New Zealand Law Society
Te Kāhui Ture o Aotearoa

17 July 2026

Brieze Va'a
Policy Advisor | Kaitohutohu Kaupapa
Inland Revenue | Te Tari Taake

By email: brieze.vaa@ird.govt.nz

Tēnā koe Brieze,

Feedback on proposed Automated Decision-Making standard for Inland Revenue

1 Introduction

1.1 The New Zealand Law Society Te Kāhui Ture o Aotearoa welcomes the opportunity to provide feedback on Inland Revenue's Automated Decision-Making (**ADM**) standard and principles.

1.2 This feedback has been prepared with input from the Law Society's Public Law and Tax Law Committees.

2 General comments

2.1 Appropriate use of ADM has the potential to streamline Inland Revenue's processes, and allow its limited resources to be redirected towards other functions that benefit New Zealand and its taxpayers. Such use of ADM would be consistent with the statutory duty of the Commissioner of Inland Revenue (**Commissioner**) to collect over time the highest net revenue that is practicable within the law having regard to the resources available to the Commissioner (among other matters).[1]

2.2 However, the inappropriate or unsafe use of ADM has the potential to undermine the integrity of the tax system, the rule of law and trust in the Government more generally. Under section 6 of the Tax Administration Act 1994 (**TAA**), the Commissioner of Inland Revenue is charged with using their best endeavours to protect the integrity of the tax system. While the proposed standard acknowledges that the implementation and use of ADM must uphold the integrity of the tax system,[2] no specific consideration is given to how the proposed principles advance the factors outlined in section 6(2) of the TAA.

1 Tax Administration Act 1994, s 6A(2).
2 Under the principles 'Improve efficiency and effectiveness of decision-making' and 'Legal, privacy and policy considerations'.

2.3 We also note with concern that the proposal is to pair a "broadly framed legislative provision" with safeguards that will be set out in an internal operational standard.[3] While this standard is based on the Ministry of Social Development's Automated Decision-Making Standard,[4] the drafting and precise scope of the legislative authorisation is unclear. Will it, for example, resemble that of section 363A of the Social Security Act 2018 (**SSA**), as amended by the Social Security (Modernisation) Amendment Act 2026? Section 363A already requires an approved standard for MSD's use of automated electronic systems, and the recent amendment will specify matters that must be addressed in the standard. However, section 363A(3), which specifies the decisions for which ADM can be used by MSD, has been repealed.

2.4 We presume that reference to a "broad authorisation" is indicative of a provision that will likely reflect the amended section 363A of the SSA and therefore contain few legislative constraints on *when* ADM can be used, requiring instead a standard that simply addresses certain considerations about the system being used. Even then, it is the existence of a standard that is mandated, and not particular safeguards or constraints, or even a requirement to remove or minimise certain risks.[5]

2.5 The Law Society is therefore concerned by the proposal to grant Inland Revenue broad authority to use ADM while the constraints on that use remain self-administered and amendable without Parliamentary scrutiny. In our view, core safeguards (including suitability criteria, human review rights, and transparency) should sit within primary legislation, with operational details left to the standard.

2.6 ADM processes also raise rule of law concerns where it implicitly allows a level of unlawful actions to occur, or reduces accountability and transparency around how the law has been applied to a particular set of circumstances.[6] The Law Society therefore supports the introduction of safeguards to enable appropriate and responsible use of ADM within Inland Revenue (but note our preference above for core safeguards to be prescribed in primary legislation).

2.7 The proposed principles are a credible starting point. However, when assessed against section 6 of the TAA, which frames the integrity of the tax system around taxpayer perceptions, taxpayer rights, the Commissioner's care and management responsibilities, and confidentiality, there remain some gaps and concerns that should first be addressed. These include:

(a) a definitional loophole that would allow high-risk processes to be exempt from the proposed ADM standard entirely;

(b) the absence of a remediation regime for decisions made by a defective ADM system; and

---

3 Page 1 of the consultation document.

4 Page 1 of the consultation document.

5 The Law Society has previously raised concerns about MSD's approved standard. See submission on the Social Security Amendment Bill 2025.

6 We discuss the rule of law implications of ADM at [9.16] of our 2025 report, *Strengthening the rule of law in Aotearoa New Zealand*, available here: www.lawsociety.org.nz/assets/Law-Reform/Rule-of-law/Strengthening-the-rule-of-law-in-Aotearoa-New-Zealand.pdf.

(c) failure to reflect the full extent of risks and disadvantages associated with the use of ADM, which relies on rapidly developing technology and social licence.

2.8 We discuss these and other matters in more detail below. We first turn to issues arising from each proposed principle, and then address other issues which are not discussed in the consultation document, and require further consideration.

3 Principle 1: improve efficiency and effectiveness of decision-making

3.1 Principle 1 appears to be an objective rather than a safeguard, and the combination of efficiency and effectiveness into a single principle has the potential to weaken this safeguard (for example, where measures to improve efficiency might be at odds with measures to improve effectiveness). We recommend restructuring this principle in a way that clearly identifies its purpose, scope, and safeguards.

4 Principle 2: accuracy and reliability

4.1 The 'accuracy and reliability' principle appropriately acknowledges the importance of providing clear, relevant and accessible guidance for taxpayers who are required to input or provide data to be used in ADM.

4.2 However, the requirement to produce "expected results at all times"[7] may be an unachievable test that invites either non-compliance or dilution in practice, particularly given the principles do not specify any requirements relating to pre-implementation testing regimes, accuracy thresholds, or data sampling to assess accuracy and reliability.

4.3 We recommend modifying this principle to explicitly require Inland Revenue to:

(a) undertake pre-implementation validation of data,

(b) document accuracy thresholds that are proportionate to the risks associated with each type of automated decision, and

(c) undertake ongoing statistical sampling of outputs to assess accuracy and reliability.

5 Principle 3: bias and discrimination

5.1 Principle 3 requires the Commissioner to identify risks relating to bias and discrimination, and to take steps to remove or mitigate those risks. This principle raises the following concerns:

The option to mitigate risks of bias and discrimination

5.2 The option to either remove *or* mitigate risks implies that some level of bias or discrimination may be acceptable without warranting human involvement. This has the potential to undermine the integrity of the tax system (for example, where it results in an individual's tax affairs being treated with greater or lesser favour than the tax affairs of others,[8] or it results in a failure to administer the law fairly, impartially and according to the law),[9] and should be reconsidered.

---

7 Page 2 of the consultation document.
8 Section 6(2)(c) of the TAA.
9 Section 6(2)(f) of the TAA.

5.3 Our preference would be to require the Commissioner to *remove* risks relating to bias and discrimination, rather than simply *mitigating* those risks. If removal of risks is not possible, the ADM should not be permitted. Otherwise, mitigation should be considered as an option only where that is legally permissible, and where any risks that are identified do not result in unlawful bias or discrimination. For these reasons, we recommend amending principle 3 to provide that:

(a) ADM systems must be designed to actively prevent unlawful bias and discrimination, and

(b) mitigation should be an option only where it will not breach administrative law rules or statutory requirements.

Consequences of human involvement

5.4 Principle 3 suggests that human involvement in a decision-making process would mean the decision-making is no longer automated, and would therefore be exempt from Inland Revenue's ADM standard.[10] This can create perverse incentives to exclude certain decisions (particularly high-risk decisions) from the proposed ADM regime by enabling nominal human involvement or sign-off on automated decisions.

5.5 International experiences and existing literature on automation bias show that humans reviewing machine outputs at volume tend to 'rubber-stamp' automated decisions without any meaningful review or engagement.[11] Processes where bias and discrimination cannot be mitigated should therefore attract more – rather than little or no – scrutiny when there is human involvement.

5.6 Consideration should also be given to extending these ADM standards to both fully and partly-automated decision-making processes (i.e., processes with some human involvement). It could also be helpful to clarify what constitutes 'human involvement', and whether it requires the human reviewer to have particular skills and capabilities, adequate time to review decisions, authority to depart from an automated decision, and the ability to monitor override rates.

6 Principle 4: legal, privacy and policy considerations

6.1 Principle 1 ('improve efficiency and effectiveness of decision-making') states that decisions made by ADM must be clear and rules-based: the Law Society considers this to be an appropriate safeguard, but notes that it appears to be at odds with principle 4 ('legal, privacy and policy considerations'), which contemplates the automation of discretionary decisions where legal and privacy risks are "identified and then mitigated, removed, or *accepted* by Inland Revenue".

6.2 The Law Society has previously expressed concerns that the automation of discretionary decision-making raises rule of law concerns and could constitute an improper delegation

---

10 This is also suggested under the 'compliance and assurance' principle at page 3.

11 See, for example: Parasuraman and Manzey "Complacency and bias in human use of automation: an attentional integration" (2010) 52(3) *Human Factors* 381; Goddard, Roudsari and Wyatt "Automation bias: a systematic review of frequency, effect mediators, and mitigators" (2012) 19(1) *Journal of the American Medical Informatics Association* 121; Skitka, Mosier and Burdick "Accountability and automation bias" (2000) 52 *International Journal of Human-Computer Studies* 701.

of discretionary powers – in a 2025 report regarding the rule of law, the Law Society observed that:[12]

> The use of ADM systems in government departments raises rule of law concerns about the improper delegation of any discretionary powers which must be exercised by the government by making some form of evaluative judgment. Where a public sector body is given an element of discretion in its decision-making, it must put its mind to the decision and not follow any policy it adopts blindly and inflexibly. The use of ADM systems, which are incapable of exercising evaluative judgement and making merit-based decisions and instead adopt a 'one-size-fits-all' approach (which fetters discretion),[13] could amount to an improper delegation of such discretionary powers.
>
> Even where discretion is built into the machine-learning algorithm, that is done by data scientists and engineers who are unlikely to have the necessary experience and training to adopt administrative enactments, or to specify how a law should be applied – this could be problematic where, for example, ADM systems are to be used to determine and impose sanctions on individuals, in circumstances where there are discretionary elements or judgment that could be exercised.

6.3 We further note that it would not be appropriate to treat *all* legal risks as risks which could appropriately be mitigated or accepted (as suggested in principle 4). There are some legal risks for which the only options should be that they are removed (i.e., with steps being taken to ensure the ADM system does not give rise to that legal risk) or, if that is not possible, the ADM system is not implemented. Administrative law principles offer guidance on the types of legal risks that should not be traded off to achieve other benefits – these include that decisions must be lawful (that is, they should correctly apply relevant legal tests, take into account relevant factors and disregard irrelevant ones, and avoid the rigid application of policies irrespective of individual circumstances), and be fair and reasonable.

6.4 If Inland Revenue opts to either mitigate or accept such legal risks, this could be seen as an informed decision to implement a system that is likely (or is capable) of making unlawful decisions. Such an approach is likely to be inconsistent with section 6(2)(f) of the TAA , which requires those administering the law to do so fairly, impartially, and according to law. While errors can be made even within a human-only decision-making system, no ADM system should be implemented if it is *known* that it will make errors, or that there is a high chance of errors being made in certain circumstances.

6.5 We note with concern that decisions to mitigate and accept the legal risks of deploying an automated methodology at scale contributed to the failure of Australia's Robodebt Scheme and to the United Kingdom's Post Office Scandal. Any suggestions to simply mitigate, rather than remove, legal risks must be approached with caution.

---

12 Law Society *Strengthening the rule of law in Aotearoa New Zealand* (June 2025) (***Rule of Law Report***) at [9.19].

13 We note here that the fettering of statutory discretion is recognised as a ground for judicial review in New Zealand – see, for example, *Criminal Bar Association of New Zealand Inc v Attorney-General* [2013] NZCA 176.

6.6 The Law Society suggests removing the reference to the automation of discretionary decision-making in principle 4. If the intention is to retain this reference despite the concerns we have raised, we recommend:

(a) specifying those who are authorised to accept residual legal risk;

(b) requiring external legal review of any proposals to automate discretionary decision-making powers; and

(c) requiring ongoing monitoring of ADM systems to ensure compliance with legal requirements (as legal risks may become apparent only *after* the system has been implemented).

6.7 Where principle 4 refers to the need to comply with the TAA and the Privacy Act 2020, we also suggest including a requirement to ensure consistency with the *Algorithm charter for Aotearoa New Zealand* (which sets out the government's commitment to algorithm transparency).[14]

7 Principle 5: transparency and accountability

7.1 The Law Society supports the inclusion of the 'transparency and accountability' principle which requires Inland Revenue to publicly disclose decisions that are made by ADM, and how the automated decisions affect outcomes. This principle will serve as an important safeguard, particularly in circumstances where a decision made by ADM disadvantages a taxpayer.

7.2 However, the Law Society is of the view this principle could go further to promote transparency and accountability, and consequently, protect the public's perception of the integrity of the tax system (as required under section 6(2)(a) of the TAA), as we discuss below.

Discouraging the use of black box systems

7.3 While principle 5 currently requires ADM systems to "provide enough information about how a decision is made to enable staff to explain decisions if needed",[15] it does not require disclosure or identification of *how* the system reached a particular decision.

7.4 ADM processes which rely on impenetrable decision-making systems are often referred to as 'black boxes'.[16] The workings of black boxes are mysterious: we can examine their inputs and outputs, but we are unable to explain how one becomes the other. In our 2025 report, we identified several rule of law concerns arising from the use of black box ADM systems within government departments:[17]

> Where black box ADM systems are used by government departments, they can create challenges in the judicial review process by making it difficult to understand how and why a certain decision was made, whether relevant powers

---

14 New Zealand Government and Stats NZ *Algorithm charter for Aotearoa New Zealand* (July 2020).

15 Page 3 of the consultation document.

16 This may be because it uses data which is not accessible to the data subject, and/or it deploys algorithms which are either similarly inaccessible, or are so complex that they cannot be reduced to a series of comprehensible rules and rule applications – see: *Rule of Law Report* at [9.19].

17 *Rule of Law Report* at [9.19].

> were properly exercised in making that decision, whether relevant factors were considered, and whether bias or discrimination in the learning input for the algorithm influenced that decision. Even where the internal workings of an ADM system can be accessed, those workings may not be accessible to others — for example, because of restrictions imposed by intellectual property rights in the algorithm, or concerns about security (such as that publicising the algorithm will allow gaming of the system by malevolent actors).

7.5 For these reasons, the Law Society has recommended that black box ADM systems should not be used in government departments.[18] We remain of this view, and consider that the 'transparency and accountability' principle could go further to discourage the use of black box ADM systems. This could be achieved, for example, by including a requirement for Inland Revenue to collect, and in appropriate circumstances, produce a full record of each decision (including all inputs into decision, decision-making criteria and reasons for decisions).

Reducing gaming risks

7.6 Our 2025 report also cautioned that parties may be able to 'game' ADM systems, for example, by selectively suppressing information or making false or irrelevant factual claims to skew the operation of the ADM system unfairly in their favour.[19] Fixed, rules-based logic applied at scale (as is likely to be the case within Inland Revenue) can be predictable, and predictability could invite structuring through refund-fraud schemes and aggressive positions that are calibrated just under automated thresholds.

7.7 The risk of this occurring could be reduced if the 'transparency and accountability' principle is modified to distinguish between requirements to disclose decision types (which would be appropriate), and requirements to disclose exact rule parameters and thresholds (which, if disclosed, could invite gaming of the system). This distinction could be recognised in a carve-out that is anchored to the confidentiality grounds identified in section 18 of the TAA.

Improving accountability

7.8 The 'transparency and accountability' principle requires the disclosure of decisions that have been made by ADM after they have been "approved by Inland Revenue". This particular safeguard does not identify who would be accountable for approving decisions. We recommend identifying any specific governance bodies or roles that would be responsible for approving the use of ADM, and requiring the creation and publication of a register of the members of the relevant ADM approval body or bodies.

8 Principle 6: human oversight

Statutory right to seek human review of automated decisions

8.1 The 'human oversight' principle requires Inland Revenue to provide a channel for challenging or appealing automated decisions, and for this channel to be clearly communicated to, and easily accessible by, those affected by the decision. The Law Society supports giving individuals the ability to appeal and challenge decisions made

18 *Rule of Law Report* at [9.20].
19 *Rule of Law Report* at [9.19].

using ADM, and recommends this be enshrined in a statutory right to seek a human review of any decision made using ADM.[20] This right could be accompanied by a deeming provision which clarifies that automated decisions are actions of the Commissioner for legal purposes: this would ensure decisions made by ADM are classed as 'disputable decisions' for the purposes of Parts 4A and 8A of the TAA.

Ensuring the effectiveness of review and appeal channels

8.2 Careful consideration should also be given to how Inland Revenue will ensure that its review and appeal channels provide an effective remedy for taxpayers who have been genuinely impacted by an incorrect decision, particularly:

(a) if the channel also receives a large volume of unmeritorious complaints from those who are simply unhappy with the (correct) outcome of a decision they learn has been made using ADM; and

(b) in light of research which indicates that human reviewers may be less rigorous in reviewing decisions made by ADM, when compared to how they review decisions made by other humans.[21]

Relationship with existing statutory rights

8.3 We also note these review channels do not appear to be mapped against existing statutory rights: many automated outputs will be disputable decisions subject to Part 4A (disputes procedures) and Part 8A (challenges) of the TAA. It could therefore be helpful to map each review channel to specific decision types, and to expressly confirm that ADM will never truncate existing statutory disputes rights.

Scope and application of the proposed standards

8.4 Principle 6 identifies the types of ADM systems that are intended to be used within Inland Revenue (and will therefore be subject to the proposed ADM standards and principles). It is unclear why these matters are identified under principle 6, as they trigger the application of *all* principles and standards. Our preference would be to include an exhaustive list of the types of ADM systems that are intended to be used within Inland Revenue as a standalone section, rather than under principle 6. This provision could then also clarify whether the ADM standards apply to decision-support tools, risk-scoring and case-selection engines (which determine who is audited, even where a human makes the final call), correspondence triage, or automated debt-collection actions (noting case selection is precisely where perceived unfairness and disparate impacts are likely to arise).

8.5 We also note the exclusion of predictive systems (also identified under principle 6) may not sit comfortably alongside the requirement to assess whether ADM has the potential

---

20 Such a provision could be inserted into the 2026/27 Annual Rates Bill alongside other provisions relating to Inland Revenue's use of ADM.

21 See, for example: "Anthropic Education Report: The AI Fluency Index" (23 February 2026) <www.anthropic.com/research/AI-fluency-index>'; Klingbeil, Grützner and Schreck "Trust and reliance on AI — An experimental study on the extent and costs of overreliance on AI" (2024) *Computers in Human Behavior,* vol 160; University of Melbourne and KPMG International *Trust, attitudes and use of artificial intelligence: a global study* 2025 (2025); Skitka, Mosier and Burdick "Does automation bias decision-making?" (Int. J. Hum. Comput. Stud., 1 November 1999).

to increase fraudulent activity (as required under principle 4), since fraud detection is typically a predictive exercise. We invite officials to give further thought to whether an exception may be warranted if ADM is to be used for this purpose.

9 Principle 7: compliance and assurance

9.1 The 'compliance and assurance' principle requires regular monitoring proportionate to the risks involved to be carried out to ensure ADM continues to produce expected results. We make the following suggestions to improve the safeguards which sit underneath this principle:

Incident response framework

9.2 Principle 7 requires Inland Revenue to correct issues that are identified, but remains silent on how it should respond to decisions that are either unlawful, or are made following a defective process.

9.3 ADM errors are usually systemic: a single flaw within an ADM system has the potential to generate thousands of incorrect assessments before it is detected. The Law Society therefore queries whether it is appropriate to require Inland Revenue to correct an issue "as soon as practicable, with Inland Revenue accepting all risks involved until this occurs" – this suggests that the Department could knowingly continue to operate a faulty ADM process over an indefinite period of time: doing so can cause significant damage to the public's perception of the integrity of the tax system.

9.4 We therefore recommend including a requirement to create an incident response framework, which clearly identifies how Inland Revenue would be expected to respond to any large-scale errors or ADM failures. This framework should, among other things, require:

(a) the proactive identification of affected taxpayers;

(b) notification of the error or failure;

(c) correction of past decisions;[22] and

(d) identification of defined suspension triggers.

More regular reviews

9.5 Principle 7 requires Inland Revenue to undertake compliance reviews at least once every three years, or more frequently. While this principle contemplates the possibility of more regular reviews, the specific reference to a three-yearly review is likely to establish a benchmark expectation.

9.6 It would be preferable for this benchmark to be set at a lower level (say, one year) in the early days of implementation, with the possibility of less regular reviews further down the track if appropriate. This would enable the earlier identification of issues (including those which arise from high-risk uses of ADM) and encourage more prompt responses –

---

22 We note here that the amendment power in section 113 of the TAA would be the natural vehicle for making such corrections.

all of this would go some way towards protecting the integrity of the tax system, and the public's perception of the tax system.

Versioned recordkeeping

9.7 Principle 7 does not require Inland Revenue to keep 'versioned' records of ADM rules that are in force at a given time. We recommend including such a requirement to ensure compliance with the Public Records Act 2005, and to enable Inland Revenue to defend challenges about what rules were applied to historic decisions.

Equitable outcomes for Māori and other groups

9.8 Principle 7 requires consideration of whether ADM processes are producing disproportionate, unexpected or inequitable outcomes for Māori "where this can be assessed lawfully and appropriately". The inclusion of this caveat risks watering down this obligation to ensure ADM does not produce disproportionate, unexpected or inequitable outcomes for Māori. We therefore recommend deleting the phrase "where this can be assessed lawfully and appropriately" from this safeguard.

9.9 We also recommend requiring Inland Revenue to meaningfully engage with Māori before *and* after implementing an ADM system (that is, during the design stage, as well as post-implementation monitoring of the ADM system).

9.10 It is also important to ensure that ADM systems do not result in disproportionate outcomes for other taxpayer groups (for example, children or disabled people). We invite officials to consider whether principle 7 should be broadened to ensure these groups are also protected from disproportionate, unexpected or inequitable outcomes.

10 Other matters

10.1 In addition to the issues canvassed above, the Law Society has also identified several other matters which are not addressed in the proposed principles, which we discuss below. We consider it would be appropriate to include a further, overarching principle of 'proportionality and responsible adoption of ADM' (or similar), which addresses each of the points we have identified below. Alternatively, officials could look to address these matters in additional standalone principles which will sit alongside the principles identified in the consultation document.

Security considerations

10.2 The proposed principles do not address safety and security issues arising from the use of ADM despite the fact that ADM systems routinely deal with sensitive taxpayer data, and . The Public Service AI Framework includes the following principle which focuses solely on safety and security:[23]

> Public Service AI systems should treat the security of customers and staff as a core business requirement, not just a technical feature (security-by-design). They should minimise risk to individual or national safety and security under normal use, misuse or adverse conditions.

---

23 New Zealand Government *Public Service Artificial Intelligence Framework* (July 2025) (**Public Service AI Framework**).

> The Public Service should ensure traceability of data, apply a robust risk management approach and work collaboratively with commercial and security colleagues in the procurement and assurance of AI tools.

10.3 We recommend including a similar requirement in the proposed standards, which:

(a) is consistent with the confidentiality requirements set out in section 18 of the TAA;

(b) safeguards against safety and security issues arising from the use of ADM; and

(c) addresses (for example) security-by-design, data traceability, and vendor/procurement assurance.

Environmental impacts

10.4 The proposed principles and safeguards do not recognise or seek to mitigate any adverse impacts that ADM may have on the environment. In contrast, these matters are addressed in several other documents relating to the use of ADM and AI – these include:

(a) The Public Service AI Framework: the very first principle of this framework ('Inclusive, sustainable development'), which provides that public service AI systems should "contribute to inclusive growth and sustainable development through a focus on innovation, efficiency and resilience, and on reducing economic, social, gender and other inequalities and protecting natural environments".

(b) The Cabinet Paper which sets out the Government's approach to work on Artificial Intelligence, which notes that "future work could consider issues of emissions, sustainability, and AI use".[24]

(c) The AI Forum's AI Principles, which call on stakeholders to design, develop and use AI systems to promote (among other things) the wellbeing of New Zealand's environment, including in the area of sustainability.[25]

10.5 It could be useful to include an additional safeguard which requires consideration of the environmental impacts of any ADM system Inland Revenue proposes to use (including in terms of its energy use, computational resources, and carbon footprint). Such an assessment could be undertaken as part of any cost-benefit or proportionality analysis prior to the implementation of an ADM system.

Social licence

10.6 A 2024 Cabinet paper on New Zealand's approach to AI notes that, in 2023, over 66% of New Zealanders were extremely or very concerned about AI being unregulated, used for malicious purposes, or having unintended consequences, and goes on to warn that "social licence is hard to build but easy to lose". We suggest this is of particular importance in the area of tax administration, given the historical origins of taxation and representative democracy going hand in hand.

---

24 Hon Judith Collins KC *Approach to work on Artificial Intelligence* (June 2024) at [33].

25 AI Forum New Zealand *Trustworthy AI in Aotearoa AI Principles* (March 2020) at page 4.

10.7 While the proposed standards include transparency and accountability principles, and acknowledge that public perception of ADM will evolve over time,[26] they do not engage with the concerns identified in the 2024 Cabinet paper, or directly address the need to have and maintain social licence to use ADM, and to foster public acceptance and trust of ADM systems used by Inland Revenue.

10.8 Given section 6(2)(a) of the TAA requires the government to protect public perceptions around the integrity of the tax system, we consider it would be appropriate to include an additional safeguard which requires Inland Revenue to understand whether the public finds Inland Revenue's use of ADM acceptable, and whether Inland Revenue holds a social licence to use ADM. This would require (among other things) establishing mechanisms for seeking public input regarding the use of ADM within Inland Revenue (noting there is currently no mechanism for gauging public sentiment or trust over time beyond complaints channels).

10.9 We also note that while 'shifts in public expectation' are identified as a trigger for reviewing the standards, these could be elevated to a standalone consideration that applies to the use of ADM on an ongoing basis.

11 Next steps

11.1 We hope this feedback is helpful. Please feel free to get in touch with me via the Law Society's Senior Law Reform & Advocacy Advisor, Nilu Ariyaratne (Nilu.Ariyaratne@lawsociety.org.nz) if further discussion would assist.



David Campbell
President

[26] At page 1.