management.

The deployment purpose is also a critical factor in the risk analysis. The deployment of a Generative AI system to create ideas for floral arrangements is unlikely to have the same risk profile as one used to sort CVs for job interviews.

Once these risks are identified as significant in a particular scenario by a firm, measures should be taken by the firm to address those risks.

In regard to the potential for ‘error’, we suggest a model below for dealing with imperfection that is likely to also be suitable for deployment in addressing bias and discrimination risks.

Dealing with AI Imperfection

As incredibly powerful as they are proving to be, Generative AI systems, particularly at this point in their development are imperfect.

Many of the ways that these imperfections manifest will no doubt be addressed over time, including the potential for ‘hallucinations’, and the inclusion of data from erroneous, misleading or otherwise undesirable sources in the creation of responses.

As we discuss below, AFMA holds that in most cases these risks can be appropriately managed through a combination of the actions of both the developers of Generative AI systems and the firms that deploy them.

As such there is no need to let ‘perfect’ be the enemy of ‘good’ in AI regulation. As a jurisdiction we can safely deploy AI and enjoy its benefits while minimising its risks.

To make this a reality the current approach to regulation of machine systems will need to adjust.

Currently, regulation expects perfection from machine systems for the factors within their control. This is not unreasonable as for most typical machine systems they follow pre-determined logic which is well-understood, predictable and stable. This is less the case with Generative AI (and other DL-ANN systems).

Generative AI systems are, strictly speaking also deterministic, except insofar as they may artificially introduce randomisation into their processes to benefit the creativity of their output. The introduction of this randomisation as well as the inherent difficulties with ‘explainability’ create a pseudo-indeterminism.

This means that practically speaking firms cannot be sure how systems will respond to every conceivable possible scenario.

It is worth noting, however, that firms can face similar risks with employees. The same sensible risk management policies and procedures should be adaptable for application to DL-ANN systems.

The table below shows how some of these controls might be adapted:

Employee control measure	DL-ANN AI agent control measure
Check for good education	Check for training against known-good data sets (potentially in addition to unstructured data)

	Check manufacturer for information about the quality and security of AI agent
Interview to check for character and fit	Test output in various scenarios to check for fit
Use manager to supervise work output	Use a separate system to monitor output
Consequence management for undesirable output	Agent management protocol (e.g. retrain/decommission) for undesirable output
Provide ongoing training to drive continuous improvement and ensure that employee remains up to date	Provide ongoing training to drive continuous improvement and ensure the AI agent remains up to date

The model we propose might be considered an ‘employee model’ for safe AI-agent deployment.

Third party AI

In the Australian context, while we expect significant local development of AI technologies over time, given the global nature of the technology, many of the core technologies will be developed, deployed and delivered by third parties in remote jurisdictions.

Local firms will have, as a practical matter, little to no influence on the core development of these technologies. As a jurisdiction we are likely to be a relatively modest global customer and the regulatory framework should accommodate this reality.

AI system developers are unlikely to accede to requests from firms in jurisdictions such as Australia to reveal training technologies and data sets used in the creation of AI agents.

The regulatory framework should aim to be supportive of remote development and delivery of AI systems of this type.

- c) emerging international approaches to mitigating AI risks;

International approaches are early in their development and deployment and none appear fully formed at this time.

While the EU has moved rapidly to implement its AI regulations, and has a reasonably comprehensive package of measures, there are concerns it may not be optimised for economic growth³.

³ <https://www.euronews.com/next/2024/03/22/could-the-new-eu-ai-act-stifle-genai-innovation-in-europe-a-new-study-says-it-could>

The US so far has presidential Executive Orders which endeavour to support innovation and support the development of safety and security standards, and some US states have implemented further measures to target various concerns including privacy⁴.

The UK has released a “pro-innovation” white paper on AI regulation. The paper supports using regulators’ domain specific knowledge to implement principles in a non-legislated way at first, before legislating at a later stage. The aim will be to create a ‘cross-cutting’ AI regulatory framework.

The principles are:

- Safety, security and robustness
- Appropriate transparency and explainability
- Fairness
- Accountability and governance
- Contestability and redress

While the thought leadership is welcomed and well-intentioned, there is no data yet on its impacts and outcomes.

Australia should benefit from the outcomes produced by these various regulatory experiments by being a follower in AI regulation except where specific concerns have been clearly identified and can be remedied by regulatory intervention.

- d) opportunities to adopt AI in ways that benefit citizens, the environment and/or economic growth, for example in health and climate management;

To benefit economic growth, excessive or untargeted regulation must be avoided. As a Copenhagen Economics study found about Generative AI in the EU:

Given that GenAI is still a nascent technology, it is also important to ensure that any regulation is fit for purpose and does not lead to a dampening of competitive conduct, including entry and expansion of new players. A survey of smaller firms (which are less capable of bearing the costs of regulation) found that regulation ranks second only to financing as the main barrier for EU startups in the AI segment more generally.⁵

- e) opportunities to foster a responsible AI industry in Australia;

Out of scope for AFMA’s submission.

⁴ <https://www.bclplaw.com/en-US/events-insights-news/2023-state-by-state-artificial-intelligence-legislation-snapshot.html>

⁵ <https://copenhageneconomics.com/wp-content/uploads/2024/03/Copenhagen-Economics-Generative-Artificial-Intelligence-The-Competitive-Landscape.pdf>

- f) potential threats to democracy and trust in institutions from generative AI; and

Out of scope for AFMA's submission.

- g) environmental impacts of AI technologies and opportunities for limiting and mitigating impacts.

AFMA holds that a holistic view of the environmental impacts is appropriate – for example efficiency gains for the economy increases the common wealth and thereby can free up capital for direct spending and investment in the reduction of the economy's environmental impacts in other areas. As is already the case for non-AI data centres, which are already significant consumers of energy, we expect the net gain from these efficiencies for the environment will be significant.