Submission to the OAIC: Automated Decision-Making Transparency Obligation (APP 1) Issues Paper

Submitter: 40 South Pty Ltd (ACN 698 132 188), 40south.au **Date:** 15 June 2026 **Submitted to:** privacyreformtranche1@oaic.gov.au **Publication:** This submission may be made public. It contains no confidential information.

About 40 South

40 South is an Australian compliance-technology company based in Wyoming, New South Wales. We work with organisations in regulated industries on the operational visibility and evidence they need when they use AI systems in their business.

We make this submission because the ADM Transparency Obligation intersects directly with that work. It is a policy contribution, drawn from our practical experience of how Australian organisations deploy and oversee AI in decisions that affect people.

General observations

We strongly support the OAIC's direction of travel in the Issues Paper. Three points before the question-by-question response.

Disclosure accuracy depends on operational visibility. APP 1.3 already requires the privacy policy to be clearly expressed and up to date. The ADM disclosures under APP 1.7-1.8 sharpen this: an entity cannot truthfully describe the kinds of personal information used in its ADM, or the kinds of decisions made, unless it can see what its AI systems are actually doing. We encourage the OAIC to frame disclosures as living statements that must be supported by an operational basis, not as point-in-time statements written from generic templates.

Scope should be broad and behavioural. The economic and social value of the transparency obligation depends on capturing AI uses that materially shape outcomes for individuals, not only systems labelled "automated decisions" by their operators. We support the broad reading of "computer program", the inclusion of advisory and recommender AI under "substantially and directly related", and the inclusion of third-party AI under "arranged for". A narrow reading would push ADM into the most accountability-light corner of an entity's operations: tools labelled "advisory", systems run by third parties, and outputs framed as suggestions to a human.

Synergy with adjacent obligations matters. The Issues Paper rightly observes the synergy between APP 1.7 and APP 8 (cross-border disclosure), CPS 234 (third-party information security), the FOI Information Publication Scheme, anti-discrimination law, and the General Insurance Code. The OAIC's guidance will land best if it explicitly cross-references how an entity's ADM transparency disclosures interact with these obligations, so APP entities aren't left to reconcile them in isolation.

Responses to questions

Question 1: Substantially and directly related to making a decision

Ranking of the listed factors (1-5 where 5 is most important):

- Degree of reliance on the ADM system output: **5**
- Integration of ADM into the decision-making workflow: **5**
- Nature of the output (advisory vs determinative): **4**
- Ability and likelihood of human override over an ADM decision: **4**
- Transparency and explainability of outputs: **3**

Additional factor we propose: the **operational and evidentiary basis** an entity has for understanding what the ADM system is actually doing in production. Where an entity has limited insight into the system's inputs, outputs, or behaviour, it should err toward treating the system as in scope, because it cannot demonstrate that the system is not substantially related to a decision.

Free text comments. The "advisory vs determinative" distinction is useful but should not be treated as decisive. The literature on automation bias is consistent: humans defer heavily to algorithmic suggestions, especially under time pressure, cognitive load, or with junior staff. An ADM output framed as a "recommendation" can be substantially related to a decision even where a human nominally signs off. The OAIC's NACIA edge case is a clean example. We would encourage the guidance to set the default toward inclusion where the ADM output is presented to the human decision-maker prior to their decision and is plausibly relied upon, rather than requiring evidence of actual reliance, which is hard for individuals or regulators to observe after the fact.

We would also encourage the OAIC to caution against entities relying on workflow design choices (e.g., requiring a human to click "accept" on an AI recommendation) as a basis to exclude the ADM from scope. The substantive question is whether the ADM materially shaped the decision, not whether a human pressed a button.

Question 2: NACIA edge case (generative AI with human oversight)

Yes, NACIA's described use of GPTea is substantially and directly related to making decisions.

Reasons.

- The chatbot is being used to summarise candidate profiles and provide eligibility recommendations. Both activities materially shape what the human decision-maker sees and considers. A staff member's understanding of the candidate is partly constructed by the chatbot's summary; the recommendation is then a focal anchor for the decision.
- Aged-care eligibility is a benefit granted under legislation and clearly falls within the APP 1.9(d)(i) and (iii) examples (benefits, significant services or support).
- "Human oversight" in the workflow does not in itself remove the chatbot from scope. A human can have nominal oversight and still rely heavily on the AI output. The question is whether the AI substantively contributed to the decision, and on the facts described, it did.
- The case is closely analogous to the EU GDPR guidance treatment of advisory systems that significantly shape human decisions: such systems are within Article 22 even where a human is in the loop, where the human role is not meaningful, independent oversight.

If the OAIC were to treat NACIA's GPTea use as out of scope, almost any AI use in regulated

industries could be brought out of scope by adding a perfunctory human-approval step. That would defeat the obligation's purpose.

Question 3: Factors increasing the likelihood that a decision affects rights or interests

Ranking (1-5):

- Sensitive information: **5**
- Vulnerable persons: **5**
- Financial outcomes: **5**
- Intrusive practices: **4**

Additional factors we propose:

- **Reversibility of the decision.** Decisions that are difficult or impossible to reverse (e.g., a withheld treatment, a missed employment opportunity, a publicly recorded refusal) should be treated as more likely to significantly affect rights or interests.
- **Scale of repetition.** A decision pattern applied at scale across a population (e.g., differential pricing, targeted advertising) is more likely to significantly affect rights or interests, even where any individual instance is modest, because the aggregate effect on access, opportunity, or pricing across a group is material.
- **Information asymmetry.** Where the individual cannot reasonably know they were subject to ADM, the likelihood of significant effect should be treated as higher because the individual cannot exercise other rights of inquiry or contestation.
- **Cumulative effect with other ADM.** Many individuals are now subject to multiple overlapping ADM systems (credit, insurance, employment, advertising). The OAIC may wish to flag that significance should be assessed in context, not in isolation.

Question 4: Classes of person considered vulnerable

We suggest the OAIC adopt a contextual definition of vulnerability rather than an exhaustive class list. The Australian Consumer Law and the Banking Code of Practice already use a context-led concept of vulnerability, which travels well to ADM.

For illustration, we would encourage the guidance to name at minimum:

- Children and young people under 18.
- Individuals experiencing financial hardship, including those in formal hardship arrangements with creditors or utilities.
- Individuals experiencing family or domestic violence.
- People with a disability, mental health condition, or chronic illness that affects access to services.
- First Nations Australians, particularly in remote communities where service access is already constrained.
- People with limited English proficiency.
- People in elder care, hospice, or aged-care settings.
- People exiting custody, refugees, and asylum seekers.

The reason for naming these groups explicitly is that ADM models trained on majority-population data routinely under-serve or mis-classify these cohorts, and the consequences of mis-classification

are often disproportionately severe. The OAIC's research finding that 89% of Australians believe they should be told when ADM is used is, in our view, even more pronounced for these groups.

Question 5: What constitutes a significant service or support

Ranking (1-5):

- Access to essential banking and credit services: **5**
- Financial assistance: **5**
- Access to essential utilities (energy, water): **5**
- Access to telecommunications: **4**
- Physical assistance/support: **5**
- "In-home services": **5**
- Access to educational services: **4**

Additional examples we propose:

- Access to housing, including rental tenancy decisions and social housing.
- Access to insurance, including the offer, pricing, exclusion, or cancellation of any cover relevant to financial security (general, life, health, income protection, lenders mortgage insurance).
- Access to government services, including welfare, immigration, and identity verification.
- Access to employment, including hiring, promotion, performance management, and termination.
- Access to legal services and dispute resolution.
- Access to public services delivered through digital channels (e.g., MyGov, healthcare portals).

We encourage the OAIC to frame "significant service or support" by reference to the practical importance of the service to the individual's life, not the commercial classification of the service by the entity. A subscription-based grocery service serving a person with mobility limitations may be a significant service to that person even if it is not significant to the general population.

Question 6: Other legal frameworks the OAIC should consider

We would draw the OAIC's attention to:

- **APRA Prudential Standard CPS 234** (Information Security), particularly Section 15 (third-party information security capability) and the 30 April 2026 letter from APRA on AI risk management. The CPS 234 framework already requires entities to maintain a view of AI suppliers, model behaviour, and incident triggers. Alignment between CPS 234 evidence and ADM disclosure is meaningful for APRA-regulated entities and worth signposting in the guidance.
- **The EU AI Act, Article 86**, which establishes a right to an explanation of individual decision-making in high-risk AI systems. Australian APP entities operating internationally will need to reconcile APP 1.7 with Article 86 in practice, and consistent OAIC framing helps.
- **General Insurance Code of Practice (Part 8, claims handling)** and the parallel Banking Code of Practice obligations on hardship and decision review. Both create individual rights of review that depend on the entity knowing whether ADM contributed to the decision.

- **State-based privacy reform**, particularly the WA Privacy and Responsible Information Sharing Act 2024 commencing 1 July 2026, which gives stronger ADM rights (notice, information, human-intervention) than the Commonwealth obligation. The OAIC's guidance should anticipate a multi-jurisdiction reality for federally operating entities.
- **The Australian Human Rights Commission's Technology and Human Rights work**, particularly on algorithmic discrimination.
- **The Robodebt Royal Commission report and recommendations**, which remain the most consequential Australian guidance on the harms of unaccountable automated decision-making.
- **The Commonwealth Ombudsman's Automated Decision-Making Better Practice Guide (March 2025)**, which is operationally detailed and would help private-sector entities reading the OAIC's guidance understand what a mature ADM governance posture looks like.

Question 7: Daintree differential pricing edge case

Would inflated pricing in the Daintree edge case significantly affect Marco's interests?

Yes, in our view. Differential pricing based on inferred wealth or postcode is a paradigmatic case of ADM materially affecting an individual's access to goods, even where the dollar amount on any single transaction is modest. The aggregate effect is what matters: Marco is paying a premium for the same good as Charles on the basis of personal information he was required to disclose to use the service at all.

At what amount does differential pricing become significant?

We would discourage the OAIC from setting a fixed percentage threshold. A 5% differential on an essential good is more significant than a 50% differential on a discretionary luxury. A more durable framing is whether differential pricing has the effect, on average across the affected population, of restricting access, limiting choice, or transferring economic surplus from the consumer to the operator on the basis of personal information.

That said, if the OAIC needs an operationally usable threshold, a differential of more than 10% from the median price offered to comparable consumers, applied systematically, would in our view comfortably clear the significance bar.

Ranking of factors (1-5):

- The dollar amount of differential pricing: **3** (relevant, but secondary).
- The significance of the product or service being purchased: **5** (primary).

We would add a third factor we consider important: **whether the differential is applied on the basis of inferred characteristics from personal information**. Pricing that varies on the basis of objective, contemporaneous transaction inputs (e.g., shipping distance, order size) is different in kind from pricing that varies on the basis of inferred wealth, demographic, or behavioural profile. The latter is what the obligation is concerned with.

Question 8: Manscon / Opaque Window targeted job ad

Yes, it should be considered a decision that Rafqa did not receive the Manscon job ad.

Reasons.

- APP 1.9(a)-(b) explicitly include refusing or failing to make a decision, and refusing or failing to do a thing. Opaque Window's algorithm decided who to show the ad to and, by extension, decided who not to show it to. Rafqa's exclusion is a decision in the operative sense.
- The EU GDPR guidance treats exclusion from an opportunity through algorithmic targeting as a decision with similarly significant effect, and Recital 71 cites e-recruiting practices as an example.
- The case engages anti-discrimination law (the Sex Discrimination Act 1984) as well as the ADM obligation. The synergy is precisely the point the OAIC raises in the Background section: a transparent ADM disclosure is what allows an affected individual to exercise other rights, including anti-discrimination remedies.
- A holding that algorithmic exclusion from an opportunity is not a "decision" would create a perverse incentive: entities would prefer ADM-driven exclusion over explicit refusal, because the former would carry weaker transparency obligations than the latter. That cannot be the policy intent.
- Manscon, as the entity that engaged Opaque Window and benefited from its targeting, has "arranged for" the ADM under APP 1.7(a) and should bear the disclosure obligation, even though Opaque Window operates the underlying system.

Question 9: "Arranged for"

We strongly support a broad reading of "arranged for" that places accountability with the entity that procures or directs the use of the ADM, irrespective of who operates the underlying system. The Issues Paper's framing (procurement, contracting, permitting employees to use AI tools, case-management auto-escalation) is right and should be retained.

Scenarios where further guidance would be valuable.

- **Generic enterprise AI tools used at scale within an organisation.** Many entities have rolled out general-purpose AI assistants (chat interfaces over LLMs) for staff use, with no specific decision-making task in mind. When staff use those tools to draft customer correspondence, performance reviews, refund decisions, or eligibility assessments, those uses can become ADM in scope. The guidance should address how entities should approach disclosure where the ADM use is emergent rather than designed.
- **Third-party SaaS providers that embed AI mid-stream.** Many enterprise SaaS products have introduced AI features (CRM lead scoring, ticket triage, claims fraud scoring) without their customers actively procuring those features. APP entities may be subject to ADM scope by virtue of supplier feature releases. The guidance should clarify how the procuring entity discharges its obligation where it does not control whether AI is used in the supplier's pipeline.
- **AI gateways and orchestration layers.** Some entities procure AI orchestration platforms that route requests to different models depending on the task. The guidance should clarify how disclosure should treat the upstream model identity when the entity does not directly select the model.
- **Foundation-model providers used through APIs.** Entities calling commercial foundation-model APIs (OpenAI, Anthropic, Google, AWS Bedrock) for decision-relevant tasks have "arranged for" the ADM, but they do not control the model. Guidance on what entities are expected to disclose about the model component is needed; we suggest model family and provider should be sufficient, since model version drift is constant.

- **Staff-initiated use of consumer AI products.** A staff member using a consumer AI chatbot on a personal device to draft a customer-affecting decision should be treated as the entity having "arranged for" the ADM where the entity has not actively prohibited the practice. This recognises the realities of shadow IT and creates an incentive for entities to maintain a clear position on staff use of AI.

In each case, the unifying principle is that the entity benefiting from the ADM is best placed to maintain visibility of it and best placed to disclose accurately. Allowing entities to disclaim accountability for ADM operated by third parties would significantly undermine the obligation.

Question 10: Extent of disclosure

We support the OAIC's framing in all five bullet points (plain language, structured for further inquiry, appropriately tailored, logically organised, framed for challenge).

Additional considerations we propose for the guidance.

Currency of disclosure. APP 1.3's requirement that the privacy policy be "up to date" is particularly demanding for ADM, because AI use within an organisation typically evolves on a faster cycle than annual policy refresh. We encourage the guidance to specify that an entity should disclose, at a minimum, the categories of personal information used and the kinds of decisions made by ADM **as the entity could reasonably know them**. This implies entities should have an operational basis for knowing what their AI is actually doing, not only a record of what they intended to deploy. Where an entity is not in a position to know, the guidance should expect the entity to take steps to improve its visibility, not to fall back on generic language.

Granularity. We support the OAIC's caution against overwhelming detail. We would propose the guidance frame granularity by reference to whether the individual could reasonably understand:

- What categories of personal information about them are being used (e.g., demographic, financial, health, behavioural, contact, identifiers).
- What categories of decisions are being made or supported (e.g., pricing, eligibility, ranking, content selection, fraud assessment).
- Where the individual can go for further information or to contest a decision.

Disclosure beyond this baseline should be available to consumers who request it under the obligation's "structured to enable further inquiry" requirement, rather than crowded into the privacy policy itself.

Contestability. The OAIC's fifth bullet (framed to allow challenge) is, in our view, the most important. A disclosure that does not give the individual a clear path to inquiry, review, or human reconsideration is not meaningful transparency. We encourage the guidance to require every ADM disclosure to identify the contact point for inquiries and, where another framework provides review rights (administrative law, anti-discrimination, codes of practice), to cross-reference it.

Layered disclosure. Drawing from the UK ICO and Alan Turing Institute guidance cited in the Issues Paper, we encourage the OAIC to endorse a layered approach: a short, plain-language ADM summary in the privacy policy itself, with a deeper "how we use AI" page accessible for users who want more detail. This reflects how people actually engage with privacy notices.

Evidence of accuracy. We encourage the OAIC to note in the guidance that, while the obligation is to disclose in the privacy policy, the OAIC may have regard to the entity's operational basis for

the disclosure when assessing complaints or compliance. This nudges entities away from generic disclosure templates and toward grounding disclosure in operational visibility.

Commercial-in-confidence carve-out. We agree that commercial-in-confidence material about ADM systems is excluded under the Explanatory Memorandum, but we encourage the OAIC to clarify the carve-out narrowly: it should cover proprietary model architecture, training data, and similar IP, but should not cover the categories of personal information used or the categories of decisions made. Allowing entities to claim commercial-in-confidence over the latter would undermine the obligation entirely.

Disclosure of who decides. Where ADM is "arranged for" via a third party, we encourage the guidance to require the privacy policy to name the categories of third parties involved (e.g., "we use AI-based candidate screening tools provided by a third-party recruitment platform"), without necessarily naming the specific supplier.

Closing

Thank you for the opportunity to make this submission. We would welcome the chance to provide further technical or operational input as the guidance develops, particularly on how operational evidence systems can support entities to make and maintain accurate ADM disclosures.

For follow-up correspondence:

- 40 South Pty Ltd
- Email: hello@40south.au