27 July 2023

By Upload.

Dear Sir/Madam

Re: Safe and responsible AI in Australia

AFMA welcomes the opportunity to provide comment on the Department of Industry, Science and Resource's Discussion Paper *Safe and responsible AI in Australia*.

Artificial Intelligence (AI) is currently developing at an extremely fast pace. For certain applications outside of AFMA's remit and largely in the consumer space there may be potential for regulatory limitations to avoid harms, particularly to the most vulnerable.

In the wholesale markets, and in more generally in financial markets, however, we suggest a more cautious approach to regulation that moves slowly, leverages international regulatory outcomes, and is careful to preserve the potential efficiency gains that AI is enabling.

The risks of early regulation in such a fast-moving area are skewed much more towards creating impediments to industry and innovation rather than being empowering of them. We note in our response below some of the proposals, although nominally risk based, appear unnecessarily burdensome. Experience has shown that even where regulatory initiatives intend to support innovation, such as sand box schemes, these typically have marginal utility.

We trust our response is of assistance and would be pleased to continue to contribute to the development of policy in this area.

Yours sincerely

Damian Jeffree
Senior Director of Policy

Australian Financial Markets Association
ABN 69 793 968 987
Level 25, Angel Place, 123 Pitt Street GPO Box 3655 Sydney NSW 2001
Tel: +612 9776 7907 Email: secretariat@afma.com.au Web: www.afma.com.au

Definitions

1. Do you agree with the definitions in this discussion paper? If not, what definitions do you prefer and why?

AFMA supports the use of ISO definitions. The use of standard international definitions will assist with integration and the establishment of common approaches to AI and its regulation, an outcome that is efficient for business.

Potential gaps in approaches

2. What potential risks from AI are not covered by Australia's existing regulatory approaches? Do you have suggestions for possible regulatory action to mitigate these risks?

Most of the risks of AI are already be addressed through the principles-based sector specific regulation, and the general regulations. These channels can leverage established processes in managing new and emerging risks through the publication of guidelines etc. The potential for scalability of AI should not be assumed to create additional issues. In most cases it will be the outcomes that count, and these are typically well covered by the combination of existing specific and general regulations such as the Privacy Act (currently under review).

For financial advice, fiduciary or best interest duties provide sufficient safeguards to address ethical concerns raised in the paper. The current regulatory structures are outcomes-based and provide good coverage of relevant issues. Limited roll out of AI approaches in advice and industry feedback suggest the regulations may not be sufficiently accommodative of automated solutions at present. We expect transparency practices around such applications will evolve over time. We note that financial advice related rules are currently being reviewed by Treasury.

AFMA sees value in the use of industry standards as a sensible first step in emerging areas. Industry standards are more likely to keep pace with the rapid developments in AI. Tailored technical standards can be built for specific AI use case as AI evolves. Where consensus builds and the benefits outweigh the costs, in some circumstances it may be appropriate to move some standards into regulation. This regulation particularly any legislated elements should be principles-based.

In wholesale markets we are not currently aware of any regulatory gaps for AI as outcomes are covered by existing legislation.

3. Are there any further non-regulatory initiatives the Australian Government could implement to support responsible AI practices in Australia? Please describe these and their benefits or impacts.

Given the breadth of laws that already potentially impact AI there is likely benefit in relevant industry regulators supporting increased awareness of the existing regulatory limitations on AI.

Support by regulators for awareness of international industry standards is also likely to be of benefit increased conformance with these standards.

The Government should continue work to understand the risks posed by developments in AI and assist government agencies who are regulating specific sectors respond under existing regulatory frameworks in ways consistent with national principles.

4. Do you have suggestions on coordination of AI governance across government? Please outline the goals that any coordination mechanisms could achieve and how they could influence the development and uptake of AI in Australia.

We see some parallels between the coordination of AI governance and the regulation of cyber security.

An opportunity was lost in relation to cyber security when multiple regulators created their own nationally and internationally inconsistent security standards that applied different levels of security for the similar data.

AFMA suggests that every effort should be made to avoid a similar outcome in relation to AI.

The Government should align programs and strategies with OECD AI Strategies and other international approaches to facilitate consistency globally.

Regulators should be discouraged from drafting their own unique rules. Where sector specific rules are necessary these should be minimized and draw on common principles set at the national level, and where at all possible be internationally consistent.

We see value in consolidating some of the AI specific elements of law and maintaining a guide for implementors of AI systems to assist them ensure their implementations are compliant with all relevant requirements.

Where AI specific laws are implemented, they should be consistent with global standards as they emerge. Regulators can then adapt these laws to industry specific cases with their guidance.

Support should also be given for pilot and trial programs.

Responses suitable for Australia

5. Are there any governance measures being taken or considered by other countries (including any not discussed in this paper) that are relevant, adaptable and desirable for Australia?

At a general level AFMA suggests Australia should leverage the experience (as it emerges) of foreign jurisdictions in respect of their approaches to AI governance. We expect there will be the opportunity

to see how different approaches function in jurisdictions on the technological leading edge before considering what might be optimal locally.

We expect sensible regulatory approaches to emerge over time as the risks are better understood.

As an immediate step we suggest increasing awareness of voluntary standards such as NIST's AI Risk Management Framework. Common approaches with leading jurisdictions are likely to result in lower costs and increased international compatibility.

Target areas

6. Should different approaches apply to public and private sector use of AI technologies? If so, how should the approaches differ?

AFMA supports a risk-based approach for both public and private sector use. While there are many commonalities in the risks of public and private sector use of AI technologies there can be differences. For example:

- Large data sets of sensitive information can increase risks that AI will train on private information and risk its release or misuse.
- Public sector actors and, more generally, governments may have more coercive powers over firms and individuals in their jurisdictions than private sector actors. This greater power can increase the risks of adverse outcomes for individuals and firms.
- Consistent with other limitations sovereign governments place on themselves and their agencies, decision making (and potentially other activities) by, or assisted by, AI in Government as an exercise of state power must be bound by additional protections designed to limit the potential for abuse of power. For example, AI should be bound by the model litigant rules the Government imposes on its agencies.
- In some cases, private sector actors will have good commercial reasons for distinguishing between individuals or firms. Governments and public sector actors will less often have such commercial reasons.

A risk-based approach that focuses on outcomes is suitable regardless of public or private usage. Where risks of harm are greater, and this may be the case for some of the reasons outlined above, then a more cautious approach may be appropriate.

7. How can the Australian Government further support responsible AI practices in its own agencies?

To support fair use of AI in its interactions with firms and individuals we suggest the Government look to the NSW Government's Artificial intelligence assurance framework.

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

As a general finding AFMA's experience is that technology specific regulations are rarely optimal as they can limit innovation, whereas principles-based regulation, typically focussed on outcomes, allows greater flexibility for firms to innovate and provide services efficiently and at the lowest possible cost.

9. Given the importance of transparency across the AI lifecycle, please share your thoughts on:
    a. where and when transparency will be most critical and valuable to mitigate potential AI risks and to improve public trust and confidence in AI?
    b. mandating transparency requirements across the private and public sectors, including how these requirements could be implemented.

In a wholesale context transparency is not the same issue as it might be in retail. In a retail context transparency can, for example, be appropriate to ensure that retail investors are aware of the source of recommendations. Transparency should only ever be required where warranted by the risk profile.

We reject the view that transparency is an unalloyed good as suggested by Questions 9 (a) and 9 (b) and we caution against a default mode of regulating for transparency.

Allowing firms to maintain intellectual property as confidential assists the commercial development of new technologies including AI.

This early in the development of a new field of technology such as AI the regulatory approach should be more cautious than suggested by the discussion paper lest innovation be stifled, and benefits denied.

10. Do you have suggestions for:
    a. Whether any high-risk AI applications or technologies should be banned completely?

In areas within AFMA's purview we currently do not see the need for any complete bans.

We expect that outside AFMA's purview there may well be areas that governments might determine bans are appropriate based on their prioritisation of matters such as privacy, safety, governance and other concerns.

    b. Criteria or requirements to identify AI applications or technologies that should be banned, and in which contexts?

There may be many disparate areas where AI applications or technologies should be banned, but the reasons for these bans are likely to vary significantly. They might include reasons of privacy, child protection, national security, the protection of political discourse, etc. As such it is difficult to generalize about why certain applications or technologies should be banned.

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

Usage of AI will increase as products and services become more available that use AI.

Government action might therefore be best focused on assisting firms utilize AI.

Implications and infrastructure

12. How would banning high-risk activities (like social scoring or facial recognition technology in certain circumstances) impact Australia's tech sector and our trade and exports with other countries?

We do not see a need for bans within AFMA's industry areas.

13. What changes (if any) to Australian conformity infrastructure might be required to support assurance processes to mitigate against potential AI risks?

AFMA views this area as highly technology dependent. Explainability, of the type we have come to expect from machine-based decision-making processes may have to evolve.

Neural network-based systems work on similar principles to human brains. As such we might reasonably expect similar arrangements to ensure conformance.

With parallels to the requirements placed on humans AI might require exposure to standardised data sets (the machine equivalent to training courses) and testing on outcomes (the machine qualifications equivalent).

Risk-based approaches

14. Do you support a risk-based approach for addressing potential AI risks? If not, is there a better approach?

AFMA supports a clear separation between wholesale and retail regulation of AI. Retail (including consumer and retail investor) use of AI has a different risk profile compared to wholesale uses. There are fewer risks involved in firms using AI to scan data for investment opportunities for example, than interacting with individuals via online chat bots.

AFMA supports risk-based approaches to AI, but at this stage in the development of the technology it is not yet appropriate to be prescriptive except in the most sensitive cases.

For example, we note that the table provide in the paper suggests that "Users must be trained" where the use case is "Use of AI in computer chess systems" or "AI-enabled chatbots". We think this is likely excessive. Similarly, the table suggests that for "Use of AI in safety-related car components" that it "Must have meaningful human intervention at specific points and final decision made by human/s" is more likely to *decrease* safety than *increase* safety in many circumstances.

A risk-based approach will evolve over time but is unlikely to be well-tuned ahead of time by regulatory preemption.

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 when implemented properly will provide the best cost to benefit ratio. The limitations are in the accurate a priori assessment of risk. A posteriori data will improve these estimates over time.

16. Is a risk-based approach better suited to some sectors, AI applications or organisations than others based on organisation size, AI maturity and resources?

AFMA holds that a risk-based approach is generally better. Where stakes are higher as with other areas of regulation there is likely more of a case to ensure risk assessments are accurate and up to date.

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

While we support the elements listed in Attachment C as potentially contributing to a sensible risk-based approach, as noted above, we have concerns around the proposed mandated application proposed in Box 4.

At this early stage we suggest working with industry to help voluntary risk-based approaches emerge, leveraging off international experience.

18. How can an AI risk-based approach be incorporated into existing assessment frameworks (like privacy) or risk management processes to streamline and reduce potential duplication?

This is an important but complex question. Where possible duplication of similar principles across multiple regulatory schemes should be avoided in preference of a single regulatory approach. Supporting industry assessments into existing regulatory requirements in a consistent manner should also contribute to achieving a streamlined outcome.

19. How might a risk-based approach apply to general purpose AI systems, such as large language models (LLMs) or multimodal foundation models (MFMs)?

This will go to the data on which the model is trained and the purpose for which it is deployed and the associated risks. An LLM/MFM that assists customers choose paint colours that is trained mainly on colour trends should have no or very little requirements applied to it.

20. Should a risk-based approach for responsible AI be a voluntary or self-regulation tool or be mandated through regulation? And should it apply to:
    a. public or private organisations or both?
    developers or deployers or both?

For private organisations in AFMA's industry sectors voluntary tools will generally be preferable to preserve innovation and in AFMA's areas of industry at present we believe that this is all that is required at present.