frazerwalker

14 June 2026

privacyreformtranche1@oaic.gov.au

Office of the Australian Information Commissioner

GPO Box 5288

Sydney NSW 2001

Frazer Walker Submission to the Office of the Australian Information Commissioner

Please find attached our submission to the Automated Decision-Making Transparency Obligation (APP 1) with responses to the consultation questions. Frazer Walker would welcome the opportunity to discuss this submission, or to assist the OAIC further as it develops the guidance.

Confidentiality: This submission contains no confidential information and no personal information. FrazerWalker consents to its publication by the OAIC.

[redacted],

Partner

Frazer Walker

M: [redacted]

E: [redacted]

ABOUT FRAZER WALKER

Sydney-based independent management & technology consultancy

Frazer Walker is a Sydney-based independent management and technology consultancy, established in 2002. We help Australian and New Zealand organisations realise their strategies through practical and high-quality advice and delivery.

Our work integrates four practices delivered as a single offering:

1. Data — data strategy, governance, management, analytics, and AI readiness.
2. Artificial Intelligence — AI strategy, governance, ML operations, and assurance.
3. Technology — technology strategy, system selection and modernisation, governance, and hands-on delivery.
4. Managing Transformation — portfolio and program management, independent assurance, and execution-focused delivery.

We are vendor-independent and practitioner-led, with deep specialist knowledge in the Insurance sector. We have served more than 85 organisations across Australia and New Zealand and maintain evergreen Professional Service Agreements with our long-term clients.

We are members of the International Association of Privacy Professionals (IAPP) with CIPT accredited staff.

Level 16, 175 Pitt Street, Sydney NSW 2000

T +61 2 9283 4825

E enquiries@frazerwalker.com

W www.frazerwalker.com

EXECUTIVE SUMMARY

We welcome the opportunity to respond to the Office of the Australian Information Commissioner's (OAIC) Issues Paper, and we support both the objective of the automated decision-making (ADM) transparency obligation and the OAIC's development of guidance ahead of its commencement on 10 December 2026.

Our perspective is formed from practical exposure to how ADM is procured, configured and operated within regulated financial-services environments, in particular the increasingly common situation in which ADM capability is embedded within international third-party core systems and platforms that the obligated entity neither built nor fully controls. Four themes run through our responses:

1. **Human oversight is not an exemption by itself.** Whether a system is "substantially and directly related" to a decision should be founded on its actual causal influence on the outcome, not on a nominal "human-in-the-loop" (HITL) label.
2. **Disclose not only the existence of ADM, but also its assurance.** A responsible entity should be able to show that it has assessed the tool and its vendor, against a recognised framework for bias and transparency, and should disclose that assurance, much as organisations now disclose SOC 2 or ISO/IEC 27001 for information security. While such frameworks can assure governance and bias-assessment methods, it will prove much harder to assess "fairness."
3. **Third-party and supply-chain ADM is where the obligation is most likely to fail in practice.** The "arranged for" boundary, vendor assurance and the commercial-in-confidence carve-out will require specific guidance.
4. **Transparency is a necessary precondition for other protections, but not by itself sufficient.** Disclosure is what enables anti-discrimination and consumer-law remedies to operate; the threshold to disclose should not be confused with the threshold for unlawfulness.

We have answered all questions, with greater depth on those closest to our expertise. Question numbers follow the OAIC submission form, with the Issues Paper question number shown in brackets. Where a ranking is requested, we have rated each item by the weight we consider it should carry (5 = most relevant, 1 = least relevant).

RESPONSES TO THE CONSULTATION QUESTIONS

Question 1 - Substantially and directly related to making a decision [Issues Paper Q1]

Factor / example	FW weighting (1–5)	Rationale
Degree of reliance on the ADM output	5	This is core to the system's causal influence on the decision outcome.
Integration of ADM into the decision-making workflow	5	If the decision cannot realistically be made without the system's output, the connection is substantial and directly related.
Ability and likelihood of human override	5	Although the Human in the loop ability is there to override the decision, the "likelihood" of this happening with ADM reduces over time, becoming a theoretical override rarely exercised. This then renders oversight not meaningful.
Nature of output (advisory vs determinative)	4	Relevant, however the "advisory" label understates the reality of ADM.
Transparency and explainability of outputs	3	While important for disclosure, an AI 'black box' can still wholly drive a decision.

Reliance, workflow integration and the likelihood of override are the most probative factors because they go to the system's actual influence on the outcome. We urge the guidance to read "ability and likelihood of human override" with emphasis on *likelihood*: the appropriate measure is the observed override or disagreement rate, not the mere existence of an override control.

"Advisory versus determinative" should be treated with caution. A system characterised as "decision-support" can be determinative in practice. For that reason we propose an additional factor: the practical conditions under which human review actually occurs. The established human-factors literature identifies two effects that systematically erode the value of human oversight:

- Automation bias: the well-documented tendency to over-rely on automated outputs, accepting them even when incorrect and the error is detectable, because deference is locally rational when the system is usually right.
- The vigilance decrement: the degradation of sustained attention on low-event monitoring tasks. The rarer the error requiring intervention, the faster and more severe the decay in the reviewer's ability to detect it.[1]

[1]See, for example, the foundational work on the vigilance decrement (Mackworth, 1948) and the wider human-factors literature on sustained attention in low-event monitoring.

Compounding these, reviewer time pressure, decision volume and skill decay further reduce the substance of oversight.[2]

Accordingly, the guidance should direct assessment to whether the human exercises genuine, independent judgement or substantively ratifies the system's output. The greater the reliance, and the weaker the practical conditions for independent review, the more clearly the program is "substantially and directly related" to the decision.

Question 2 - Fictional edge case: generative AI with human oversight (NACIA / GPTea) [Issues Paper Q2–3]

Response: Yes.

GPTea performs two functions, both of which are substantially and directly related to the eligibility decision.

- The recommendation is plainly a key factor facilitating the human decision (substantial) and is directly connected to it (direct).
- The summarisation is also substantial. By selecting and compressing the candidate profile, the model determines what the officer sees and how it is framed. This control of the information frame (an anchoring effect), combined with automation bias, means the system materially shapes the outcome even where a human formally decides.

The assurance that staff "always have oversight and make the final decision" is precisely the framing that should not, on its own, remove the system from scope (see Question 1). Generative AI raises the significance further: such systems are non-deterministic. The same profile may yield materially different summaries or recommendations on different occasions, and can produce fluent but incorrect or fabricated content (hallucinations). Affected individuals cannot assume consistency or accuracy, which strengthens the case for disclosure.

The cohort (elderly people seeking life-changing aged-care support, many of whom will be vulnerable) and the nature of the decision (access to a significant service) reinforce our view. NACIA's use of GPTea is substantially and directly related to making the decision; the ADM obligation applies, and NACIA should disclose the kinds of personal information used and the kinds of decisions made.

Question 3 - Factors increasing the likelihood a decision affects rights or interests [Issues Paper Q4]

Factor / example	FW weighting (1–5)	Rationale
Sensitive information	5	Heightens both the potential for harm and the reasonable expectation of disclosure.

[2]On over-reliance on automation, see, for example, Parasuraman & Manzey (2010), "Complacency and bias in human use of automation."

Factor / example	FW weighting (1–5)	Rationale
Vulnerable persons	5	The same decision can have a materially greater impact (see Question 4).
Financial outcomes	5	Direct, measurable effect on a person's circumstances.
Intrusive practices	4	Indicative of significance, particularly where profiling is involved.

We recommend the guidance also recognise the following factors:

- irreversibility or difficulty of reversing or contesting the outcome;
- opacity: the individual's inability to understand or challenge the basis of the decision;
- essentiality of the good, service or support affected; and
- downstream and third-party effects. A decision may rest on the characteristics or behaviour of others (for example, inferences drawn from people in the same postcode or cohort), so that an individual is affected by data that is not their own, as per the credit-limit example in the Issues Paper.

Question 4 - Classes of vulnerable persons [Issues Paper Q5]

We recommend a 'functional' style test which is more contextual, rather than a fixed list of vulnerable persons. A person is vulnerable where they face a heightened susceptibility to harm from a decision, and/or a reduced capacity to detect, understand or contest this.

Classes of persons that will commonly meet the test include children and young people; older people; people with disability (including cognitive or psychosocial disability); people experiencing financial hardship; people with low digital or financial literacy or access; culturally and linguistically diverse communities; Aboriginal and Torres Strait Islander communities; and people in crisis (for example, family violence or acute health events).

We think the 'functional' test will be better as exhaustive list will date and may create future service gaps. The idea of a principle of "reasonably foreseeable heightened impact" is more suatainable.

Secondly, ADM can itself create situational vulnerability. For example, a person adversely flagged by an 'black box'/opaque system, or targeted because of inferred financial distress (such as in the high-interest-loan example in the Issues Paper). Thus vulnerability should be assessed in light of the decision, not only on the person's fixed characteristics.

Question 5 - Significant service or support [Issues Paper Q6]

Factor / example	FW weighting (1–5)	Rationale
Essential banking and credit services	5	Access to credit and payments is foundational to all citizens.
Financial assistance	5	Direct bearing on a person's means.
Essential utilities	5	Energy and water are essential to health and safety.
Physical assistance / support (incl. healthcare)	5	Significant for health and wellbeing.
Telecommunications	4	Increasingly essential to access services and employment.
Educational services	4	Significant, with longer-term effects on opportunity.
'In-home services'	4	Significant for those who rely on them, particularly vulnerable groups of people.

Insurance as an illustration of the significance test.

Consistent with our view above that significance is best assessed functionally rather than from a closed list, we offer insurance as an example of the principle rather than as a new category. The Explanatory Memorandum already cites a life-insurance contract, and the Issues Paper addresses the General Insurance Code of Practice (GICOP). Key insurance decisions such as underwriting, eligibility, pricing and claims, bear directly on the significance factors the obligation raises, such as financial circumstances, rights under a contract, and access to protection. They are increasingly automated within insurers' core systems and therefore transparency is mandatory.

With regards Insurance, we should add that it should capture *material* decisions only, not every insurance-related calculation. This is the "substantially related" threshold discussed at Question 1.

Secondly, recognising these decisions as significant engages only the transparency obligation. It does not question or constrain risk-based pricing, which is a core element of insurance and which remains lawful and regulated. As we have raised in Question 7, disclose is not the same as unlawful. We make this point as an independent adviser working across the insurance supply chain, in support of the obligation's consumer-transparency purpose rather than any particular commercial interest. Housing and tenancy (for example, automated tenancy screening) and employment and income decisions illustrate the same principle.

International practice reinforces our view of an effect from the decision rather than an "essential service" tag. Comparable regimes reach insurance through an effects test: the European Union and the United Kingdom capture any solely-automated decision with a legal or similarly significant effect, including underwriting, pricing and claims, rather than by designating insurance "essential." Where a jurisdiction does single insurance out for heightened treatment, it is life and health insurance specifically: the EU AI Act classifies AI used for risk assessment and pricing in life and health

insurance as high-risk, while general (property and casualty) insurance is not listed and claims handling is treated as a grey area. We therefore suggest the guidance treat material insurance decisions as significant by their effect, while recognising that life and health insurance attract heightened scrutiny internationally. The United States takes a different approach, regulating insurer automated decision-making largely through insurance-specific regimes. For example, the NAIC Model Bulletin on the use of AI by insurers, adopted across many states in the USA.

Question 6 - Other legal frameworks the OAIC should consider [Issues Paper Q7]

- The General Insurance Code of Practice (GICOP) and the Life Insurance Code of Practice (LICOP), and the broader financial-services bodies requirements, such as ASIC and APRA requirements, the Design and Distribution Obligations, and the unfair contract terms regime. These should also be considered given how often ADM sits within regulated financial products.
- Commonwealth and State/Territory anti-discrimination law, particularly indirect discrimination, which is directly engaged by the use of proxy variables.
- The Australian Consumer Law (misleading or deceptive conduct; unconscionable conduct), which is particularly relevant to differential pricing.
- The Attorney-General's Department's work on government use of ADM, and the findings of the Robodebt Royal Commission on the harms of opaque automated decision-making and the value of transparency and contestability.
- The EU AI Act (high-risk classification) and GDPR Article 22 and Recital 71, already noted in the Issues Paper, and the WA Privacy and Responsible Information Sharing Act 2024 (Principle 10), for national coherence.
- The Australian Government's Policy for the responsible use of AI in government and the DTA transparency-statement standard, so that public-sector privacy-policy disclosures and AI transparency statements remain consistent.
- Australia's Voluntary AI Safety Standard and the December 2025 National AI Plan (including the new Australian AI Safety Institute), which align domestic practice to AS ISO/IEC 42001 and the NIST AI Risk Management Framework. As the National AI Plan confirms there will be no standalone Australian AI Act in the near term, the privacy-policy disclosure becomes one of the few binding transparency hooks and should be anchored to these convergent frameworks (see Question 10).
- International AI-assurance standards as reference points: ISO/IEC 42001 (AI management system), ISO/IEC TR 24027 and ISO/IEC TS 12791 (assessing and treating bias), and ISO/IEC 27701 and ISO 31700 (privacy information management and privacy-by-design).

International approaches to ADM disclosure.

The OAIC's obligation can be calibrated against several comparable regimes, which are converging on a common disclosure core even as their enforcement models differ:

- European Union: GDPR Articles 13 to 15 require disclosure of solely-automated decision-making with legal or similarly significant effects, including meaningful information about the logic involved and the significance and envisaged consequences. The Court of Justice (SCHUFA) confirmed this means more than a generic "we use AI." The EU AI Act adds high-risk obligations, including for life and health insurance risk assessment and pricing, from August 2026.

- United Kingdom: the Data (Use and Access) Act 2025 replaced Article 22 with new Articles 22A to 22D (in force February 2026), shifting from a prohibition to a permission-plus-safeguards model: transparency before the decision, a right to obtain human review, and a right to contest it.
- Singapore: the Personal Data Protection Act's notification and accountability obligations, supplemented by the PDPC's advisory guidance on AI decision systems, encourage disclosure of the system's function and the measures taken to achieve fairness; financial-services AI carries additional sectoral guidance.
- United States: there is no federal rule. California's CCPA automated-decisionmaking-technology regulations (effective 2026, with significant-decision obligations phasing in from 2027) require a pre-use notice covering the decision, the personal information used, and opt-out and access rights, with insurer decision-making largely regulated separately by insurance authorities.

What is common for all the regimes listed above is the existence of the ADM, the need for meaningful information about its logic, its significance and likely consequences, and a pathway to a human. These four elements we think, are a useful benchmark for the OAIC's guidance. A disclosure built with these elements in mind would satisfy the Australian obligation while remaining portable across jurisdictions.

Question 7 - Fictional edge case: differential pricing (Marco / Daintree) [Issues Paper Q8–9]

Has Marco's interest been significantly affected?

In our view, differential pricing of this kind *can* significantly affect an individual's interests. However we would caution against framing the test as a fixed dollar amount or percentage. On these facts the price to Marco is nearly three times that paid by another customer for the same product, set using his address as a proxy for wealth, and he cannot afford the book at the inflated price. That is capable of significantly affecting his interests, particularly because the pricing turns on personal information used by a computer program.

Should there be a fixed percentage threshold? No.

We do not support a fixed threshold (such as 5% or 20%). A fixed threshold means very different things across products and across individuals' circumstances; it invites avoidance by pricing just below the line; and it substitutes a mechanical test for the contextual judgement the provision calls for. We recommend significance be assessed against four factors:

1. the essentiality and importance of the product or service;
2. the magnitude of the differential relative to the baseline price and to the individual's circumstances (a relative, not absolute, measure);
3. whether the differentiating input is a proxy for a sensitive or protected attribute. Note that geographic location frequently correlates with protected and historically disadvantaged groups, so postcode-based pricing carries elevated risk; and
4. transparency and avoid-ability: whether the individual can know about, understand or avoid the differential.

Importance of the two listed factors (scale 1–5)

Factor / example	FW weighting (1–5)	Rationale
The dollar amount of the differential	2	This is meaningful only relative to the baseline price and the individual's circumstances.
The significance of the product or service	5	The critical element of whether interests are significantly affected.

Key point: keep 'disclosable' separate from 'unlawful'.

For a *transparency* obligation, we would encourage the OAIC to set a low trigger and to keep two questions separate: whether a practice is significant enough to require disclosure is different from, and lower than, whether it is ultimately lawful. Transparency is the mechanism that allows an individual to discover differential pricing and to invoke other frameworks such as anti-discrimination law and the Australian Consumer Law. Treating "disclosable" as equivalent to "unlawful" would set the threshold too high and defeat the obligation's purpose. Differential pricing for a discretionary product such as a book will often be less significant than for an essential good or service; but where the pricing variable is a proxy for a protected attribute, even a discretionary purchase may warrant disclosure because of the discrimination risk it signals.

Question 8 - Fictional edge case: discriminatory targeted job ad (Rafqa / Manscon) [Issues Paper Q10]

Response: Yes.

The statutory definition expressly provides that "making a decision includes refusing or failing to make a decision" and "doing a thing includes refusing or failing to do a thing." The algorithm not providing the information to Rafqa is the system *failing to do a thing*, in this case, failing to surface the opportunity to her. This omission is substantially and directly related to an outcome affecting her interests.

The effect is on access to employment opportunities, an example expressly given in the Explanatory Memorandum. The prioritisation operates on gender, a protected attribute, and is a measure evaluating personal aspects of the individual (cf. GDPR Recital 71). Critically, the fact that Rafqa never sees the advertisement and may never know it existed, is the reason the transparency obligation matters. Decisions that take the form of invisible suppression are the hardest for an individual to detect or contest; reading "decision" to include "failing to do a thing" is essential to capture algorithmic targeting and suppression.

This raises the "arranged for" question addressed at Question 9. In our view Manscon has arranged for the targeting (it chose to advertise via Opaque Window's algorithmic promotion) and bears the disclosure obligation; Opaque Window operates the program. Guidance should confirm that an entity which elects to use a platform's algorithmic targeting feature to reach or screen individuals has "arranged for" that ADM, notwithstanding that it did not build or fully control the algorithm.

Question 9 - 'Arranged for': scenarios requiring guidance [Issues Paper Q11]

Yes. In our advisory experience, the "arranged for" boundary is the area most in need of guidance, because ADM is increasingly embedded within third-party core systems, SaaS platforms and general-purpose models that the obligated entity neither built nor fully understands. We recommend the guidance address:

1. Layered supply chains - where a vendor's system itself relies on a sub-processor's or third party's model. The guidance should clarify how far the obligated entity's duty to identify, assess and keep oversight extends down the chain.
2. Off-the-shelf versus configured systems - where the entity sets the decision rules, thresholds or weightings within a procured system, it is in substance arranging for the ADM even though the underlying software is a vendor's. This should be confirmed.
3. Platform targeting features (as at Question 8) - using a platform's algorithmic targeting or ranking to reach or screen individuals constitutes arranging for ADM.
4. General-purpose and generative models accessed via API - where an entity embeds a third-party model into a decision process, it cannot fully characterise the model's behaviour, and the model is non-deterministic. The guidance should address how such entities can meaningfully disclose the kinds of personal information used and decisions made despite limited visibility.
5. Contractual detail - entities will need transparency information from vendors to comply. The guidance should set expectations that procurement and contracts secure these disclosures, and address what an entity should do where a vendor will not provide them.

Vendor assurance: the practical mechanism.

Where ADM is procured, the most workable way for an entity to satisfy itself that the tool has been assessed for bias, fairness and transparency is recognised third-party assurance of the vendor. ISO/IEC 42001 - the AI-governance management-system standard, and the direct analogue to ISO/IEC 27001 for information security - is now third-party certifiable and is rapidly becoming a procurement signal in its own right. We recommend the guidance encourage entities to obtain and evidence such assurance, and to flow the corresponding transparency information down through contracts, exactly as security certifications already are. We expand on how this should feature in disclosure at Question 10.

The commercial-in-confidence carve-out is a material risk.

We draw particular attention to the risk that the commercial-in-confidence exclusion becomes a loophole. In a market where ADM is supplied by vendors, there is a clear incentive to over-claim commercial-in-confidence to avoid disclosure, which would hollow out the obligation. We recommend the guidance confirm that (i) the exclusion is to be read narrowly, limited to genuinely proprietary technical detail such as source code or model weights; and (ii) it does not excuse disclosure of the *kinds* of personal information used or the *kinds* of decisions made which is the substance the obligation actually requires.

Question 10 - Extent of disclosure [Issues Paper Q12]

We support the five characteristics the OAIC has proposed and suggest the following additions:

1. System type: disclose, at a useful level, whether the ADM is rule-based or uses machine-learning or generative AI. This matters because such systems are non-deterministic and may not produce reproducible or fully explainable outputs; individuals should not be left to assume consistency, or that an explanation fully accounts for an outcome.

2. The assurance framework applied: disclose the recognised framework against which the entity (and its vendor) has assessed the tool for bias, fairness and transparency, and the assurance status (for example, vendor ISO/IEC 42001 certification). This is the transparency analogue to disclosing SOC 2 or ISO/IEC 27001 for information security: a recognised, comparable trust signal. It should be framed as "assessed against [framework]" rather than as a guarantee of fairness (see below).
3. The real role of human oversight: where an entity relies on human involvement, the disclosure should honestly describe that involvement rather than asserting a bare "all decisions are reviewed by a person." Given the automation-bias and vigilance considerations at Question 1, an unqualified assurance of oversight can mislead.
4. Layered disclosure: a short, plain-language headline with the ability to access further detail, consistent with the ICO / Alan Turing Institute model referenced in the Issues Paper.
5. Currency and versioning: disclosures should be dated and updated when the ADM or its inputs materially change.
6. Signposting contestability: although APP 1 imposes a transparency obligation and not a contestability right, we encourage the guidance to recommend that entities signpost the avenues available to seek human review, further information, or to complain (including, as relevant, the General Insurance Code of Practice, the Australian Human Rights Commission and review bodies). This is an area where the Australian obligation is narrower than WA Principle 10 (which includes a right to request human intervention) and the GDPR; voluntary signposting would partly close that gap and serve the obligation's purpose of enabling individuals to act.
7. Guard against the commercial-in-confidence carve-out undermining meaningful disclosure (see Question 9).

An illustrative direction: anchoring disclosure to recognised frameworks.

We do not suggest the OAIC mandate any single standard. But to make the assurance disclosure above concrete and comparable across entities, the guidance could point to the frameworks that are already converging internationally and domestically. By way of illustration:

- AI governance: ISO/IEC 42001 - the first certifiable AI management-system standard, structured around documented governance, risk assessment, testing and human oversight. Like ISO/IEC 27001 for security, it certifies that the system is managed, not that any model is accurate or fair.
- Bias assessment and treatment: ISO/IEC TR 24027 and ISO/IEC TS 12791, which set out recognised methods for measuring and treating bias across the AI lifecycle.
- Privacy: ISO/IEC 27701 (privacy information management) and ISO 31700 (privacy-by-design), which integrate with the above through the shared ISO management-system structure.
- Domestic anchor: Australia's Voluntary AI Safety Standard, explicitly aligned to AS ISO/IEC 42001 and the NIST AI Risk Management Framework, which already frames supply-chain information-sharing and procurement expectations.

One caution is essential. These frameworks assure *process and bias-assessment method*; they do not certify an outcome as "fair." Unlike information security, which has a relatively singular objective, fairness is plural: the common fairness metrics are mutually incompatible, and the choice between them is a value judgement rather than a technical one. The OAIC should therefore ensure any such disclosure is expressed as "assessed against [framework] using [recognised methods]" and should be wary of permitting a "fairness certified" claim that could mislead individuals.

LOOKING AHEAD: A LAYERED ADM ASSURANCE LABEL

As a forward-looking measure, we encourage the OAIC and the Australian AI Safety Institute to explore a voluntary, layered assurance label for ADM; something that is a concise, comparable indicator (for example, the sensitivity of the personal information used and the significance of the decision, with an optional governance-maturity tier) that sits at the head of a privacy-policy disclosure and links through to the full detail.

Such a label would let citizens and regulators gauge an ADM use immediately, much as energy ratings, vehicle-safety stars and cyber-trust marks do, and would operationalise the assurance-disclosure approach described above.

To be credible it should be voluntary and integrated to existing standards: AS ISO/IEC 42001, the Voluntary AI Safety Standard, and comparable schemes such as Singapore's AI Verify, rather than creating something new and novel in itself. It should also remain an index to ADM assurance, noting the difficulty of certifying 'fairness'. As systems change and evolve over time, the ADM should be re-assessed appropriately. We would welcome the opportunity to contribute to that work and have set out a fuller concept separately to this submission.