

28 July 2023

Lucas Rutherford
General Manager, Technology Strategy Branch
Department of Industry, Science, and Resources
Via webform.

Dear Mr Rutherford,

RE: Safe and Responsible AI in Australia Discussion Paper

The Financial Services Council (FSC) welcomes the opportunity to make a submission in relation to safe and responsible artificial intelligence (AI) in Australia.

The FSC is supportive of a risk-based approach to governing AI use in Australia with only the most obviously harmful AI products banned. Australia's approach to AI governance should be using a pro-AI approach which encourages innovation in Australia. This in turn will encourage more and more consumers to interact with the technology and build trust in its use cases.

A risk-based approach to the governance of AI will allow low-risk AI to be utilised by organisations in a relatively unfettered way, whilst requiring more of higher risk uses. What uses of AI are to be considered high and low risk should be the subject of a second, more fulsome consultation once the general approach is agreed following this consultation.

This approach will create a more competitive landscape for innovation where those with the appropriate resources to develop and manage high-risk use cases can do so with appropriate guard rails whilst those that employ low-risk AI solutions can do so without expending significant resources.

Further, a risk- and principles-based approach, taking the best of those implemented in the European Union and United Kingdom, should create certainty and encourage innovation within the AI landscape and should be scalable to adequately capture future innovations. The governance framework should also be consistent with existing regulatory interventions so as not to create confusion or duplicate regulation for entities that already have a strong governance framework, such as financial services.

Summary of Recommendations

1. Any governance framework should be considered holistically within the existing scheme of non-AI specific regulations to ensure any new arrangement plugs gaps where appropriate rather than duplicating or even contradicting existing requirements.
2. Australia's approach to AI regulation should be informed by a set of principles that focus on risk-based regulation, encouraging innovation, protecting consumers, creating certainty, and improving public trust in AI.
3. Any governance framework should be easy to comprehend and risk-based, with reference to both impact and likelihood, so that it does not create unnecessary red tape and a confusing overlap of regulations.
4. The FSC is supportive of a governance framework that creates trust in AI through transparency requirements such as disclosure obligations where decisions are being made by artificial intelligence. However, this should be appropriately balanced with intellectual property and competition considerations.

5. The Government could encourage innovation of trusted and safe AI through the provision of a voluntary regulatory sandbox in which organisations could test their products in a safe way before going to market. This would provide certainty that the product met the required regulatory standards and build trust amongst consumers that the product has been appropriately tested.
6. The Government should provide clarity to business that R&D tax incentives extend to the development of AI products and processes.
7. The FSC supports a risk-based approach to AI governance that creates certainty for innovation of the AI landscape.

About the Financial Services Council

The FSC is a peak body which sets mandatory Standards and develops policy for more than 100 member companies in one of Australia's largest industry sectors, financial services. Our Full Members represent Australia's retail and wholesale funds management businesses, superannuation funds, and financial advice licensees.

The financial services industry is responsible for investing more than \$3 trillion on behalf of over 15.6 million Australians. The pool of funds under management is larger than Australia's GDP and the capitalisation of the Australian Securities Exchange and is one of the largest pools of managed funds in the world.

AI in Financial Services

AI has myriad uses in the Australian financial services context and many organisations are considering ways to responsibly incorporate AI into their systems and processes. Use cases seen in other jurisdictions around the world include assessing financial risks, crafting investment portfolios, and reducing and preventing fraud and combating financial crime. Other, non-financial services specific uses include helping manage customer service for simple interactions and improving efficiencies within the workplace.

Equally, FSC members understand that AI presents a risk to their businesses, particularly in relation to fraud, scams, and cybersecurity. The sophistication of AI forgeries make it even harder for organisations like superannuation funds to check and confirm the identity of their members and FSC members are keen to see appropriate risk-based mitigation measures put in place in this regard.

FSC members are fully alive to the risks associated both internally and externally of using AI within their organisations. The FSC is therefore overall supportive of the work of government to provide certainty for industry with regard to AI use, encouraging innovation in the future and providing organisations with appropriate risk-based guidelines for the use of AI technology in Australia.

Potential Gaps In Approaches

Question 4. Do you have any suggestions on coordination of AI governance across government?

In addition to any potential artificial intelligence specific governance arrangements, it should be considered that financial services such as the provision of financial advice is already heavily regulated in Australia and there may be existing legislation and regulation that adequately captures the use of artificial intelligence in certain contexts.

It is recommended that any new governance arrangements work cohesively with existing regulatory frameworks to ensure that operators have certainty and there is no duplication or inconsistency of regulation.

In addition, and as acknowledged by the Discussion Paper, there is significant crossover with the work of other Government agencies such as work being undertaken on the Privacy Act and Digital Platforms, and even the substantial work occurring across government in relation to scams, fraud, and

cybersecurity. The FSC recommends government consider crafting an overarching strategy that fully considers the full gamut of the work being undertaken across government and the existing governance framework. The role of any further governance should be to plug gaps, where appropriate, with risk-based and scalable measures rather than creating a whole new standalone framework.

Given the scope of AI innovation, it should be acknowledged that full-scale legislative change of an emerging area is going to take time. The Government should be careful in its approach, ensuring that appropriate guardrails are put in place in a timely manner but also ensuring that regulation is fit for purpose and not rushed through out of necessity.

It should also carefully consider other incoming legislative arrangements that will have an impact on the AI landscape, such as the privacy review and the incoming Cybersecurity Strategy, which will have significant overlap with what is proposed in any AI regulatory scheme.

RECOMMENDATION 1

Any governance framework should be considered holistically within the existing scheme of non-AI specific regulations to ensure any new arrangement plugs gaps where appropriate rather than duplicating or even contradicting existing requirements.

Responses Suitable for Australia

Question 5. Are there any governance measures being taken or considered by other countries that are relevant, adaptable, and desirable in Australia?

Many jurisdictions around the world have approached AI regulation with a principles-based lens. The benefit of this approach is that it provides a guiding star for all preceding legislation and regulation that may follow. Several of these jurisdictions choose to craft these principles with a pro-AI framework, one that considers the innovation possible through AI whilst balancing the risks appropriately.

It should be noted that no jurisdiction is sufficiently far along in their regulation of AI to make reasonable judgments about the impact that individual regulatory schemes have had in relation to their intent. That said, there are positive aspects of various international approaches which can be drawn on to support the Australian approach.

The FSC is supportive of an approach to AI regulation that is risk-based and pro-innovation. The responsible use of AI technology necessitates guardrails that have reference to the risks of the activities involved and place appropriate mitigation measures in place to protect Australians, but still allows for the use of the technology in a way that is safe and responsible.

Conversely, the FSC is not supportive of measures that blanket ban the use and development of AI technologies, except in extreme circumstances, as this would only seek to hinder innovation and cost organisations from competing globally. A risk-based approach to AI regulation means that AI uses that are deemed the highest risk should have an appropriate level of regulatory oversight to mitigate those risks.

One of the most noteworthy approaches to AI regulation at the moment is that of the EU who are seeking to draft regulation that captures EU participants use of AI. While seen as prescriptive, the EUs approach steps up regulation commensurate with the risk of AI, with the risk categories defined within the Act. The benefit of the EU approach is that it provides certainty with regard to the requirements associated with putting a specific AI based product on the market. This gives both organisations a remit to create and market products within the regulatory scope they are most comfortable complying with, and provides a level of surety that their product will not be the subject of regulatory interventions after it goes to market. It also improves trust for consumers by providing a clear framework against which AI products have been tested.

A drawback of the EU model is that it places AI products into buckets based on risk but does not appear to give weight to the likelihood that that risk would eventuate. As risk is a product of impact and likelihood, it is important that both of these factors are weighed into the definition of AI risk categories.

Of note as well is the United Kingdom's (UK) approach to AI regulation. This approach seeks to ensure the UK remains an AI and science superpower and is underpinned by a series of cross-sectoral principles which guide the UK Government's AI regulation going forward. These principles ensure AI regulations are:

- Context-specific and based on use and impact;
- Pro-innovation and risk-based, focussing on high-risk concerns rather than hypothetical or low-risk concerns;
- Coherent and easy to navigate; and
- Proportionate and adaptable starting with lighter touch regulatory options in the first instance.¹

These principles have a focus on the benefits of AI whilst still maintaining consumer protections through a risk-based focus. It also ensures that innovation is encouraged and regulation scalable, creating certainty for operatives.

The Strategy's approach to governance of AI technology seeks to:

- Provide certainty for the UK AI ecosystem;
- Improve public trust in AI;
- Increase responsible innovation; and
- Maintain the UK's position as a global leader in AI.²

One of the criticisms of the UK's model is the requirement for individual agencies to enforce AI within their own patch. This creates the potential for under-regulation, over-regulation, and significant confusion amongst industry operators and consumers. The FSC is supportive of a co-ordinated approach to AI regulation, that does not leave individual agencies and regulators to create their own AI specific regulation. Instead, regulation should allow for AI to be captured by existing frameworks in a technology neutral way, with AI specific regulation only plugging gaps where necessary.

The FSC is also supportive of models that focus on building capability and consumer trust, thereby supporting innovation, and encouraging accountability and transparency. For example, the Singaporean National Artificial Intelligence Programme in Finance supports Singaporean financial institutions to research, develop, and deploy AI solutions.³ This principles-based approach has a focus on helping organisations develop their capabilities and utilise AI solutions in a safe and transparent way. The objectives of the Programme include:

- Increasing productivity through the adoption of AI;
- Creating new jobs through increased AI innovation activities and upskilling in AI-related competencies; and
- Improving societal acceptance of AI through sound AI governance.⁴

The FSC is supportive of a principles-based approach to AI regulation and is particularly supportive of principles that ensure that regulation provides room for innovation, is scalable with incoming innovations, and is risk-based. A regulatory framework that provides this certainty would also be cohesive and easy to navigate.

RECOMMENDATION 2

¹ The Office for Artificial Intelligence. (2022). *Establishing a pro-innovation approach to regulating AI*.

² The Office for Artificial Intelligence. (2021). *National AI Strategy p. 14*.

³ Monetary Authority of Singapore. (2021). *National programme to deepen AI capabilities in financial services*.

⁴ Ibid

Australia's approach to AI regulation should be informed by a set of principles that focus on risk based regulation, encouraging innovation, protecting consumers, creating certainty, and improving public trust in AI.

Target Areas

Question 8. In what circumstances are generic solutions to the risks of AI most valuable? And in what circumstances are technology-specific solutions better.

The FSC is supportive of an AI governance framework that is easy to comprehend, is risk based, and is scalable for future innovations in the sector.

As such, general governance solutions that provide definitional boundaries within which existing and future types of AI products can be captured, would provide more certainty for industry in developing these innovating use cases. These definitional boundaries would likely be based on risk, with reference to both impact and likelihood of occurrence.

The FSC acknowledges that industry specific codes and regulations may be required for certain use cases but again encourages government to consider the overall framework of proposed governance measures to ensure that multiple layers of governance do not create overlap or confusion.

RECOMMENDATION 3

Any governance framework should be easy to comprehend and risk-based, with reference to both impact and likelihood, so that it does not create unnecessary red tape and a confusing overlap of regulations.

Question 9a. Given the importance of transparency across the AI lifecycle, please share your thoughts on where and when transparency will be most critical and valuable to mitigate potential AI risks and to improve public trust and confidence in AI?

FSC members are supportive of a transparency framework that places appropriate risk-based guardrails for the provision of notices and even explanations where they are warranted, especially in relation to automated decision-making processes. However, these transparency requirements should also be considered with a lens of competition and ensuring that commercial products and intellectual property (IP) are appropriately protected.

Building trustworthy AI will undoubtedly create a safe and innovation-friendly environment for users, developers and deployers. However, increasing trust in AI powered products and services requires a high level of trust from consumers however, equally, companies adopting AI technologies must be confident that their competitiveness is maintained.

From an organisation perspective, this has two lenses, protecting organisational IP but also understanding what goes into training the AI products that they are building on.

The creators of AI models should be required to disclose important information, such as training models, so that users understand fully any potential biases, shortcomings, or issues of interpretation that the AI may present. This will also help non-organisational consumers to understand the underlying assumptions and AI may make.

For products that use AI, either built on an open API, such as ChatGPT, or a bespoke system, there are considerations that need to be made to protect intellectual property rights and shield organisations from anti-competitive behaviour. Striking the right balance is important.

As such, appropriate guidelines related to, for example, automated decision-making processes should be crafted with industry consultation. These guidelines should focus on providing certainty for customers about when decisions are being made using an AI technology and provide appropriate expectations about how and when decisions can be reviewed by a human. However, these guidelines

should be scalable to different levels of risk and not ultimately circumvent the efficiencies that artificial intelligence is adopted to create.

These guidelines could also consider matters such as prompting of AI. Outside of the information used to train the AI, much of the output of AI is dependent on the input given to prompt the program. The way that a prompt is formed speaks to the quality of the output, as well as bringing out further internal biases. A prompt may be framed negatively and therefore provide a negative outcome. Organisations could be required to have clear internal policies, based on national guidelines about the use of AI in their products, including the way in which they prompt their AI solutions when making decisions. These could be required to be disclosed in higher-risk operations such as automated decision making.

Further, transparency of AI should not necessarily extend to having to disclose proprietary algorithms or other matters of a commercially sensitive nature.

This should not, however, absolve organisations from responsibility for their AI based products. There should be a positive duty on producers of AI based products to be able to explain adequately, when queried, how an AI product came to a certain conclusion. This may be with reference to the training model, or through other means. In order to produce AI based products, organisations should be able to prove that they have sufficient understanding and control of the AI to justify its presence in the market.

RECOMMENDATION 4

The FSC is supportive of a governance framework that creates trust in AI through transparency requirements such as disclosure obligations where decisions are being made by artificial intelligence. However, this should be appropriately balanced with intellectual property and competition considerations.

Building Trust

Question 11. What initiatives or government action can increase public trust in AI deployment to encourage more people to use AI?

One initiative that would encourage safe innovation is the provision of certainty through a voluntary regulatory sandbox for AI development. This would allow organisations to test AI in a safe and controlled environment, with the assistance of the appropriate regulators, to ensure that AI met the requirements of the risk-based regulatory framework before it went to market.

This would serve the dual purpose of encouraging innovation and building trust by ensuring that AI based products safely met the requirements of the legislative framework. Consumers could interact with the AI based products with confidence, knowing that they were appropriately tested under the supervision of appropriate experts. This process would also provide certainty to organisations building the products that they met legislative standards.

RECOMMENDATION 5

The Government could encourage innovation of trusted and safe AI through the provision of a voluntary regulatory sandbox in which organisations could test their products in a safe way before going to market. This would provide certainty that the product met the required regulatory standards and build trust amongst consumers that the product has been appropriately tested.

Another approach to encouraging innovation would be through the provision of explicit research and development tax incentives for AI uses.

As acknowledged by the Department of Industry, the R&D tax incentives guard against an organisations hesitation to develop new technology because it may not result in significant ROI and

further encourages innovation by guarding against some of the losses that the spillover of their efforts will have on their competitors.

Providing clarity that tax incentives extend to AI uses would encourage organisations to develop and innovate within the AI space.

RECOMMENDATION 6

The Government should provide clarity to business that R&D tax incentives extend to the development of AI products and processes.

Risk Based Approaches

Question 14. Do you support a risk-based approach for addressing potential AI risks?

The FSC is supportive of a risk-based approach to governing AI use in Australia. Any alternative governance models could create significant overreach and red tape for lower risk users, creating an uneven playing field for potential innovation where only the most well-resourced users can afford to navigate the governance framework and develop AI technologies.

RECOMMENDATION 7

The FSC supports a risk-based approach to AI governance that creates certainty for innovation of the AI landscape.

Question 15. What do you see as the main benefits or limitations of a risk-based approach? How can any limitations be overcome?

A risk-based approach will encourage innovation and provide an appropriate level of oversight for higher risk activities whilst allowing lower risk activities to flourish without excessive red tape. An understanding that not all AI uses carry significant potentials for harm will also allow consumers to become more comfortable with the use of the technology leading to increased trust and further stimulate innovation.

When considering a risk-based approach, it is important too to consider the usability of the framework to ensure that both users and developers of the technology are provided with as much certainty as possible in relation to the AI product being used or created. A strong governance framework that is scalable to future development would ensure that this certainty is achieved.

From a limitations perspective, it is important to acknowledge that AI technology is changing and updating daily. Any regulatory approach will, to some extent, forever be playing catch up with evolving technology. While innovation should be encouraged, regulation should be flexible enough to adequately capture potential new uses of AI into the future, so that constant regulatory updates are not required.

Question 17. What elements should be in a risk-based approach for addressing potential AI risks? Do you support the elements presented in Attachment C?

The FSC's position in relation to the elements presented in Attachment C of the Discussion Paper are outlined below.

<i>Possible Elements</i>	<i>FSC Response</i>
Impact assessments	The FSC is supportive of impact statements with publishing required for medium- to higher-risk uses. Where an AI technology clearly has a low-risk, requiring significant resources, such as the preparation of an individualised impact assessment, may create an unreasonable barrier to entry.

<i>Possible Elements</i>	<i>FSC Response</i>
Notices	The FSC is supportive of consumers being required to be notified of when and how AI might be used to make decisions and any recourse customers might have for human intervention. This should have reference to the level of risk associated with the specific activity. <i>See response to question 9a above.</i>
Human in the loop/oversight assessments	The FSC is supportive of establishing guidelines regarding when human oversight is required. These should have reference to the risk of the activity being undertaken. <i>See response to question 9a above.</i>
Explanations	The FSC is supportive of explanations being required for consumers where AI has been used to make a decision and an unfavourable outcome has been reached. <i>See response to question 9a above.</i>
Training	The FSC is supportive of guidelines that help employers work with their employees to understand and better use AI technology to create efficiencies in their workplace. The FSC does not believe mandated training is appropriate for all employees and individual employers should be able to create their own framework with reference to the guidelines published by the Government. The FSC is supportive of developers being required to be suitably qualified.
Monitoring and Documentation	The FSC is supportive of minimum requirements for monitoring and documentation of AI uses within organisations for medium- and higher- risk uses.

Conclusion

The FSC is supportive of a risk-based approach to AI regulation. This risk-based approach should consider key principles such as encouraging innovation in the AI landscape as well as ensuring consumers are adequately protected.

If you would like to discuss anything contained in this submission, please do not hesitate to contact me.

Yours sincerely,

Kirsten Samuels
Policy Manager, Superannuation and Innovation