Automated Decision-Making Transparency Obligation (APP 1)

Issues Paper Submission — Response by RecordPoint

Submitter	RecordPoint
Contact	[redacted]
Date	June 2026

Question 1

Substantially and directly related to making a decision

In your view, what are the relevant factors to enable an entity to assess whether a computer program substantially facilitates and is directly connected to a human decision-maker's decision? For example:

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

Please rank the factors using the above mentioned (1-5) ranking system and include any additional factors.	Ranking the listed factors (5 = strongly agree this is a relevant factor): Degree of reliance on the ADM system output — **5** Integration of ADM into the decision-making workflow — **5** Ability and likelihood of human override — **4** Transparency and explainability of outputs — **4** Nature of the output (advisory vs determinative) — **3** Additional factors we consider relevant: Whether the program shapes or limits the information/options the human sees (ranking, shortlisting, triage, summarisation) — **5** Evidence of genuine human engagement vs automation bias (does the human routinely accept the output?) — **5** Whether the output is presented as a default or pre-filled recommendation — **4**
Free text comments	In our view the assessment should be **functional, not formal**. The decisive question is whether the program's output materially influences the outcome in practice — if it does, it is substantially and directly related, however it is labelled. We caution against over-weighting the advisory/determinative distinction. An 'advisory' recommendation that humans routinely adopt, or a summarisation step that controls which facts the decision-maker ever sees, is in substance determinative. Reliance and workflow integration are the strongest indicators because they describe what actually happens; the ability to override should be judged by its likelihood in practice, not its theoretical availability.

	A program that determines what reaches a human — by filtering, ranking or prioritising — is substantially and directly related even where a human makes the nominal final call. Entities can only assess these factors reliably if they maintain a current inventory of where ADM sits in their decision workflows and keep records of how outputs are used. Transparency therefore depends on an entity first knowing its own systems.

Question 2

Substantially and directly related to making a decision

Fictional Edge Case: Generative AI with human oversight

The National Aged-Care Insurance Agency (NACIA) makes decisions about eligibility for elderly people trying to access life-changing aged-care support. To make eligibility assessments quicker, NACIA staff use a generative AI chatbot, GPTea, to summarise candidate profiles and receive recommendations about eligibility. NACIA's human staff always have oversight and make the final decision about eligibility based on the recommendation provided by the chatbot.

Would you consider NACIA's described uses of GPTea to be substantially and directly related to making decisions? (Please provide yes/no)	**Yes.**
Can you provide your reasons for your response?	GPTea performs two functions that are each substantially and directly related to the eligibility decision: it **summarises the candidate's profile** (controlling and framing the information the human relies on) and it **recommends an eligibility outcome**. Both are key factors facilitating the decision and have a direct connection to it. The presence of human oversight does not remove the obligation. The test is whether the program substantially facilitates the decision — and a recommendation plus a curated summary plainly do, particularly in a high-volume process where automation bias makes routine adoption of the recommendation likely. Generative AI summaries can also omit, distort or fabricate material facts, so the summarisation step alone can decisively shape the outcome. These eligibility decisions also significantly affect rights or interests (access to life-changing support, often for vulnerable older people). Treating 'a human makes the final decision' as a carve-out would let entities keep decision-shaping AI off the books — the opposite of the transparency the obligation is intended to deliver. This use should be disclosed.

Question 3

Meaning of significantly affect rights or interests

What factors do you consider increase the likelihood that a decision could affect an individual's rights or interests? For example:

- Sensitive information
- Vulnerable persons
- Intrusive practices
- Financial outcomes
- Other

Please rank the examples using the above mentioned (1-5) ranking system and include any additional examples.	Ranking the listed factors (5 = strongly agree this increases the likelihood): Sensitive information — **5** Vulnerable persons — **5** Financial outcomes — **5** Intrusive practices — **4** Additional factors: Use of inferred or derived data (decisions based on profiling rather than information the person supplied) — **5** Opacity to the individual / inability to detect that a decision was made (silent exclusion) — **5** Irreversibility or persistence of the effect — **4** Cumulative or aggregate effects (small effects repeated or applied at scale) — **4** Effects driven by other people's data (decisions based on the conduct of similar cohorts) — **4**
Free text comments	These factors are best understood as **amplifiers of consequence and of the individual's inability to protect themselves**. Sensitive information, vulnerability and material financial effects heighten consequence; intrusive and profiling-based practices heighten both consequence and opacity. The most under-recognised factor is opacity. Where a person cannot tell that an automated decision was made about them — for example a silent exclusion — they cannot seek review, which is precisely the gap transparency is meant to close. We encourage the OAIC to treat 'rights or interests' functionally and to recognise that significance is contextual: an effect that is trivial for most may be severe for some. Reliable assessment again depends on the entity knowing what personal and inferred information feeds each decision.

Question 4

Meaning of significantly affect rights or interests

What classes of person do you consider to be vulnerable for the purposes of considering a decision's expected effect on a person's rights or interests, and why?

Can you provide your reasons for your response?	We recommend a **contextual, non-exhaustive** approach rather than a closed list. Groups that will frequently warrant heightened consideration include: children and young people; older people; people with disability (including cognitive or intellectual disability);

	people experiencing financial hardship, insecurity or over-indebtedness; people with limited digital, financial or English-language literacy; people from culturally and linguistically diverse backgrounds; Aboriginal and Torres Strait Islander peoples; people experiencing family or domestic violence, homelessness or other crisis; and people in a position of acute dependence or power asymmetry relative to the decision-maker. But vulnerability is **situational as well as categorical**. A person not ordinarily vulnerable can be made vulnerable by the specific decision and context — a medical emergency, a moment of financial distress, or a decision that turns on a sensitive, inferred attribute. Guidance should therefore direct entities to assess vulnerability against the decision, asking who is most affected and least able to understand or contest the outcome, rather than relying solely on membership of a predefined class. This matters because the same ADM system may have a trivial effect on the general population and a severe effect on a vulnerable subset. The obligation should be read with that subset in mind.

Question 5

Meaning of significantly affect rights or interests

What do you consider to be a significant service or support, for the purposes of considering whether a decision affects access to a significant service or support? For example:

- physical assistance/support
- financial assistance
- access to educational services
- 'in home services'
- access to essential banking and credit services
- access to telecommunications
- access to essential utilities
- other

Please rank the services/support using the above mentioned (1-5) ranking system and include any additional examples.	Ranking the listed examples (5 = strongly agree this is a significant service or support): Physical assistance/support (incl. healthcare, aged and disability care) — **5** Financial assistance — **5** Access to essential banking and credit services — **5** Access to telecommunications — **5** Access to essential utilities — **5** Access to educational services — **4** 'In home services' — **4** Additional examples: Housing and tenancy — **5** Employment and access to employment opportunities — **5** Government benefits / social security — **5** Insurance — **4**

	Digital identity and access to essential digital platforms — 4
Free text comments	'Significant service or support' is best defined **functionally — by the consequence of being denied, degraded or delayed access —** not by the category label. The unifying test is whether the service or support is one an individual reasonably needs to participate in economic and social life, maintain health and wellbeing, or meet essential needs. On that test, banking and credit, utilities, telecommunications, healthcare and physical/aged/disability care, financial assistance, housing and employment are clearly significant. Telecommunications in particular is now a gateway to almost all other essential services — emergency, government, banking and healthcare — and should be treated as essential infrastructure. Education and in-home services are significant where access materially affects life chances or independent living. We recommend the guidance to frame significance around impact and essentiality, supported by examples, rather than an exhaustive list, so that it remains durable as services digitise.

Question 6

Meaning of significantly affect rights or interests

Are there any other legal frameworks or policies that the OAIC should have regard to in considering the meaning of 'rights' and 'interest'?

Free text comments	We support the OAIC's approach of reading 'rights' and 'interests' consistently with the frameworks already identified — anti-discrimination law, administrative law and procedural-fairness jurisprudence, the FOI Act and Information Publication Scheme, the General Insurance Code of Practice, the Administrative Review Tribunal Act and Corporations Act, and comparable overseas instruments. We suggest the OAIC also have regard to: • the Australian Consumer Law, including unfair-practices and consumer-protection concepts relevant to differential pricing and targeting; • the National Consumer Credit Protection framework and the Banking Code of Practice; • the Consumer Data Right, and sector-specific codes (e.g. telecommunications and energy); • Australia's international human rights commitments — including the ICCPR, the Convention on the Rights of the Child and the Convention on the Rights of Persons with Disabilities — as they inform the meaning of 'rights'; • the Online Safety Act, where children are affected; and • Australia's emerging AI assurance settings, including the Voluntary AI Safety Standard and the proposed mandatory guardrails for AI in high-risk settings, alongside the Digital

	Transformation Agency's AI policy and transparency standard. Reading the obligation in harmony with these frameworks will give entities a coherent compliance picture and reduce duplicative or inconsistent disclosure. RecordPoint tracks these regimes through its public Global AI Regulations Tracker and would welcome the OAIC mapping the ADM obligation to the broader landscape.

Question 7

Meaning of significantly affect rights or interests

Fictional Edge Case: Differential pricing

E-commerce company, Daintree, allows subscribers to order books to their home. To make an account, Marco was required to disclose his address, being the luxurious Capital Crescent in Sydney CBD. Marco is surprised that his desired book, The Art of Inheritance, was priced at $100 when his book-club buddy, Charles, from a country town, paid only $35.

Marco is shocked to read an article reporting that Daintree uses differential pricing to inflate prices for individuals from wealthy postcodes. Marco is a freelance artist, so he does not have enough disposable income to purchase the book.

Restricting access to financial products or healthcare would clearly be a decision which significantly affects the interests of an individual. **In the above Edge Case, would you consider the inflation of prices for a consumer product to have significantly affected the interests of Marco?**	**No — not on these facts.** Pricing a discretionary consumer product such as a book differently does not, on its own, significantly affect Marco's rights or interests. The statutory test (APP 1.7(b)) asks whether the decision could reasonably be expected to *significantly* affect rights or interests. The Explanatory Memorandum confirms the effect must be 'more than trivial' and have the potential to 'significantly influence the circumstances of the individual'. Critically, the Memorandum's own differential-pricing example is tied to the provision of, or access to, **significant goods or services**, and the statutory examples in APP 1.9(d) concern statutory benefits, contractual rights, and access to a **significant** service or support. A single book is a discretionary purchase, not a significant good or service; paying more for it — or declining to buy it — does not significantly influence Marco's circumstances. This aligns with the EU position the OAIC already draws on. Under the GDPR profiling guidelines, differential pricing reaches a 'significant effect' only where, for example, prohibitively high prices effectively bar a person from certain goods or services; ordinary personalised pricing in typical cases does not. Under the EU AI Act, consumer retail pricing is not among the 'high-risk' essential services use cases (limited to public benefits, creditworthiness/credit scoring, life and health insurance pricing, and emergency dispatch), so it would not attract that Act's transparency or explanation duties. **In principle, however, differential pricing can be in scope** — where it concerns significant or essential goods or services (such as credit, insurance, healthcare, utilities or housing), or where prohibitively high personalised prices effectively bar a person from accessing them. The decisive questions are the significance

	of the good or service and the seriousness of the effect, not the mere fact that personal information was used. The use of a residential address as a wealth proxy remains a legitimate consumer-protection and privacy concern, but it is not, by itself, what brings a decision within this transparency obligation.
If yes, at what amount would differential pricing be considered to have significantly affected the interests of an individual? **5% price difference, 20% price difference, or other (please specify amount)**	**Other — significance does not turn on a price-difference percentage at all.** No fixed figure (5%, 20% or otherwise) should be determinative. The threshold question is whether the good or service is significant, and whether the differential is serious enough to bar or materially impair access to it. On that test, a discretionary item such as a book falls short at any percentage, while differential pricing of an essential good or service can be significant even at a modest percentage — particularly for a person in vulnerable circumstances. Significance is therefore best assessed on the combination of (i) the significance or essentiality of the good or service, (ii) the seriousness of the effect (does it effectively bar or materially impair access?), (iii) whether the pricing uses personal, sensitive or inferred data or proxies for protected attributes, and (iv) the individual's circumstances. If the OAIC nonetheless adopts a presumptive figure for workability, it should apply only to significant or essential goods and services, and to personalised pricing that uses personal data — not as a universal trigger. A blanket low percentage applied to all consumer pricing would over-capture ordinary commercial practice.

In judging whether differential pricing significantly affects rights or interests, how important would you consider the following factors on a scale from 1 to 5 (with 1 being not important and 5 being very important):

The dollar amount of differential pricing	3 — relevant context, but neither necessary nor sufficient. A large percentage on a trivial product is still not significant, and a modest one on an essential good or service can be.
The significance of the product or service being purchased	5 — the decisive factor. Whether interests are significantly affected turns principally on how significant or essential the good or service is, consistent with the Explanatory Memorandum's reference to 'significant goods or services'.

Question 8

Meaning of making a decision

Fictional Edge Case: Discriminatory targeted job ad

Engineering firm, Manscon, is hiring for graduate engineers. To try to target the best candidates, Manscon advertises open roles on job site, Opaque Window. Opaque Window algorithmically promotes this job posting to candidates it thinks are best suited. One metric it uses is gender of account holders. Since most engineers at Manscon are male, Opaque Window prioritises promoting the job to male engineering graduates.

Rafqa is a female engineering graduate searching for a job on Opaque Window. She does not receive the Manscon job ad, which limits her employment options.

Should it be considered a 'decision' that Rafqa did not receive the Manscon job ad? Please provide a Yes/No response	**Yes.**
Can you provide your reasons for your response?	APP 1.9 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. Opaque Window's algorithm determined — using gender as a metric — not to surface the Manscon ad to Rafqa. That determination is a decision (a failure to do a thing) that affects her rights and interests by limiting her employment opportunities, and it engages anti-discrimination law. The fact that the effect is an omission (an ad not shown) rather than an overt rejection does not take it outside the obligation. On the contrary, silent exclusion is the more concerning case, because Rafqa cannot detect or contest a decision she never saw. We would also note the 'arranged for' dimension: Manscon arranged for Opaque Window's targeting to fill its roles, so Manscon should not be able to disclaim responsibility merely because a third party operates the algorithm. Treating algorithmic non-delivery as outside 'decision' would create an easy avoidance route for some of the highest-impact automated effects.

Question 9

Meaning of 'arranged for'

Are there any scenarios relevant to the meaning of 'arranged for' that you consider require guidance? If so, why?

Can you provide your reasons for your response?	Yes. Several scenarios warrant guidance, because most entities will 'arrange for' ADM that they neither build nor operate themselves. Priority scenarios: • Embedded AI features in general-purpose software (CRM, HR, marketing, productivity tools) that quietly introduce scoring, ranking or recommendation into a decision workflow — entities may not realise they have arranged for ADM. • 'Shadow AI' — staff using generative-AI tools that are not formally sanctioned but nonetheless shape decisions (e.g. drafting assessments, summarising applicants). • Layered supply chains and sub-processors, where the operator of the model is several steps removed from the entity making the decision. • Configuration versus default use — where an entity tunes, prompts or sets parameters for a third-party model (greater responsibility) versus using it out of the box.

	• Agentic systems that can take action or transact on the entity's behalf. The guiding principle should be that 'arranged for' follows functional responsibility: the entity that chooses to deploy ADM into its own decision-making, and benefits from it, should bear the disclosure duty even if another party operates the model. This mirrors the OAIC's own examples — procuring a screening tool, or directing staff to use an AI tool. To meet this in practice, entities need to maintain a current inventory of the AI/ADM systems they use — including embedded, third-party and shadow systems — and to conduct procurement due diligence (data rights, model-change notifications, logging access). This is the discipline RecordPoint helps regulated organisations operationalise: discovering AI in use, maintaining an authoritative AI inventory, and capturing the approvals, risk reviews and audit trail that make a disclosure substantiable. Guidance that makes 'arranged for' turn on who deploys and benefits from the ADM, and that expects a maintained inventory, would give entities a clear and workable test.

Question 10

Extent of disclosure

When making disclosures about ADM, the OAIC's view is that entities must ensure that meaningful and accessible information is provided to consumers. In doing so, the OAIC considers that entities should strike an appropriate balance to ensure that disclosures are:

- clearly articulated in plain language and easy to understand;
- structured to enable consumers to request further information, where required;
- appropriately tailored – sufficiently specific to be meaningful, while avoiding overwhelming levels of detail;
- organised so that similar information is grouped in a logical manner; and
- framed in a way that allows the information and the decision to be challenged or contested.

Having regard to the above considerations, is there anything additional that entities should disclose to provide effective and meaningful disclosures about ADM?	We strongly support the OAIC's plain-language, layered and contestable approach. We suggest entities also be encouraged to: • Maintain an internal AI/ADM inventory as the precondition for any meaningful disclosure — an entity cannot accurately disclose ADM it has not identified. Disclosure quality is downstream of data and AI governance maturity. • Disclose at the right moment — providing layered, 'just-in-time' notice at or near the point of interaction, with the ability to drill down, rather than relying solely on a static privacy policy. • Convey not just the kinds of personal information used, but a meaningful, plain-language account of the role the ADM plays

	in the decision, its key limitations, and how to seek human review or contest the outcome. • Be specific about third-party and embedded ADM (the 'arranged for' point), so individuals understand when a decision about them is shaped by a system the entity did not build. • Keep disclosures accurate and current — versioned and updated whenever the underlying system materially changes — so the privacy policy never drifts from what the systems actually do and avoid over-stating AI capabilities. Underlying this is one caution: disclosure is necessary but not sufficient. Transparency is only credible if it sits on top of genuine accountability — records, data lineage and audit trails that let an entity substantiate what it discloses and reconstruct, explain and (where needed) review an individual decision. A disclosure that cannot be evidenced is assertion, not accountability. We therefore encourage the OAIC to frame disclosure as one link in an accountability chain, and to signal that good disclosure presupposes that entities know their data, know their systems, and can show their work.