InsurtechAustralia
INSURANCE INNOVATION + COLLABORATION

Discussion Paper: Safe and responsible AI in Australia

Submission to Department of Industry, Science and Resources

Insurtech Australia for and on behalf of its members

About this Submission

Insurtech Australia (**Insurtech**, **we**, **us**, **our**) welcomes the opportunity to engage with the Department of Industry, Science and Resources (**Government**) to lodge a submission on its Discussion Paper, 'Safe and responsible AI in Australia' released on 1 June 2023 (**Paper**).

This submission was drafted by Insurtech in consultation with our members (each, a **Member**). In developing this submission, our Members participated in a roundtable session to discuss their attitudes towards Artificial Intelligence (**AI**) generally, the extent of their adoption of AI in their respective organisations, as well as their collective view on the preferred approach to AI regulation in Australia.

We acknowledge the support and contribution of law firm DLA Piper in respect of the development of this submission and their facilitation of the roundtable discussion between us and our Members.

About Insurtech Australia

Insurtech is Australia's leading not-for-profit industry association for insurance technology (**insurtech**) and insurance innovation. Our mission is to make Australia a world leader in insurtech and insurance innovation by supporting and growing the Australian insurtech community, including insurtech startups, insurers, hubs, accelerators and investors, and advocating on behalf of our Members.

Executive Summary

Insurtech considers that AI is playing, and will continue to play, a key role in facilitating innovation in the insurance and insurtech sectors, through enhancing the efficiency of internal and organisational processes and assisting in the development of more sophisticated and efficient commercial offerings to better meet customer needs.

Already, Members are deploying (and exploring the potential deployment of) AI across various business functions, including organisational risk assessment, fraud detection, pricing, policy generation, claims management and customer experience, using various forms of AI.

More broadly, Insurtech considers that AI will play a key role in growing Australia's digital economy and sees strong potential in AI becoming a driver of innovation and economic growth in Australia generally.

Preferred regulatory approach

Insurtech believes that a decentralised, sector-based and principles-based approach to the regulation of AI would best serve the needs of both public and private sector stakeholders in Australia, enabling the regulatory approach in different sectors to be tailored to the specific risks and use cases unique to those sectors, thereby allowing for more nuanced and targeted regulation. Further, we consider that a principles-based approach that is not overly prescriptive, will allow a more dynamic regulatory approach that caters to the rapidly changing nature of AI.

Alternatively, a centralised, 'generic' or risk-based approach to regulation (an approach that has been taken in jurisdictions such as the EU and Canada), while providing consistency in the regulatory treatment of AI across sectors, could create a rigid and onerous compliance environment that could stifle innovation. We think this would be an undesirable outcome for Australia, as an innovation-based economy that wishes to become a global technology powerhouse.

Further, as some existing legislation may already apply to AI, we consider that adding additional 'overarching' legislation that is 'sector-agnostic' could introduce undue complexity where legislative regimes overlap and could lead to an increased compliance burden for organisations, and an opacity as to best practice.

In order to implement a decentralised approach to AI regulation, we consider that Government should first conduct a gap analysis of existing regulation, in order to identify where existing frameworks need to be updated to more-effectively contemplate AI.

Regardless of which legislative approach is taken, we think that collaboration between the public and private sectors will be crucial to creating an effective regulatory regime. Regulations (especially when it comes to innovation) cannot be created in a vacuum, and given the rapid development of AI technology, any guardrails established must be informed by current domain expertise. As such, we think that it is imperative to establish dialogue between industry and government so that relevant voices can be heard.

Guidance and education

As noted in the Paper, one of the key impediments toward the adoption of AI in Australia has been a lack of public trust and confidence in AI technologies – we think that this is primarily due to a lack of knowledge about the technology itself, as well as a general focus on the potential complications of using it, particularly with respect to applications that have the potential to affect the lives of Australians.

In addition, we think that there is a significant knowledge gap in Australian industry and the general public, as to how existing legal frameworks may be applied to AI (for example, in relation to data security and privacy, discrimination, consumer protection etc). We consider that these issues may be mitigated or resolved through the development of public and industry education and guidance about the technology itself, its benefits and challenges, as well as how the existing legal framework caters to it (and in particular, how consumer rights are protected in connection with their interaction with AI).

To that end, we support the development of practical guidance and educational initiatives surrounding the use of AI to educate the public, generate trust and confidence in AI and provide industry and consumers with an understanding of how the law applies.

We think that this would not only increase AI adoption, but would also be conducive towards encouraging the legal and responsible development and deployment of AI.

Detailed Responses

Section 1: Definitions

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

While Insurtech does not object to the definitions used in the Paper *per se*, we consider that the definitions are, in some cases, narrow and do not adequately capture key features of AI. If it is Government's intent to create definitions that capture all-manner of AI-driven technologies, then we consider that more work will need to be done to develop definitions that are fit to form the basis of any AI regulation.

Further, given that AI and its application is constantly evolving, we think that any definitions that incorporate examples of AI output (for example, the definition of 'Generative AI models') should be clear that the examples are non-exhaustive.

We also consider that any definitions that are adopted should be consistent with the equivalent definitions used by our key trade partners, in order to achieve, as far as is possible, harmonisation across regulatory regimes and to allow insurance and insurtech businesses to expand and operate in such jurisdictions without the need to recalibrate its regulatory approaches.

Definition of AI

We think that the proposed definition of AI omits important references to some inherent qualities of certain subsets of AI, including 'autonomy' and ability to self-calibrate and learn. While these qualities do not apply to certain types of AI, such as Large Language Models (**LLMs**), which are static models that generally do not learn beyond their base training data, most definitions of AI used in the insurtech space refer to mimicry of human cognitive function and accordingly, we think that consideration of these capabilities is an important aspect of any definition of AI.

Further, we think that these aspects are important to specifically reference, as, if these aspects are not contemplated, this raises the risk that the definition will capture traditional mathematical and statistical models (for example, logistic regression) that have been used for decades in the banking and insurance sectors, which we assume is not the intent.

Additionally, the proposed definition of AI contemplates that all outputs are 'predictive', which is not necessarily the case for certain applications that do not predict anything as such, but rather, create connections between data points in order to generate output based on input.

We think that a better approach would be for the definition of AI to refer to 'inference' made by AI (similarly to the definition used in the AI Act in the EU), as opposed to 'prediction', as functionally, AI draws inferences based on thematic connections between different data points in its training data sets.

In addition, the examples of outputs provided in the definition of AI do not include 'actions' which could therefore exclude, by definition, AI-enabled robotics.

Definition of Machine Learning

We consider that the proposed definition of Machine Learning does not contain the appropriate level of detail for it to form the basis of any sort of regulatory definition, as it does not contemplate the process by which the relevant algorithm derives its learnings, and only seems to contemplate the results or learnings themselves.

Definition of Generative AI

Similarly, the proposed definition of Generative AI models does not make any reference to how the relevant model is trained (the input and processing of training content and data sets); rather, it merely refers to its function of generating novel content.

Section 2: Potential Gaps in Approaches

Question 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?

The below response is not intended to be an exhaustive recitation of all gaps in the current regulatory approach. Rather, the below sets out the key risks relevant to the insurance and insurtech sectors, from Insurtech's and our Members' perspective.

AI Ethics Principles

The regulatory regime in Australia, as it relates to AI, includes the AI Ethics Principles, and while these principles may address some or all of the risks posed by AI-driven technologies (including those considered below), they are voluntary principles that do not carry the force of law. Accordingly, we have not referenced those principles in answering this question, but instead have focussed our answer more on the enforceable steps that Government could take to mitigate these risks (which, for completeness, could include making the AI Ethics Principles enforceable in respect of the development, deployment and use of AI).

Development risk

While some of the potential risks from AI are addressed to some extent by existing laws on privacy, consumer protection and anti-discrimination, and may be useful in providing consumers with remedy if those laws are infringed as a result of a consumer's interaction with AI (for example, where personal information is used to train an algorithm without the relevant individual's consent, where required, or where interaction with an AI-driven automated decision making (**ADM**) application produces a discriminatory outcome), we consider that the fundamental gap is that there are no laws which specifically regulate how AI is developed and used (arguably with the exception of the Consumer Guarantees under the Australian Consumer Law).

Accordingly, it appears that the current regulatory framework as it relates to AI can only help consumers after harm has been suffered and does not take a preventative approach to consumer harm by regulating the development and use of AI in the first instance.

Government could choose to mitigate this risk by:

- introducing laws or regulations (including mandatory standard and/or codes of conduct) that place guardrails around the development of AI, based on the potential for the relevant use case to cause harm (as per the EU approach); or
- mandating transparency requirements for AI products that are consumer-facing, not only in respect of the manner in which the relevant application is used, but also how it is developed and how it operates. This could enable consumers to make informed choices about their interaction with AI. For example, requirements to disclose the use of AI at any stage of the customer journey (which may also include a requirement to provide consumers with an ability to opt-out of AI use) or a right for consumers to request such transparency.

However, we do recognise the difficulties in mandating transparency requirements in respect of certain types of AI application that are trained by a 'black box' method that makes it very difficult to provide transparency on, for example, how the relevant application learns and/or makes decisions. We also recognise that for this to be an effective means of regulation, there would need to be associated accountability, security and privacy requirements (see below).

In traditional modelling, where one can articulate the causal relationships between data points that contribute to a particular outcome or decision, transparency is easy to provide. In the AI context however, where an algorithm operates autonomously and human oversight over the 'thought process' is minimal, transparency may be more difficult to provide.

Because AI models cannot think in terms of the causal link between data sets (one may be able to explain how a model works (e.g., it was trained using a neural network algorithm), but may not be able to explain, for instance, the weighting given to certain data points in decision making, or even the depth

of the examination of those data points by the algorithm), this makes it hard to provide transparency in respect of the workings of the AI model, how it makes decisions or how it produces output.

Ultimately however, we consider that while transparency requirements in some cases may be sensible and practicable, legislating in respect of the development of AI itself would likely stifle innovation and compromise the ability of developers and product offerors to freely-develop AI-driven products and accordingly, we do not endorse this approach. Further, if a 'blanket' approach to AI regulation is taken, we think such an approach would likely be inadequate in keeping pace with the development of the technology.

Data risk

Given the way training data for AI algorithms is generally procured (for example, by data scraping), there is a risk pertaining to the completeness, accuracy, currency and overall quality of that data, which creates a knock-on risk that outputs from AI models could be deficient, discriminatory, biased or otherwise incorrect, and therefore, unreliable. While the negative effects of sub-standard data quality may be addressed to an extent by existing law (for example, discrimination laws may provide a consumer with remedy where they have been discriminated against as a result of an AI decision that was made by an algorithm that had been trained on low quality data), we think that existing regulation does not cover-off this risk sufficiently, as there are no laws that require the quality of data used in training an AI algorithm to be validated.

While this regulatory gap may be mitigated by mandating quality assurance processes for training data sets being used to train AI algorithms, this may not be a reasonable or practical solution, especially considering the often immense quantity of data used to train an algorithm. Further, such requirements would require a level of human oversight which may devalue the use of AI-driven automation in the first place.

Further, in respect of generative AI specifically, given that many platforms that host generative AI models require the power of graphics processing units to enable their functionality, and because most of these capabilities are in the USA or the EU, this creates a data sovereignty risk in respect of the way in which data would need to be transferred and housed outside Australia to allow the relevant platform to operate. In respect of this risk however, we consider that, as it relates to consumer protections, this risk may be adequately mitigated under the existing privacy regime in Australia (see below).

Privacy risk

While the *Privacy Act 1988* (Cth) and the Australian Privacy Principles (together, the **Privacy Laws**) regulate the collection, use and disclosure of personal information, the existing regime does not place any specific limitation on the extent to which personal information may be used to train AI algorithms, or that determine the ingestion of such personal information to be a 'disclosure' for the purposes of the Privacy Laws (for example, if an AI model ingests data which is then housed on an external server), noting however, that the Privacy Laws otherwise apply to these issues, particularly in respect of the indirect collection and the notification and retention of personal information.

To mitigate this risk, Government could consider introducing specific regulation in relation to how personal information may be collected and used in the AI context (and relevantly, the means of collection, how long it may be retained for and if/when it needs to be deleted or destroyed). This could be through codifying existing determinations and guidance from the Office of the Australian Information Commissioner (**OAIC**) into the Privacy Laws. Examples of such regulation could include:

- mandating the use of synthetic data (derived from personal information and other sources) to train AI algorithms – while the problem still exists that the generative AI model needs to ingest the personal information to create the synthetic data, the personal information itself would not be used in the same way; and
- implementing regulation to limit or prohibit data scraping of personal information for the purposes of training AI algorithms.

While these would seem to be sensible regulatory measures to protect the integrity of consumers' personal information, we think that they should be implemented in a way that does not stifle innovation. Further, we consider that any such measures should be completed by the development and dissemination of clear guidance from the OAIC as to the application of the Privacy Laws to AI.

Security risk

As has always been the case, the sophistication of cybercrime runs parallel with development in technology – AI is no different. Malicious actors may be able to, for example, remotely alter the parameters of algorithms or corrupt training data for criminal or other nefarious purposes. Accordingly, for organisations that develop or use AI, this is a real risk, and it is a sensible approach that organisations invest in security solutions and threat detection and response infrastructure.

In order to curb this threat and incentivise industry to continue to innovate, we think that Government should consider amending tax and R&D laws to allow for concessions for organisations investing in security solutions and infrastructure. Further, we think that there should be clarity over recommended security protocols through the Australian Cyber Security Centre and the development of guidance for organisations.

Question 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.

We think that the implementation of non-regulatory initiatives is imperative to help boost consumer trust and confidence in AI by framing consumer perception of AI as a technological proliferation that is designed to improve the customer experience and help grow the Australian economy through innovation.

Such initiatives could include:

- public education and guidance, media involvement and consultation. For example, the development of a public AI toolkit, as has been implemented in Singapore;
- the development of opt-in-based industry codes (and a public register of organisations that have signed up the codes), which would not only be dynamic enough to adapt to evolving use cases, but would be developed on a sector-specific basis. While Australia already has the AI Ethics Principles for the development of AI, this is a voluntary framework and we think that such codes should be enforceable for those who opt-in; and
- the creation of an industry-agnostic AI body that acts as a conduit between government and industry in respect of how AI is developed, used and regulated (see also our response to Question 4 below).

We think that it is imperative that Government canvasses feedback broadly from industry players of all sizes, to ensure that non-regulatory initiatives and materials are fit-for-purpose and speak to the use cases in the relevant industry sectors.

Question 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.

Fundamentally, we do not see it as the role of Government exclusively to coordinate thinking on AI governance and regulation, and we consider that this is a role that will be best played by industry to inform how AI should be regulated based on the risks and use cases peculiar to particular industry sectors. Of course, open and transparent dialogue with Government and law makers is imperative, and a governance regime as between industry and government should certainly be established, but industry will hold the most informed view as to how the regulatory framework should be shaped to contemplate AI, while ensuring that innovation is encouraged and facilitated.

To assist this, an industry-led governance body could be established to help support a coordinated and coherent response to AI governance in Australia as between government and industry. However, we question the extent to which such a body would have the capacity and capability to speak on behalf of the entire business community and accordingly, we consider that sector/industry groups and regulators will still need to be heavily involved in coordinating AI governance efforts.

Section 3: Responses suitable for Australia

Question 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?

UK's decentralised approach

We think that the decentralised, sector-specific approach to AI regulation proposed by the UK is favourable in the Australian context, as it allows specific sectors to regulate AI based on the unique risks and use cases in the relevant sector. While this could mean that a regulator could choose to take a 'hard line' on AI regulation in respect of its sector's particular risks, we think it is also likely that the relevant regulator would wish to foster innovation in the relevant sector as much as is possible.

The key difficulty with this approach, however, is the question as to how to harmonise the regulatory treatment of general-purpose AI that is used in several different sectors.

EU risk-based approach

With regards to the EU's proposed risk-based, legislative approach, we think that there is merit from a consumer protection perspective in giving specific high-risk applications, such as social scoring mechanisms and real-time biometric and emotional monitoring, specific treatment (including, where appropriate, complete prohibition). However, generally, we are not in favour of this approach as we consider that an overarching legislative regime would be too 'rigid' and would likely stifle innovation.

Guidance and tools

It would be useful for industry and the public to have access to guidance and practical tools (along the lines of Singapore's 'AI Verify' toolkit) to assist in navigating the development of AI technologies.

Practical tools that can assist organisations in making relevant and informed decisions about procuring, developing and deploying AI systems (and to keep up with developments in their industry) would be particularly relevant and useful in the context of the insurtech industry, as well as in the insurance sector more broadly.

Public registers

The Amsterdam government has adopted a searchable public AI register (**Amsterdam Register**), which provides information on the algorithmic systems it uses in public service, the data sets used, the manner in which that data is processed, the details of the human oversight and risk management frameworks employed, and even provides an explanation of the potential for the relevant application to create discriminatory outcomes.

While we consider that this is a desirable initiative for the public sector, we do not see this as being appropriate for industry.

Section 4: Target Areas

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

We think that public sector uses of AI technologies should be held to the same, if not a higher standard than those of the private sector, particularly as a certain level of transparency is generally expected by governments. Further, transparency by the public sector, we think, would be impactful in improving public trust and confidence in AI generally.

If Government wished to implement initiatives to regulate public sector use of AI, then it could:

- take a similar approach to the UK, whereby public sector entities are required to adhere to a 'Algorithmic Transparency Recording Standard' (**Transparency Standard**) that requires public sector entities to provide information about the algorithmic tools they use, and why. However, if Australia was to implement such an initiative, we think that the scope of any sort of transparency requirements will need to be carefully considered, and balance the consumer's right to know, with the technical practicality of providing information about public sector AI use; and
- implement a public register similar to the Amsterdam Register.

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

Please see our responses to Questions 5 and 6 above in respect of the Amsterdam Register and the Transparency Standard.

Question 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.

We assume that the Paper's references to:

- 'generic solutions' means the industry-agnostic, risk-based approach to AI regulation (like that which has been proposed in the EU) that seeks to regulate AI in a general manner, regardless of its specific use cases; and
- 'technology-specific solutions' means the decentralised, sector-specific approach (like that which has been proposed in the UK) that seeks to regulate AI based on its specific use cases, including in respect of consumer interaction.

On the basis of this assumption, we consider that the appropriate solution to addressing AI risk will depend largely on where the relevant application is in its lifecycle. For example, we think that generic solutions to AI regulation may be best-suited to AI applications that are in the early stages of development, with few real-world use cases, or are still in pre-production. We think that such solutions could include:

- guardrails around development based on the intent of the product and any potential use cases;
- mandated human oversight and quality assurance, even the appointment of a designated 'AI Officer' (much like the Data Protection Officer mandated by the GDPR); and
- requirements to keep detailed documentation on the development process, applications and training methodologies, coupled with audit powers conferred on regulators.

As the relevant application evolves and use cases for that application become clearer, then more-technology-specific solutions may be implemented. For example, the relevant application could receive more-specific regulatory treatment once it becomes a consumer-facing ADM application.

Question 9:	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? • mandating transparency requirements across the private and public sectors, including how these requirements could be implemented.

We consider that upfront and proactive transparency is the best way in which public trust and confidence in AI can be fostered, as we think that the public will want to understand when, why and how they interact with AI, as well as the unique ways in which they may be affected by their use of AI, as early as possible in the customer journey.

In respect of mandating transparency requirements, while we think that such requirements may be appropriate in some circumstances (see also our response to Question 2 above), the key questions to be answered of any proposed approach relate to:

- the scope of such requirements and their application – for example, whether the law is framed as a consumer right to request transparency, as well as a requirement on providers to provide it (including the degree to which it is to be provided);
- how such requirements are implemented; and
- whether these requirements are legislative or regulatory, principles based, technology-based or sector-based.

In any case, we consider that such requirements:

- should not be so onerous to the point of stifling innovation and obstructing or preventing smaller players from entering the market;
- cannot attempt to adopt a one-size-fits-all approach; and
- should be subject to an appropriate implementation period, given that there will likely be a need for organisations to adopt technical solutions to comply with any such requirements.

We think that a requirement similar to that which is in place under the Privacy Laws (which requires up-front disclosure in respect of the ways in which the regulated entity handles a person's personal information), whereby organisations are required to provide an overview of uses of AI in its systems and processes (especially customer-facing processes), would be a suitable approach that is not overly-burdensome.

Question 10:	Do you have any suggestions for: • whether any high-risk AI applications or technologies should be banned completely? • criteria or requirements to identify AI applications or technologies that should be banned and in which contexts?

As canvassed above, we consider that a decentralised, sector-specific approach to AI regulation is generally a better solution to limit the potential risks and complications of AI, without unduly stifling innovation, as compared to a risk-based approach via legislation.

In our view, it is the specific application of AI-driven technologies that predominately engenders risk, not the underlying technology itself and accordingly, decentralised regulation is preferable (as compared to any broad, blanket restriction), as it allows for particular use cases to be examined on an individual basis for their potential to cause harm.

Notwithstanding the above, we recognise that certain applications of AI technology may have an inherently high risk of harm to individuals and accordingly, we consider that such high-risk applications should be banned or restricted depending on the degree to which those applications have the potential

to cause harm. An example of an application that we consider should be closely considered in terms of the regulatory approach to it, is the use of machine learning techniques to manipulate visual and audio content for the purpose of replicating someone's likeness without their consent (otherwise known as 'deep synthesis' or a 'deepfake'). This type of AI-driven application has obvious potential to be used for all manner of nefarious purposes and accordingly, we consider that the development and deployment of this type of use case should be subject to some form of 'harm minimisation' regulation.

However, Government should be careful to ensure that such regulation is implemented in a way that it does not ban or restrict the underlying technology that powers the high-risk application (thereby stifling innovation in respect of that underlying technology), but only the high-risk application itself.

For example, in the EU, the AI Act does not seek to ban the use of AI-driven biometric identification systems entirely, but rather, limits its use in the following ways:

- it may only be used on an ex-post facto basis for the prosecution of serious crimes; and
- it may not be used in real-time in public spaces at all.

Further, the AI Act places a ban on practices that have significant potential to manipulate persons through subliminal techniques, practices that exploit vulnerable groups and AI-based social scoring – these are all initiatives that seek to regulate the application of the technology, not the technology itself.

We support a similar approach to the regulation of high-risk AI applications in Australia.

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

Please see our responses to Questions 5 and 6 above in respect of the Amsterdam Register and the Transparency Standard. To this end, we agree with Government's view that greater public trust and confidence is required in order to increase AI adoption and we consider it appropriate that the Government 'leads from the front' in this respect.

Further, we consider that regulation cannot be created in a vacuum. Rather, what is needed is a collaborative public/private approach which provides a forum for inputs from both industry and regulators. Accordingly, the Government should explore initiatives that allow opportunities for the public, the private sector and Government to knowledge share and collaborate in respect of AI regulation.

From a public education perspective, Government should also consider creating guidance and public education initiatives to educate the public on how AI technology is currently being used in their day-to-day lives. This would create a better-informed public on the risks and benefits of AI. We set out some of these initiatives within our response to Section 3.

Section 5: Implications and Infrastructure

Question 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 exports with other countries?

Currently, AI use in the insurance and insurtech industries is primarily focussed on areas including organisational risk assessment, fraud detection, pricing, policy generation, claims management and customer experience. Accordingly, so called 'high-risk' use cases such as social scoring or facial recognition are less relevant to the insurance and insurtech industries.

However, we consider that such 'high-risk' use cases may become relevant to the insurance industry in the future in a number of ways. Using the examples provided in this question:

- social scoring applications could be used to gauge the risk profile of a potential insured, when deciding whether to insure the relevant individual; and
- facial recognition could be used to assess the mood and attitude of a customer during a customer service interaction, to enable an insurer to improve its customer service processes.

We do not think it is necessarily the case that such use cases will always be high risk – for example, we consider that the facial recognition example provided above is an example of a more innocuous application of facial recognition technology.

Accordingly, while we think that the ethical considerations pertaining to high-risk use cases must remain paramount, we do not consider that banning such technologies completely is a sensible regulatory approach, and doing so would certainly stifle innovation in the technology and insurance sectors, via the inability to apply these technologies to more low-risk use cases. Our view is that a decentralised approach to regulation is preferable, due to its ability to regulate specific use cases that may be high-risk, as opposed to the underlying technology itself.

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

We consider that industry should lead the development of conformity infrastructure and relevant AI technical standards, and while we acknowledge that there are AI risk management standards by ISO and the National Institute of Standards and Technology (**NIST**) in the market currently, we do not think that these standards are sufficient by themselves.

Additional conformity efforts will need to take place across all stratum of industry, not just the end of industry that is sophisticated enough to warrant the adoption of standards such as those of ISO and NIST. We consider that Government should adopt a graduated approach to conformity infrastructure, based on factors such as an organisation's size or annual turnover.

This would reduce the regulatory burden for smaller organisations and give them flexibility to innovate, whilst still establishing a mainstream requirement for conformity and technical consistency. However, we note that the inherent flaw in this approach is that use of AI is not unique to large organisations and therefore smaller AI users could fly under the regulatory radar, even if they are using AI in a way that carries material risk.

Section 6: Risk-based Approaches

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

We recognise the consistency and uniformity benefits of a broad risk-based approach for regulating AI, however, we think that such an approach is inherently rigid. Further, as technology evolves, legislation will need to be constantly updated to contemplate the new risks that may be presented by such evolution. Accordingly, we do not support a generic approach to restricting use of technology, except potentially at early stages in the AI development lifecycle (as per our answer to Question 8 above).

As conveyed earlier in this submission, we think that a decentralised approach would allow for greater innovation and flexibility to consider a broad range of specific use cases. Further, a decentralised approach also allows for greater isolation of risk so that specific attention and effort can be focused where regulation is most-needed, based on an assessment of the sector-based risks raised by a particular application.

Given that existing legislation such as the Privacy Laws, Australian Consumer Law and various anti-discrimination laws already regulate certain risks related to AI to a certain extent, one approach may be to create legislation that seeks to incrementally fill gaps in the existing legislative and regulatory frameworks, thereby creating a light-touch approach to AI regulation that remedies shortcomings in the current regulatory regime, without introducing unnecessary complexity or regulatory inconsistency.

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

We consider that the key benefits of a risk-based approach are that it:

- provides a comprehensive and harmonised regime for AI governance, minimising the need to navigate various regulatory regimes that may conflict with one another;
- focuses on the risks related to the way in which AI is applied;
- seeks to place guardrails around the development of AI in the first instance; and
- allows for greater oversight of general-purpose AI technologies that may impact multiple sectors.

However, we consider that the key limitations of a risk-based approach are that it:

- could be difficult to assign and future-proof general risk categories and define in what circumstances AI systems should be regulated or banned;
- would not allow sufficient ‘flex’ to permit or prohibit specific uses cases of technology based on context;
- would need to be constantly updated as new AI applications and technologies are developed; and
- could create duplication and overlap between existing general and sector-specific regulation.

We consider that the above limitations can be overcome by adopting a sector-led approach that accounts for existing regulatory regimes (as is being proposed in the UK), but with elements of the risk-based approach whereby the stringency of regulation depends on the risks posed by the specific use case in that sector.

Further, this should be supplemented with practical initiatives such as guidance resources, AI toolkits and regulatory sandboxes to help ease regulatory burden and support innovation.

Question 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?

We have assumed that this question seeks to understand whether a generic, risk-based approach to regulation would be better suited towards sectors that are larger and have greater AI maturity and resource capacity.

Notwithstanding that we support a decentralised approach to regulation, we consider that larger organisations would be better equipped to navigate a risk-based approach to regulating AI due to their size, sophistication and resourcing capacity. If a generic, risk-based approach is adopted, the danger for smaller businesses is that they may lack the proper organisational risk management frameworks to manage compliance. This could have the effect of stifling the innovation efforts of smaller market players, from where much of the AI innovation in industry begins.

We also believe that a generic approach to industry regulation would be better suited towards industries and sectors that are not already subject to specific regulatory frameworks. To the extent that existing regulatory frameworks already impose restrictions on AI usage (explicitly or otherwise), then additional regulation may cause confusion, if it overlaps with existing regulatory requirements.

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?

In our view, Attachment C covers several important aspects of risk-based AI governance and overall, we are supportive of the elements described in Attachment C. Introducing risk assessments, notices and human oversight (but only as a short-term measure until human intervention is not needed) are all valid steps that will create guardrails for the safe development of AI.

In addition, we consider that the development of policy frameworks will also be an important aspect of any risk-based approach for addressing AI risks.

Question 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?

We consider that, as existing legislative frameworks already touch on and regulate AI to a certain degree and in respect of certain risks, any approach to incorporating AI regulation to existing frameworks should begin with a gap analysis to understand the areas that are not already addressed.

Following this, we consider that a light-touch approach should be taken where gaps in existing regulation are identified and then addressed. Such gaps could be addressed by, for example:

- the development of formal guidance from the OAIC in relation to the application of the Privacy Laws to AI;
- clarifying the role of privacy impact assessments in identifying and controlling the risk of privacy harm to individuals in connection with the development, deployment and use of AI; and
- incorporating into an organisation's privacy policy and personal information collection notices, disclosures about automation and AI being used in the relevant organisation's services – however, this would only be appropriate for entities that are subject to the Privacy Laws.

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

We consider that, for general purpose AI systems, their risk level and the subsequent stringency of the associated requirements should be determined based on the nature of their specific application in the relevant context.

As stated earlier in this submission, the generic approach to regulating a technology itself, as opposed to specific use cases, can unduly stifle innovation and place undue burden on potentially-innocuous applications of that general technology.

In our view, any risk-based approach to AI should focus on the risk of specific applications and use cases, rather than the risk potential of the technology itself.

Question 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: • public or private organisations or both? • developers or deployers or both?

We consider that a mixture of enforceable (but light-touch) regulation, supplemented by principles-based guidance, may be appropriate to regulate AI in Australia.

As it is expected that AI will proliferate across the public sector and industry, any guardrails established should apply to both the public and private sectors. At the same time however, due to the nature of the public sector and the need to engender public trust and confidence in AI, we consider that public sector uses of AI in particular, should be subject to probity and transparency requirements.

Such regulation should also apply to both developers and deployers alike, so as to capture the entire lifecycle of AI. This would serve to prevent the development of AI with high harm potential, as well as restrict the improper use of AI technology.

Regardless of the regulatory methods adopted however, we consider that the correct approach is one that does not stifle innovation.