21 April 2022

Department of the Prime Minister and Cabinet
Digital Technology Taskforce
PO Box 6500
Canberra, ACT 2600

Email: digitaltechnologytaskforceinbox@pmc.gov.au

Dear Sir/Madam

Response to Digital Technology Taskforce Issues Paper: Positioning Australia as a Leader in Digital Economy Regulation - Automated Decision Making and AI Regulation

The Actuaries Institute ('Institute') welcomes the opportunity to respond to the Issues Paper "Positioning Australia as a Leader in Digital Economy Regulation – Automated Decision Making and AI Regulation" released on 13 March 2022.

The Institute is the sole professional body for actuaries in Australia and has been actively involved in public discussions of digital economy regulation, including:

- The Australian Human Rights Commission Discussion Paper "Human Rights and Technology Discussion Paper", from March 2020 with the Institute's recommendations referenced several times in the final report;
- Treasury's Consultation on the Strategic Assessment of the Implementation of the economy-wide Consumer Data Right and implications for the Superannuation Sector from September 2022
- Institute Members involvement piloting the Australian Government's AI Ethics Principles Framework during 2020-2021**; and**
- The Institute CEO's Paper at the joint ABS/RBA Economic Implications of the Digital Economy Conference in March 2022.

We would also like to draw your attention to work the Institute is undertaking with The Australian Human Rights Commission to produce an anti-discrimination guidance resource for actuaries regarding the use of artificial intelligence (AI) in relation to insurance pricing and underwriting which we expect to be published in the second half of 2022.

In responding to the broad scope and specific questions of this Issues Paper, we have provided high level comments, drawn from the Institute's previous submissions and thought leadership publications. We first provide comments outlining our perspectives on regulation in general and how this might relate to the areas in question. We then respond to the questions set out in the Issues Paper, and in some cases combine questions of similar theme, for brevity.

We welcome any opportunity to discuss this important topic with you in more depth.

Institute of Actuaries of Australia
ABN 69 000 423 656
Level 2, 50 Carrington Street, Sydney NSW 2000, Australia
t +61 (0) 2 9239 6100 **f** +61 (0) 2 9233 3446
actuaries@actuaries.asn.au | www.actuaries.asn.au

‘Good’ Regulation, and Regulatory Uncertainty

The Institute’s principles guiding our response to policy development includes a commitment to ‘good’ regulation. We consider that excessive or unnecessary regulation can obstruct a market from functioning efficiently and can undermine the ‘public interest’. Some key elements of ‘good’ regulation include:

- proportionality – between the regulatory solution and the problem that it intends to solve; and
- appropriate regulatory tools – using self-regulation where possible while acknowledging that prescription can sometimes be appropriate.

The fundamental point made by this principle – and in Institute publications and positions on regulation – is of the importance of balance and proportionality. Poorly designed or poorly targeted regulation can create as many problems as it seeks to solve. Hence, we make the general suggestion that the Taskforce should ensure that any proposed regulation in this area is appropriately targeted and proportionate to the problems it seeks to solve.

An important outcome of ‘good’ regulation is to provide clarity and certainty to institutions and their employees of their expected conduct, and to the public of the conduct they might expect of those institutions. Professionals, such as actuaries operating as employees or advisors, can then have confidence in the work they are performing, and clarity in the rules that apply.

The Issues Paper suggests that the recent popularisation of Artificial Intelligence (AI) and Automated Decision Making (ADM) has led to situations where existing regulation is unclear or potentially absent.

Whilst it is often suggested that AI and ADM are new, in many industries they have been present for a long time. For example, Actuaries Institute (2022)[1] notes the use of data-driven decision processes in insurance underwriting and pricing at least as far back as 1762 (this is an early example of ADM – using less modern tools than many of today’s systems – applied to at least part of a significant decision process). We suggest that in contemplating reforms, it may be useful to consider the regulatory history of industries (including insurance) which have long used ADM.

We agree with the Taskforce that in the presence of AI and ADM, it has become apparent that there are some uncertainties and gaps in regulation in some areas. We suggest that instances of this may arise from three quite separate issues:

1. A situation where there is no adequate regulation for the context of the AI/ADM system; or
2. A situation where there is some existing regulation for the context of the AI/ADM system, but it is unclear how that regulation applies to AI/ADM; or
3. A situation where there is some existing regulation for the context of the AI/ADM system, but it is unclear how that regulation applies to many situations, including but not limited to AI/ADM.

These are separate matters, requiring different responses. To create ‘good’ regulation, the first step must be to identify the precise nature of the situation requiring remedy.

Australia generally operates with broad, principles-based regulation in key areas of the economy. Whilst this does not mean the first situation above will be entirely avoided, it suggests

[1] Actuaries Institute (2022) “Big data and the digital economy: Benefits and pitfalls in the insurance industry”. https://www.actuaries.asn.au/research-analysis/thought-leadership/big-data-and-the-digital-economy; page 4.



that generally there is likely to be some existing regulation which applies to a situation in which AI/ADM is to be applied.

In many cases, we suggest that uncertainty in regulation will not be limited to AI/ADM, but AI/ADM will instead reveal to us a more general gap in our regulation. This is the third situation above, which we suggest is likely to be commonly encountered. For example, and as we detailed in our response to the Human Rights Commission in 2020, there are well-understood conflicts and competing definitions in considering what is intended by prohibitions on indirect discrimination[2]. These conflicts are not new. AI and ADM do not create them, rather the recent discussion of AI and ADM has revealed these issues to us with greater clarity. An appropriate response would be to correct this existing uncertainty, by issuing suitable guidance around indirect discrimination including but not limited to situations involving AI/ADM.

We elaborate on these general comments further in response to specific questions.

1. What are the most significant regulatory barriers to achieving the potential offered by AI and ADM? How can those barriers be overcome?

4. Are there specific examples where regulations have limited opportunities to innovate through the adoption of AI or ADM?

Any lack of clarity of existing regulation can be a barrier to innovation – particularly in sectors like financial services where the costs of non-compliance can be high. Unclear regulation is not 'good' regulation, and this carries costs. A lack of certainty over the law can – and does – result in projects that may be of value to both the community and to industry being abandoned, for fear of breaching that unclear regulation.

Whilst case law and regulatory guides often exist to provide some clarity over the interpretation of high-level regulatory principles, these are often backwards looking, or written with traditional, human-centric decisioning systems in mind. Hence many existing regulations are likely to be unclear in the context of AI/ADM. Most important to note is that automated systems involve data, numbers, and mathematics. This includes systems involving unstructured data which humans may not intuitively recognise as being mathematical, but which an AI/ADM system must encode into a mathematical form to utilise. This essential role of mathematics in AI/ADM means a level of precision is imposed on the AI/ADM decision process which human decision processes may not traditionally contain. This means that high level 'principles based' regulation written in words may not always be precise enough to authoritatively guide the design of such systems. Guidance to interpret such regulation written only in words may again not always be specific enough to authoritatively guide the design of AI/ADM. It is essential that guidance which is intended to be applied in a mathematical context such as AI/ADM avoids such uncertainty. This does not necessarily mean that guidance needs to itself contain mathematics (though this may be an appropriate approach in some situations), but at the least the guidance given should not lead to ambiguity when translated into the mathematics of an AI/ADM system. Use of greater precision in guidance may have the added benefit of providing greater clarity on the expected aggregate outcomes of human decision systems.

For example, the concept of 'fairness' is embedded in many principles-based regulations around the world today. An example in Australia is the requirement under the Corporations Act 2001 (Cth) for financial services companies "to do all things necessary to ensure that the financial services covered by the licence are provided efficiently, honestly and fairly"[3]. But what precisely is intended by a requirement of 'fairness', in a given context? Recent academic

[2] Actuaries Institute (2020). "Actuaries Institute Response to Human Rights and Technology Discussion Paper" https://content.actuaries.asn.au/resources/resource-ce6yyqn64sx3-2093352434-18053; page 2.

[3] Corporations Act (2001) s912A(1)(a)



studies of 'fairness' in particular instances of AI/ADM have yielded important insights, most notably the incompatibility of intuitively reasonable mathematical definitions of 'fairness'.[4] This means a general requirement to act 'fairly' is not a specific enough instruction to encode into an AI/ADM system. If we must define what vague terms like 'fairness' mean for our decisions mathematically – as we must for AI/ADM – we require far more granular, specific guidance as to the correct interpretation of such words in particular situations.

Hence, we consider that the primary regulatory action required to allow greater adoption of AI/ADM is to clarify the operation of existing regulation. This guidance must be authoritative and should be specific and detailed enough to guide the design and operation of AI/ADM. We suggest a starting point should be situations where AI/ADM is likely to be introduced or is already present, and where the decision being made can potentially cause harm to consumers. This guidance must allow us to confidently translate the words in regulation into the mathematical instructions required for AI/ADM to be constructed and operated. As we say in Actuaries Institute (2022): "The law must be made clear"[5].

3. What specific regulatory changes could the Commonwealth implement to promote increased adoption of AI and ADM? What are the costs and benefits (in general terms) of any suggested policy change?

6. Are there actions that regulators could be taking to facilitate the adoption of AI and ADM?

As already suggested, the Commonwealth can act to create regulatory clarity around AI/ADM, by instructing all relevant regulators under their authority to undertake reviews of existing regulations in the context of AI/ADM and to issue guidance as needed to create that clarity. Regulators could also seek to conduct such activity without Commonwealth instruction – including regulators operating at other levels of jurisdiction (e.g. State). We suggest that industry and professional bodies will be willing and able to assist in identifying specific areas where greater clarity is needed and will be able to make suggestions as to what clarity might mean in those contexts, in response to any further consultations.

AI/ADM is powered by data. Hence actions to accelerate the openness and availability of data would likely increase the rate of adoption of AI/ADM. However, such a move is not without costs. As we note in Actuaries Institute (2022)[6], this cost may not fall evenly on the community but may be borne by those who are already considered disadvantaged or vulnerable. In contemplating the correct balance of benefits and costs, and regulation to strike that balance, we suggest a productive avenue of inquiry is to consider stronger regulation surrounding the use of particular forms or types of data in particular contexts (including but not limited to AI/ADM). Following the comments outlined previously, this could be in the form of guidance to elaborate on existing principles-based regulation of high stakes decisions such-as found in financial services.

2. Are there specific examples of regulatory overlap or duplication that create a barrier to the adoption of AI or ADM? If so, how could that overlap, or duplication be addressed?

As identified in our submission to the Human Rights Commission in 2020, there is an inherent conflict and overlapping intent between ideals of privacy, and the collection of data about protected attributes covered under anti-discrimination law to avoid indirect discrimination. We suggest that this fundamental problem is now so well understood that it is likely to be a barrier to adoption of AI/ADM in some contexts. Resolving this conflict would, therefore, allow greater

[4] See, for example, Kleinberg, J., Mullainathan, S., & Raghavan, M. (2016). Inherent trade-offs in the fair determination of risk scores. https://doi.org/10.48550/arXiv.1609.05807
[5] See earlier reference, page 16
[6] See earlier reference



adoption to occur – and would generally be of benefit to the community across a range of contexts. Since this conflict spans quite different areas of regulation, there may be a role for the Australian Law Reform Commission in helping to resolve it.

5. Are there opportunities to make regulation more technology neutral, so that it will more apply more appropriately to AI, ADM and future changes to technology?

7. Is there a need for new regulation or guidance to minimise existing and emerging risks of adopting AI and ADM?

We caution against taking the approach of the European Union (EU) to set specific regulation for AI/ADM[7]. Our submission to the Human Rights Commission outlines arguments which would also apply to this form of omnibus regulation of AI/ADM[8]. Fundamentally, this is because we agree with the proposition of the Issues Paper that regulation should be technology neutral. Specific regulation for AI or ADM is, by definition, not technology neutral and in our view is not 'good' regulation. It would almost inevitably lead to poor outcomes, including:

1. Overlaps and inconsistencies will be created, between AI/ADM regulation and other existing regulations covering the decisions to which AI/ADM is put. We suggest that the level of complexity created will be unrealistic to manage – inconsistencies will emerge.
2. Uncertainties will exist over the boundaries of AI/ADM, however these terms are defined, since this is a rapidly evolving domain. Any definition used is likely to become rapidly out of date.
3. Various forms of two-tier system will be created, where a particular form of poor outcome or harm may be disallowed if it occurs via AI/ADM but allowed if it occurs via some other means (for example a human process). Rather than create such a system, poor or harmful outcomes should be regulated consistently, irrespective of the means to those outcomes. We note that some regulations might still be far more significant in the context of AI/ADM than traditional contexts, but this is acceptable if the regulation is reasonable in the context. For example, a right to an explanation of a decision may be deemed appropriate in a particular situation, and this might be substantially more onerous for an AI system than a human one.

In some situations, there will likely be a need for new regulation in light of opportunities created by AI/ADM. For example, facial recognition technology will allow surveillance of the public by law enforcement agencies on a scale which was previously impractical. In these sorts of situations – where AI/ADM creates something genuinely new – existing regulation may be inadequate, and clarification of existing rules may be insufficient. In response to question seven (7), we suggest that if new regulation is required to manage novel outcomes created by AI/ADM, the regulation should still be written in a technology-neutral manner, not made specific to AI/ADM.

We consider the Australian AI Ethics Principles to be a useful starting point for organisations considering the use of AI/ADM in a particular context, particularly if regulation is absent or unclear. As noted above, some of our members have been involved in testing these principles in practice. To aid adoption, and in response to question seven (7), we suggest that these principles should be accompanied by more practical guidance in order that practitioners have clearer instructions to follow. The Institute has issued guidance of this form for our Members[9] which we suggest could be used to inform broader guidance which government

[7] European Commission (2021).https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52021PC0206
[8] See earlier reference, page 3-4
[9] Actuaries Institute (2020) "Information Note: Automated Decision-Making Systems", Section D



may seek to issue. However, we suggest that such ethics principles – even with more detailed guidance – is no substitute for effective regulation and guidance of a high stakes decision context.

8. Would increased automation of decision making have adverse implications for vulnerable groups? How could any adverse implications be ameliorated?

We suggest that the use of AI/ADM can both benefit and harm vulnerable communities. For example, non-English speakers may benefit from AI systems which offer cheap, automatic translations, but other AI/ADM systems may create racist outcomes for those same people. So, similarly to our response to other questions, we suggest that:

- Potential harms caused by AI/ADM systems should be regulated by general regulation focussing on harmful outcomes, not on AI/ADM specifically; and
- The operation of those regulations should be made clear, in the context of AI/ADM. This is likely to involve greater levels of mathematical precision than is traditionally found in regulation or guidance material.

AI/ADM typically involves digital technology as an enabler. If we assume that services embedding AI/ADM are better, faster and/or cheaper than traditional services, then there is a material risk that vulnerable members of the community – who may not have the same level of access to or ability to use digital services – will suffer. This is not without precedent. For example, the advent of telephone and then online banking has been an important driver of bank branch closures, which is generally accepted to have had a detrimental impact on certain members of the community (though was beneficial for others). We suggest that mass-adoption of AI/ADM is likely to create similar situations across various areas of the economy, with legacy products being withdrawn or becoming more costly to those who still require them. Whilst innovation and modernisation is part and parcel of any vibrant economy, we suggest that an appropriate regulator should be tasked with continuous monitoring of the impact of this potential in socially significant markets, to understand if market failure is occurring, or likely to occur.

Related to, but distinct from, the adoption of AI/ADM is the increasing amount of data available to drive these automated decisions. This increasing granularity and personalisation of decisions carries with it a material risk of detriment to members of society who are already vulnerable. As we already noted above, Actuaries Institute (2022) identifies many examples where this could occur[10]. Whilst this paper uses insurance examples to illustrate, particularly the unwinding of historic cross-subsidies (a theme which we explore in more detail in Actuaries Institute (2016)[11]), the problems are general in nature and will impact many sectors of the economy in similar ways. Contemplation of the impact of greater availability of data on markets and competition, including the impact on vulnerable people and any protective steps required as a result, is again likely to be a role for the ACCC.

[10] See earlier reference

[11] Actuaries Institute (2016) "The Impact of Big Data on the Future of Insurance" https://content.actuaries.asn.au/resources/resource-ce6yyqn64sx3-2093352434-53835

9. Are there specific circumstances in which AI or ADM are not appropriate?

We note various proposals to ban the use of AI/ADM in particular contexts. Perhaps most notable is the series of prohibited uses of AI within the EU's proposed regulation[12]. We observe that – aligned to comments made in the Issues Paper – such prohibitions are not technology neutral. Again, and echoing comments made earlier in this submission, we suggest that such prohibitions may be ineffective. Instead, consideration should be made to make general prohibitions of particular outcomes that are deemed unacceptable, not merely certain means to those same outcomes, via AI/ADM.

That said, we do not discount the possibility that there may be a place for specific prohibitions of AI or ADM, but we suggest these should be rare exceptions, not the rule.

10. Are there international policy measures, legal frameworks or proposals on AI or ADM that should be considered for adoption in Australia? Is consistency or interoperability with foreign approaches desirable?

We refer to comments made in response to questions five, seven and nine in relation to the proposed EU regulations on AI.

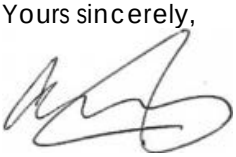
While we have no strong views either way on the desirability (or not) of international consistency or interoperability of AI/ADM laws or policies, we suggest that effort to drive international consistency over AI/ADM laws may ultimately be of limited value. This is because there is generally not international alignment of laws, frameworks or standards in the areas to which AI/ADM will be applied. For example, Australia's laws regarding conduct of financial services companies are different to the laws of the UK, or the US, or the EU, or China. AI/ADM applied to insurance underwriting or loan applications in those various jurisdictions would then have to consider compliance with each of these separate financial services laws. This will necessitate differences – potentially fundamental differences - in those AI/ADM systems across the globe. This is aside from any additional laws which might apply to AI/ADM, and any interoperability or consistency of those laws. In short: laws are generally already inconsistent around the world, and this already means that AI/ADM cannot generally be applied consistently around the world – whether AI/ADM-specific laws are implemented or not.

More critical, in our view, is to ensure consistency between any Australian AI/ADM rules and broader legislation in Australia covering the uses to which AI/ADM might be put (for example financial services laws, as highlighted earlier). If a choice is to be made, it is likely to be preferable to have consistency of rules within Australia, rather than international consistency only of specific AI/ADM rules, but inconsistency of those rules in particular areas of AI/ADM application to other Australian rules.

Further discussion

The Institute would be pleased to discuss this response with the Digital Technology Taskforce and looks forward to contributing to further discussions. Please contact Elayne Grace, Chief Executive Officer of the Institute, at elayne.grace@actuaries.asn.au for further discussion relating to this response.

Yours sincerely,



Annette King
President

[12] See earlier reference, Article 5