22 June 2026

By Email privacyreformtranche1@oaic.gov.au.

Mastercard Policy Submission: Automated Decision-Making Transparency Obligation (APP 1) Issues Paper Submission

Thank you for the opportunity to provide feedback and comments on the OAIC's development of the ADM obligation guidance.

As a global technology company operating data-driven products and services across multiple regulated environments, Mastercard supports guidance that gives effect to the transparency objective of the Privacy Act in a way that is clear, coherent and practically implementable.

This position is also consistent with Mastercard's Data & Tech Responsibility Principles, which emphasise security and privacy, transparency, accountability, fairness, inclusion, innovation and social impact as foundations for responsible data-driven innovation. In particular, those principles support clear and simple explanations of how data is collected, used and shared, and how technologies such as artificial intelligence are used to deliver intended outcomes, while maintaining robust governance, oversight and safeguards to minimise bias, inaccuracy and unintended harm.

In our view, the new obligation should be interpreted and applied consistently with the broader privacy framework, while recognising that automated decision-making may also be addressed by overlapping legal and regulatory regimes, including consumer protection, anti-discrimination, administrative and sector-specific rules. The guidance should therefore promote meaningful consumer understanding and accountability without diluting existing privacy protections or imposing requirements that are disproportionate to the APP 1 transparency purpose.

Please see our detailed responses to the specific consultation proposals in **Annexure 1**.

Annexure 1.

MASTERCARD RESPONSES TO CONSULTATION QUESTIONS

OAIC question	Mastercard response
Question 1 – Substantially and Directly Related to Decision-Making	
1. 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? Please rank factors and include any additional factors.	In Mastercard's view, a computer program should be regarded as "substantially facilitating" and "directly connected" to a human decision where, assessed holistically, it plays a **material role in determining or shaping the outcome of that decision**, rather than merely providing contextual or ancillary input. Relevant factors include whether the outcome would likely differ without the system; whether the system output constrains or frames the available options (for example, by setting eligibility thresholds, scores or rankings); and whether the output is **relied upon in practice as determinative**, even if not formally binding. This interpretation gives proper effect to the transparency purpose of APP 1.7 without narrowing the circumstances in which privacy law applies. The assessment should also consider the **nature of human involvement and the functional proximity of the system to the decision**. Where the human decision-maker exercises genuine, independent judgment — including the ability and practice of interrogating, departing from, or overriding system outputs — the system is less likely to meet the threshold. Conversely, where human involvement is limited to procedural endorsement or where the system's output is integral and routinely followed in operational workflows, it should be regarded as substantially facilitating and directly connected to the decision. Guidance framed in this way would support consistency with comparable international approaches, including UK-aligned thinking, while remaining responsive to the Australian statutory context.
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,	In Mastercard's view, use of the generative AI tool should not, without more, be treated as substantially and directly related to making a decision merely because it provides a recommendation in a decision-support process. On the stated facts, the better view is that GPTea would not be substantially and directly related to the human decision if staff genuinely review the recommendation, exercise informed and independent judgment, and can realistically reach a different outcome. In that case, the tool supports the human review process rather than materially determining the result. The key question

OAIC question	Mastercard response
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. a. Would you consider NACIA's described uses of GPTea to be substantially and directly related to making decisions?	is not whether the tool appears in the workflow, but whether it materially shapes or constrains the final decision in practice. The assessment is, however, fact-specific. If the recommendation frames the issues in a way that is routinely followed, or if human oversight is merely formal and does not involve meaningful scrutiny or a real possibility of departure, the tool may still be substantially and directly related to the decision. The guidance should therefore focus on the substance of human involvement, including the person's real ability to depart from the suggested outcome, rather than the mere existence of final human sign-off. This issue also arises in multi-entity arrangements, where one entity generates an output and another makes the consumer-facing decision. In those cases, the quality of human involvement is relevant not only to whether the decision is properly characterised as human or automated, but also to which entity should carry the transparency obligation. Where the consumer-facing entity can genuinely depart from the upstream recommendation, it is generally the appropriate entity to provide that disclosure.
Factors Increasing Likelihood of Significant Effect	
2. What factors do you consider increase the likelihood that a decision could affect an individual's rights or interests? Please rank factors and list any additional factors.	Significant services should include healthcare, housing and utilities, essential banking and credit, and education and telecommunications, in other words, services that would materially affect a person if they were not able to access them, or if they were to lose access. The guidance should preserve flexibility for context-and fact -specific assessment by the entity making the decision, while making clear that the concept is intended to capture decisions with real practical consequences for individuals. Vulnerability is relevant, but entities should not be required to collect new vulnerability data solely to decide if APP 1.7 applies. A proportionate approach should focus on reasonably identifiable indicators available in the ordinary course of business.
Vulnerable Persons	
3. 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?	Mastercard supports clear and effective transparency for individuals who may be in vulnerable circumstances. At the same time, the guidance should avoid interpretations that would require entities to collect, infer, or operationalise additional sensitive personal information solely in order to determine whether a transparency obligation is triggered. Relevant cohorts may include children, older persons, persons with disability, individuals experiencing financial distress, and persons affected by power asymmetries or digital exclusion.

OAIC question	Mastercard response
	However, vulnerability should not be treated as a closed or fixed class, as it may arise contextually depending on the nature of the decision and the individual's circumstances. Entities should not be required to infer, collect, or process additional personal information solely for the purpose of determining whether a person is vulnerable and whether the automated decision-making provisions are enlivened. Such an approach would itself create additional privacy impacts and would be difficult to reconcile with a proportionate, privacy-protective implementation of the obligation. It is important to refer back to the purpose of the automated decision-making provisions in Australian Privacy Principle 1, which are directed to transparency rather than to the comprehensive regulation of automated decision-making itself. That framing should not be understood as downgrading privacy law. Rather, it reflects the need to interpret APP 1.7 coherently within the Privacy Act as a transparency obligation that complements, and operates alongside, other privacy protections and broader regulatory safeguards. Vulnerability remains relevant in this context, particularly in informing how information should be communicated so that it is comprehensible and usable by the individual. However, the obligation should not be interpreted as requiring entities to implement differentiated or bespoke disclosures based on inferred vulnerability status, particularly where doing so would require additional data collection or complex categorisation. More substantive risks to vulnerable persons arising from automated decision-making are also addressed through other areas of law and regulation, including anti-discrimination law, consumer protection frameworks, and sector-specific obligations. The guidance should therefore reinforce that the APP 1 transparency obligation operates alongside, and does not displace, those complementary protections. This would support a coherent regulatory approach while ensuring that privacy law continues to play its full role.
Significant Services or Supports	
4. What do you consider to be a significant service or support for the purposes of assessing whether a decision affects access? Please rank	A significant service or support should be understood as one that is important to an individual's ability to participate in ordinary economic and social life, maintain safety and wellbeing, or access essential opportunities. This reflects the underlying concern that certain decisions may have a material and/or lasting effect on an individual's circumstances, rather than producing a minor or transitory outcome.

OAIC question	Mastercard response
categories and include additional factors.	In practice, this will often include services such as healthcare, significant financial assistance, housing-related support, access to essential banking and credit services, telecommunications necessary for modern participation, utilities, meaningful access to education, and in-home or personal support services. These categories are indicative rather than exhaustive, and whether a service is significant will depend on both the context in which it is provided and used, and also the person to whom it is (or is not) provided. The assessment of whether a service is "significant" should not turn solely on the category of service, but rather on the nature and degree of the effect on the individual. Relevant factors may include the extent to which the service is necessary or difficult to substitute, whether denial or restriction of access may result in exclusion from economic or social participation, the impact on the individual's financial position or living circumstances, and the availability of realistic alternatives. This approach supports a context-based assessment that is capable of accommodating a wide range of decision settings, including those in which the same type of service may have materially different implications depending on the individual's circumstances. It also promotes practical implementation by allowing entities to apply the guidance in a way that is meaningful, proportionate and operationally workable.
Other Relevant Legal Frameworks	
5. Are there any other legal frameworks or policies that the OAIC should have regard to in considering the meaning of 'rights' and 'interest'?	Yes. The OAIC should expressly recognise that the meaning of rights and interests sits alongside, and should be interpreted with awareness of, other legal and policy frameworks that address harm in automated or data-driven decision-making. These include anti-discrimination law, administrative law for government decision-making, consumer protection law, sector-specific obligations in fields such as financial services, credit and insurance, and broader privacy reform proposals including the fair and reasonable test and proposed rights relating to automated decisions. Recognising those frameworks does not narrow the role of privacy law; rather, it supports a coherent interpretation of APP 1.7 within the wider legal landscape. This is important because decisions involving automated processing may engage multiple legal frameworks simultaneously, each addressing different aspects of potential impact. Privacy law plays a central role in governing the collection, use and handling of personal information and in supporting transparency and accountability, while other frameworks may more directly address issues such as discrimination, procedural fairness, market conduct or consumer detriment in particular sectors or contexts.

OAIC question	Mastercard response
	Recognising these interactions would help ensure that the OAIC's guidance remains legally coherent and capable of operating effectively alongside existing regulatory regimes. It would also avoid placing disproportionate interpretive weight on concepts such as significance or vulnerability within APP 1 alone. In our view, the guidance should therefore describe the APP 1 transparency obligation as operating in conjunction with these other frameworks, contributing to overall accountability and consumer understanding without being expected, on its own, to address all potential risks associated with automated decision-making.
Meaning of 'Decision' (Targeted Advertising)	
6. Should it be considered a 'decision' that an individual did not receive a targeted job advertisement?	Whether the non-receipt of a targeted job advertisement constitutes a "decision" should not attract a blanket answer. The assessment should turn on two considerations. **First, the extent to which the computer program influenced the outcome.** Where the algorithm is the decisive factor in determining whether an individual is exposed to an opportunity — such that, absent its operation, the individual would have been able to view and act on the advertisement — this weighs in favour of treating the non-disclosure as a decision. **Second, the reasoning behind the exclusion.** If an advertisement is withheld because the individual does not meet an objective and unbiased requirement for the role — for example, the candidate does not hold a required university degree — the non-disclosure does not deprive the individual of an opportunity they could realistically pursue, and should not be characterised as a decision with detrimental effects. By contrast, where the exclusion is driven by subjective assessments made by the computer program — particularly assessments based on protected attributes such as gender, and made without meaningful human intervention — the withholding should be regarded as a decision that may affect the individual's rights or interests. **Applying these considerations to the OAIC's fictional edge case**, Opaque Window's algorithm uses gender as a targeting metric for the Manscon graduate engineering role. Gender is not an objective job requirement; it is a protected attribute under Australian anti-discrimination law. Because the algorithm is the decisive factor in withholding the advertisement, and the basis for exclusion is neither objective nor unbiased, Rafqa's non-receipt of the advertisement should be treated as a decision affecting her access to an employment opportunity within the meaning of APP 1.9(d)(iii).

OAIC question	Mastercard response
	Mastercard submits the guidance should direct entities to assess: (1) whether the computer program was a **decisive factor** in the outcome; and (2) whether the basis for exclusion was **objective and unbiased** or **subjective and evaluative** — with particular regard to the use of protected attributes or opaque criteria without meaningful human oversight. This contextual approach avoids over-breadth while ensuring algorithmic processes that substantively limit access to material opportunities on discriminatory grounds are properly recognised as decisions.
Meaning of 'Arranged For'	
7. Are there any scenarios relevant to the meaning of 'arranged for' that you consider require guidance? If so, why?	Yes. Further guidance would be valuable in three areas. **Multi-party ecosystems.** Many contemporary services rely on outsourced, embedded or platform-based operating models in which one entity develops or operates a tool, another procures or configures it, and a third makes the customer-facing decision. In Open Banking, for example, one participant may generate an income verification report, transaction analysis or account ownership verification, while another uses that output in a customer-facing onboarding, affordability or credit decision. The guidance should clarify that an entity has "arranged for" a computer program where it decides the program will be used in its decision-making process — including by selecting, deploying, configuring or authorising it — consistent with the Explanatory Memorandum's recognition that the entity responsible for *arranging* the program bears the obligation, not the entity that merely *operates* it. **Upstream providers and data processors.** An upstream provider that merely generates a general-purpose analytical output — without controlling how that output is applied in a downstream decision — should not be treated as having "arranged for" the computer program in relation to that decision. The guidance should go further and **expressly confirm that a data processor acting on behalf of a decision-making entity falls outside the scope of "arranged for"**. This is important for two reasons: • **Practical purpose.** A data processor does not typically present its own privacy policy to consumers. Requiring an APP 1 disclosure from an entity whose privacy policy the affected individual will never see does not advance the transparency objective underpinning the obligation. • **Legislative coherence.** Future tranches of the Privacy Act reforms are expected to introduce a formal controller/processor distinction. The OAIC's guidance should anticipate this framework and avoid creating obligations for processors that will need to be unwound once the distinction is legislated. An express carve-out now would provide certainty and ensure that entities whose outputs or contributions to decisions — such as analytics, credit

OAIC question	Mastercard response
	reporting or data enrichment services — are not inadvertently captured as having "arranged for" a program they do not control in the context of the relevant/substantial decision. **Avoiding duplication without improving transparency.** Where an arrangement does not directly affect the final decision-maker's use of the program, requiring overlapping APP 1 disclosures across multiple entities in the same chain risks creating duplication and consumer confusion without improving transparency. The guidance should direct the disclosure obligation to the entity best placed to provide meaningful information to the affected individual — ordinarily, the entity that makes or controls the customer-facing decision. Without clearer guidance on these points, there is a risk of duplicated disclosures, uncertainty about responsibility, and privacy policies that do not accurately reflect the consumer's relationship with the entity making the relevant decision.
ADM Disclosure Requirements	
8. Having regard to the OAIC's proposed principles for ADM disclosures (clarity, accessibility, proportionality, structure, contestability), is there anything additional that entities should disclose to provide effective and meaningful disclosures?	In addition to the matters identified by the OAIC, the guidance should encourage disclosures that are role-appropriate, outcome-focused and capable of being understood by a consumer without detailed technical knowledge. In our view, an effective privacy policy disclosure should explain, at a level suitable to the APP 1 context, what kinds of decisions are made or supported, what kinds of personal information are used, the broad role that automated tools play in the process, and how a person can seek more information or raise concerns. The guidance should also confirm that APP 1 does not require disclosure of commercially sensitive model details, security-sensitive detection logic, or highly granular technical specifications, provided the description remains meaningful enough to support genuine transparency and consumer understanding. This would promote practical implementation while preserving the substance of the privacy obligation.