12 June 2026

Australian Government
Office of the Australian Information Commissioner
Submitted to: privacyreformtranche1@oaic.gov.au

Re: Feedback Regarding the Automated Decision-Making Transparency Obligation (APP 1) Issues Paper (dated 18 May 2026)

To the Office of the Australian Information Commissioner:

We appreciate the opportunity to provide feedback on the Automated Decision-Making Transparency Obligation (APP 1) Issues Paper.

Epic is a global healthcare technology company working with healthcare organizations in Australia and 15 other countries. We provide an integrated software suite to manage clinical care and back-office operations. Our patient portal, MyChart, puts medical records in the pockets of patients, enhancing medical literacy, communication, and wellness.

Extensive experience developing and deploying AI across clinical and administrative settings has given Epic a deep appreciation for the practical questions that arise when defining the boundaries of automated decision-making. This consultation asks precisely the right questions, and does so with a level of specificity that is not yet common among comparable regulatory efforts internationally. We are pleased to offer the following perspectives on scope, human override, configurable software responsibilities, and layered disclosure.

We would be happy to answer any questions you may have and would welcome further collaboration on your work to advance transparency and accountability in automated decision-making across Australia.

[redacted]
Director of AI Policy

[redacted]
Director of Operations, Australia

Identifying computer programs that substantially and directly relate to making decisions

ADM Issues Paper – Question 1

The most important factor in determining whether a computer program is substantially and directly related to making a decision is the nature of its output. ADM transparency obligations should only be considered where a computer program autonomously makes a decision or recommends a single specific course of action that is not expected to be reviewed by a human for accuracy.

Computer programs that surface, compile, summarise, or display information for a human decision-maker to weigh and act upon should not be considered ADM, regardless of how consequential that information may be to the ultimate decision. Almost any computer program could be said to facilitate a decision if the standard is merely that it provides relevant information. Consider a clinical program that organises and displays a patient's known allergies to a prescribing clinician. That information will directly bear on the clinician's prescribing decisions for the patient, but the ultimate decisions on what to prescribe remain with the clinician.

Without this distinction, the scope of the obligations around ADM risks becoming unworkably broad, encompassing virtually any computer program that collects, processes, or displays information.

Ability and likelihood of human override is also a relevant factor in determining whether a computer program is substantially and directly related to making a decision. Where a human retains genuine and practical ability to override a computer program, decision-making authority remains with the human rather than the program. For the purposes of the ADM transparency obligation, the relevant question is whether override is meaningfully available to the human decision-maker.

This factor is also valuable because it allows developers to intentionally design computer programs to function as ADM or not. **A program that presents a recommendation as optional, explains the basis for it, and makes it straightforward to disregard or change with minimal friction is designed to support human judgment rather than replace it.** Conversely, a program that presents outputs as determinative, obscures its reasoning, or makes override cumbersome by design is functioning more like ADM.

Transparency and explainability of outputs is best understood as a contributor to meaningful override rather than a standalone factor. A human decision-maker is more

likely to exercise independent judgment, and to override the output where appropriate, if they understand where a recommendation came from and why. For example, organisations implementing AI tools frequently use them to operationalise policies and procedures that previously existed only on paper. When a user can see that a surfaced recommendation reflects their organisation's own established policy, rather than an unexplained model output, they are better equipped to evaluate it, apply judgment, and override it where circumstances warrant. Transparency therefore matters, but it matters because it enables genuine, informed override.

We do not recommend degree of reliance on ADM output as a determining factor. This would centre the definition of ADM on human behaviour rather than on the design and intended use of the computer program itself. Two individuals could use the same program very differently, one thoughtfully weighing recommendations against their own judgment, the other accepting the program's recommendations unquestioningly. Defining ADM by degree of reliance would therefore cause the obligation to be driven by the "worst case scenario" involving potentially untrained or unconcerned users, rather than by what the computer program is designed to do.

With respect to the triage score example OAIC provides:

We agree that a formula deriving an age from a date of birth is purely informational and should not be considered substantially related to making a decision. The triage score example is closer to the ADM line. We would stop short of concluding that it necessarily crosses it based only on the information provided, regarding its use to prioritise a human decision-maker's attention. Two additional factors should be considered.

- First, whether the computer program itself makes the triage decision or whether it simply surfaces information that users then apply to a pre-existing organisational policy. In the latter case the policy, not the computer program, substantially and directly relates to making the decision.

- Second, whether users are prevented by the computer program from calling patients out of order. If a user retains the practical ability to override the triage order, that weighs against treating the program as ADM.

Fictional edge case: generative AI with human oversight

ADM Issues Paper – Questions 2-3

We do not consider NACIA's described uses of GPTea to be substantially and directly related to making decisions. As described, GPTea is summarising candidate profiles and presenting recommendations about eligibility to human NACIA staff who retain full oversight and make final eligibility determinations. This is a computer program that organises and presents relevant information to support a human decision-maker rather than one that drives outcomes autonomously.

While GPTea's summaries may be substantial in the sense that they provide much of the relevant input information used by the human decision-makers, the program does not appear to be directly related to the decisions. GPTea is summarising existing available information, not making decisions, and nothing in the description suggests that it makes it easier to follow the recommendations than not.

A computer program that opaquely evaluates personal information and recommends outcomes based on that information without citing sources or explaining its reasoning makes it difficult for human staff to exercise genuine independent judgment and is, in our view, more likely to cross the ADM threshold. We would consider such a program a poor use of generative AI, with greater potential to process information in ways that might affect vulnerable people's rights and interests. It would not, in our view, align with APP 1 requirements for using personal information in an open and transparent way.

Impact of decision for persons experiencing vulnerability

ADM Issues Paper – Questions 4-7

The distinction between "access" and "provision" is critical to ensuring that the ADM obligation is applied only to computer programs that are truly impacting rights and interests, rather than sweeping in any program that simply provides information or policy-driven recommendations that are then acted upon by individuals who maintain decision-making autonomy.

In healthcare, for example, a computer program that autonomously determines whether a patient is eligible to receive treatment, or that makes a recommendation about eligibility that is difficult for a human to override, would likely trigger the transparency obligation. Once a patient is already receiving care, however, additional uses of computer programs such as recommending one medication over another does not affect individuals' rights and interests in the same way. A clinician using software to guide how to provide

treatment is a fundamentally different situation from a computer program determining whether a patient can receive treatment in the first place.

This distinction is supported by the language of the issue paper, which refers to admission to a country, entitlement to a housing benefit, and access to healthcare services as examples of decisions that might significantly affect the rights or interests of an individual. Each of these examples describes a threshold determination (whether someone can participate in or receive something) rather than decisions made in the course of already providing a service or benefit for which an individual has already been deemed eligible.

Fictional edge case: differential pricing

ADM Issues Paper – Questions 8-9

Differential pricing is generally considered a form of inhibiting access, since pricing is the primary mechanism by which free markets determine access to goods and services. However, attempting to set a percentage threshold on what constitutes significant differential pricing feels arbitrary and is unlikely to produce consistent or meaningful outcomes.

The significance of the product or service being priced is a far more determinative factor than the dollar amount or percentage difference alone. A modest price differential on an essential service could significantly affect an individual's rights and interests, while a substantial price differential on a discretionary consumer product may not. Guidance should therefore focus on the nature of the product or service rather than attempting to define significance by reference to a numerical threshold.

Fictional edge case: discriminatory targeted job ad

ADM Issues Paper – Question 10

A computer program that autonomously determines to whom to display an advertisement is making a decision. The more important question, however, is whether that decision is significant enough to warrant the ADM transparency obligation.

Applying the obligation to every instance of algorithmic ad targeting could have the unintended consequence of sweeping in virtually all digital advertising. A decision about who sees an advertisement for a consumer product such as shampoo does not rise to that level. A decision about who sees an advertisement for an employment opportunity is clearly more significant, given the potential impact on an individual's livelihood and career prospects.

Meaning of 'arranged for'

ADM Issues Paper – Question 11

Scenarios involving customer-configurable computer programs require guidance. Where a computer program is configurable, it is not always clear whether the software developer that created the program or the entity that configured it, the customer, is the one that has "arranged for" ADM.

We recommend that the guidance make clear that the entity that configures rules or functionality within a configurable computer program is the entity that has arranged for ADM and therefore bears the transparency obligation.

For example, a computer program for hospitals may include a clinical decision support engine which allows each hospital customer to configure rules to evaluate patient data and make recommendations. The software developer created the program, but it is the individual hospital that implements the software for use at its facility and determines what rules to configure and what actions to recommend when those rules are satisfied.

Extent of disclosure

ADM Issues Paper – Question 12

The proposed disclosure guidelines appropriately identify the need to balance comprehensiveness with comprehensibility. We recommend allowing for different levels of disclosure to exist in different locations. For example, an in-workflow disclosure might be relatively concise, prioritising readability, while including a link to more detailed information for those who wish to access it. This layered approach allows entities to communicate meaningfully with the majority of individuals without overwhelming them, while still making fuller information available to those who want it.