Privacy Act Reform for the Age of AI: Accountability, Rights, and Information

Ignacio Cofone

Professor of Law and Regulation of AI, Faculty of Law and Institute for Ethics in AI, University of Oxford

Suggested citation: Ignacio Cofone, 'Privacy Act Reform for the Age of AI: Accountability, Rights, and Information' (University of Oxford Institute for Ethics in AI White Paper, May 2026)

Introduction

The *Privacy Act*'s inadequacy for artificial intelligence (AI) is evident. The Act was built on a collection-centred model: it regulates what government institutions collect from individuals, it notifies individuals of collection, and it grants individuals access to what was collected about them. AI breaks the usefulness of each of these mechanisms. It generates personal information through inference from data not collected for that purpose, produces distributional harms that fall on populations rather than on identifiable complainants, and relocates decision-making authority from officials whose reasoning can be examined to algorithmic systems whose designers sit outside the Act's coverage. No incremental amendment of individual provisions will address these mismatches; what is required is a reorganization of the Act.

If reform similar to the *Consumer Privacy Protection Act* were to pass without parallel reform of the *Privacy Act*, the federal government will hold private companies to a more rigorous regime than it applies to its own institutions, including automated decision-making obligations. The Supreme Court has characterized the *Privacy Act* as quasi-constitutional. A quasi-constitutional statute cannot be less protective than commercial legislation.

This submission responds to the Treasury Board of Canada Secretariat's 2026 review of the Act, which sets out twenty-three proposals for legislative modernisation across six themes. The proposals under review represent the most serious federal attempt at *Privacy Act* reform in four decades. Most should be adopted substantially as proposed. Others require strengthening in specific technical respects, particularly where the drafting reflects a collection-centred model that technology has transcended. This submission addresses eleven of the twenty-three proposals, which together form the Act's accountability and enforcement mechanisms: Proposals 3, 5, 6, 7, 10, 12, 13, 14, 19, 21, and 22. It does not address proposals concerning data sharing for service integration, breach notification mechanics, institutional safeguard requirements, Indigenous data governance, or procedural harmonization of access regimes, which warrant separate treatment.

Three themes run through this submission. First, transparency and accountability obligations must attach to the inference step, meaning the algorithmic generation of personal

information from existing data, not only to collection from individuals or to the output of final decisions. The proposed definition of 'creation of personal data,' the necessity test, and the ADS transparency requirements each require amendment to achieve this. Second, the sensitivity of personal information cannot be determined solely at the point of collection when AI processing will generate sensitive outputs from non-sensitive inputs: the Act requires an inference-based sensitivity trigger alongside the collection-based sensitivity categories. Third, the enforcement regime must include direct binding order-making powers and financial penalties to be adequate for AI-scale harms. The corrective action plan under Proposal 19 is an improvement over the current ombudsman model but it is not sufficient.

Proposals 3 and 5: Monitoring and Transparency

Proposal 3: Privacy Impact Assessments as a Legal Requirement

The proposal to make Privacy Impact Assessments a statutory requirement should be adopted. The existing policy requirement under the *Directive on Privacy Practices* creates no enforceable obligation and generates no individual rights. The consequences of this are visible in institutional practice: PIAs are inconsistently conducted, their quality varies across departments, and the OPC cannot compel remediation of deficiencies identified during a PIA review even where the deficiency is clear. Legislating the PIA obligation resolves these deficiencies and gives the Commissioner a foothold from which to assess pre-deployment risk.

The proposal's central limitation is its triggering condition. It applies when a program 'uses personal data to make a decision about someone.' This formulation, while appropriate for the paper-based government the Act was designed to govern, is inadequate for AI systems that generate population-level harms without making individual decisions. A facial recognition system deployed at an airport does not make a 'decision about someone' until a positive match is flagged; the harm-generating processing, the inference of identity from visual characteristics applied to every person passing through, precedes any identifiable decision. Risk assessment algorithms that produce recidivism scores, predictive policing tools that allocate surveillance resources, and eligibility models that generate benefit denial recommendations all affect populations in ways that individual-decision-level PIA reviews cannot capture.

Australia's automated debt recovery scheme matched tax and welfare records at population scale, generating systematic over-recovery across tens of thousands of cases. The Royal Commission's 2023 report found that the scheme's errors were detectable at the design stage and were not detected because no design-stage accountability norm existed and no population-level assessment was required before deployment.[1] The triggering condition the proposal adopts would not have captured the Robodebt scheme, because the scheme did not

[1]Australia, Commonwealth, Royal Commission into the Robodebt Scheme, *Report* (July 2023) at 17–23.

make a 'decision about someone' at any single moment that would have triggered a PIA: it made decisions about everyone through a process no individual could observe.

The trigger for a mandatory PIA should be whether an AI system processes personal information at scale or in a way that could produce distributional effects on a population. This scope is consistent with the EU AI Act's risk-based classification and its pre-deployment assessment obligations: the Act requires public sector bodies deploying high-risk AI systems to conduct a Fundamental Rights Impact Assessment before the system is put into operation, covering the categories of persons affected, the risks to their rights, and the human oversight measures in place regardless of whether any specific individual decision has yet been made.[2] It is also consistent with other international examples, such as New Zealand's Algorithm Impact Assessment framework, which applies to consequential algorithms deployed by public agencies generally,[3] and the United States, where federal agencies have been required since 2024 to publish annual inventories of AI use cases with enhanced detail for rights-impacting systems.[4]

Three additional elements should be incorporated into the statutory PIA requirement. First, the PIA should mandate an assessment of whether the AI system operates accurately and without disparate error rates across demographic groups defined by protected characteristics. A PIA that addresses collection risk without assessing algorithmic accuracy cannot identify the pattern of systematic error that the Robodebt scheme represents. Second, the PIA should require institutions to verify that third-party data suppliers whose data feeds the AI system collected that data in compliance with the Act. This codifies the implicit obligation the OPC developed in its investigations of the RCMP's use of Clearview AI and of Project Wide Awake, an obligation the RCMP has contested precisely because it rests on interpretive rather than statutory authority.[5] Giving it statutory force through the PIA requirement rather than through section 4 interpretation alone resolves a six-year dispute that remains unresolved. Third, PIA summaries published under the plain-language requirement should extend to AI system documentation held by private contractors: design specifications, training data descriptions, and validation results. The *Access to Information Act*'s concept of records 'under the control of' an institution (section 4(3)) may reach some contractor-held documents, but

[2]Regulation (EU) 2024/1689 (EU AI Act), Art 27 (requiring deployers that are bodies governed by public law, prior to deploying a high-risk AI system, to perform an assessment of the impact on fundamental rights covering affected persons, risks, and oversight measures); Title III (risk-based classification and conformity assessment requirements). The Article 27 obligation applies to the first use of a high-risk system and must be notified to the relevant market surveillance authority.

[3]New Zealand, Algorithm Impact Assessment User Guide (data.govt.nz, 2023); the framework applies to consequential algorithms deployed by public agencies, not only to those making specific individual determinations.

[4] US Office of Management and Budget, Memorandum M-25-21, *Accelerating Federal Use of AI through Innovation, Governance, and Public Trust* (April 2025).

[5]OPC, *Special Report to Parliament: Investigation into the RCMP's Use of Clearview AI* (2021); OPC, *Special Report to Parliament: Investigation of the RCMP's Collection of Open-Source Information under Project Wide Awake* (2024). In each investigation, the OPC developed an implicit obligation under s 4 of the *Privacy Act* requiring institutions to verify the lawfulness of third-party data suppliers' practices; the RCMP disputed this reading while accepting the OPC's recommendations.

its application in the AI context is uncertain.[6] A statutory publication obligation in the *Privacy Act* eliminates that uncertainty and reduces the burden on individuals to identify the right records across complex procurement arrangements.

Two objections arise. The first is that the *Directive on Automated Decision-Making* already provides the relevant transparency and accountability requirements. It does not. The Directive applies to approximately ninety-seven institutions subject to the *Policy on Service and Digital* (broadly, departments listed in Schedule I of the *Financial Administration Act*, but not the full range of entities that Act governs); the *Privacy Act* applies to over two hundred and fifty. The Directive creates no individual rights and carries no enforcement consequence independent of the *Privacy Act*'s own complaint mechanism; a statutory PIA obligation creates both. Legislating the PIA obligation is not duplication of the Directive; it is extension to the full scope of the federal government and enforcement against all institutions the Act covers.

The second objection is cost and administrative burden. The Robodebt Royal Commission documented what deploying an automated system without pre-deployment accountability costs: the Australian government faced a settlement totalling approximately AUD$1.8 billion (comprising refunds of approximately AUD$746 million wrongly recovered, the write-off of further outstanding debts, and associated legal costs) arising from incorrect debt notices issued to hundreds of thousands of welfare recipients, bore the full cost of a Royal Commission, and faced documented harm including fear, distress, and in some cases self-harm and death associated with incorrect debt notices.[7] The cost of a Privacy Impact Assessment before deployment, comprising specialist review, documented necessity analysis, third-party verification, and demographic accuracy assessment, is trivial against this baseline. The discipline a mandatory PIA imposes is the discipline of specifying what the system does, on whom, and with what accuracy before deployment. Institutions with well-designed programs will satisfy this requirement without difficulty; institutions that cannot satisfy it have identified a problem the PIA is designed to detect.

> **Recommendation:** Adopt the PIA obligation with an amended trigger covering any AI system that processes personal information at scale or in a way likely to produce distributional effects; require mandatory demographic accuracy validation, third-party data supplier verification, and publication of contractor-held AI documentation as statutory PIA components.

Proposal 5: Transparency Requirements for AI and Automated Decision Systems

The proposal to establish statutory transparency requirements for AI and automated decision systems should be adopted. The *Directive on Automated Decision-Making* already articulates the relevant principles in graduated form. The proposal correctly identifies the Directive's

[6] *Access to Information Act*, RSC 1985, c A-1, s 4(3); the scope of this provision in the AI procurement context has not been judicially determined.

[7] Australia, Commonwealth, Royal Commission into the Robodebt Scheme, *Report* (July 2023) at 1--2 (describing the scheme as 'a costly failure of public administration, in both human and economic terms'); see also *Prygodicz v Commonwealth of Australia* [2021] FCA 634 (Murphy J), approving the class action settlement.

limitations: a policy instrument that applies to approximately one hundred departments, creates no individual rights, and carries no independent legal consequence cannot provide the accountability that a public-sector AI governance framework requires.[8] The Act should legislate what the Directive provides by policy.

The proposal rests on two components: an explanation right (requiring institutions to explain, upon request, how an ADS supported a decision and what personal data was used) and a human review right (allowing individuals to request reconsideration if they believe the ADS used incorrect or incomplete data). Both should be adopted. Two refinements are necessary.

The first concerns the scope of the explanation right. The proposal limits it to ADS that 'use personal data to make or support decisions that directly affect individuals.' This attaches the transparency obligation to the final decision, not to the inference step that precedes and shapes it. The Court of Justice of the European Union's 2023 decision in *SCHUFA Holding* established that the generation of a probability value from personal data, a credit score computed by an automated system, itself constitutes an 'automated individual decision' within the meaning of Article 22(1) GDPR (and therefore engages Article 22's prohibition and its attendant rights, including the right to human intervention and the right to contest the decision) where the score plays a determining role in the downstream outcome.[9] The inference step is legally significant independently of the downstream decision. The Canadian Act should reflect this understanding: the explanation right should attach whenever an ADS generates an inference, prediction, or recommendation from personal data that materially influences a consequential decision, whether or not the final determination is formally made by a human official.

This matters in government contexts. A benefit eligibility model that assigns a risk score to an application, which a human caseworker then reviews and almost always accepts, is making the decision at the scoring stage. Transparency obligations that attach only to the caseworker's formal determination provide no insight into the stage at which the consequential choice is made. The 'human in the loop' exception has weakened Article 22 of the GDPR in practice for this reason: nominal human review of algorithmic recommendations cannot substitute for genuine accountability at the inference stage.[10]

The United Kingdom has recognized this through the Algorithmic Transparency Recording Standard, which since 2024 has been mandatory for central government departments and

[8]Treasury Board of Canada Secretariat, *Directive on Automated Decision-Making* (2019, amended 2023); Teresa Scassa, 'Administrative Law and the Governance of Automated Decision-Making' (2021) 55 UBCLR 45 at 63.

[9]Case C-634/21, *SCHUFA Holding* [2023] ECLI:EU:C:2023:957 (CJEU), holding that automated generation of a credit-scoring probability value constitutes an automated decision under GDPR Art 22 even where the final lending decision is made by a human.

[10]Lilian Edwards & Michael Veale, 'Slave to the Algorithm? Why a Right to an Explanation Is Probably Not the Remedy You Are Looking For' (2017) 16 Duke L & Tech Rev 18. In *NJCM et al. v The Netherlands* (District Court of The Hague, 5 February 2020) ECLI:NL:RBDHA:2020:1878, the court struck down the Dutch government's SyRI welfare fraud detection algorithm under Art 8(2) ECHR, finding that the system's opacity made it impossible for affected individuals to understand or contest its outputs, and that the government had not demonstrated its necessity and proportionality.

public-facing arm's-length bodies: it requires each to publish structured records of algorithmic tools used in decisions with public effect, covering the tool's purpose, input data, outputs, and human review arrangements.[11] France has gone further. Under the *Code des relations entre le public et l'administration*, any individual administrative decision based on algorithmic processing must include explicit notice to the affected person, who may then request disclosure of the algorithm's contribution to the decision, the data sources used, the parameters applied and their weighting, and the operations performed in intelligible form, subject to enumerated secrecy exceptions.[12] At municipal level, Amsterdam and Helsinki have each operated public algorithm registers since 2020, documenting the purpose, datasets, governance, and oversight arrangements for algorithmic tools used in city services, showing that the disclosure is workable at different scales of public administration.[13]

The second concerns the right to human review. The proposal frames it as a mechanism to allow individuals to check whether data 'is accurate and ask for corrections if needed.' This is too narrow. Human review should require the reviewer to assess not only whether the specific inputs in the individual case were accurate, but also whether the system's general operation produced a fair outcome for that individual. Where the review reveals a pattern, specifically the same type of error occurring consistently for members of a protected group, the reviewer should be required to refer the concern to the institution's designated privacy officer for systemic investigation. Linking individual review to systemic investigation is how the human review right addresses collective harms.

The Act should also make explicit the distinction between pre-decision human oversight, which the *Directive on Automated Decision-Making* already requires at certain impact levels and which the Act should preserve, and post-decision human review, which the proposal introduces. Pre-decision oversight catches errors prospectively; post-decision review generates redress and systemic evidence retrospectively. Both are necessary and serve different purposes. The Act should specify that the post-decision right supplements rather than replaces any pre-decision oversight requirements.

The right would be exercised selectively: primarily by individuals who receive adverse and unexpected outcomes and who have already been informed of its existence through the notice obligation under Proposal 6. A system producing accurate and unbiased results will generate few review requests; a system generating many requests has identified a signal that needs systemic investigation. The cost of at-scale AI error in government, including wrongly denied benefits, unjust detentions, and misidentification, exceeds the administrative cost of review on request. Human review requirements create incentives for accuracy and bias mitigation that the current framework, which has no equivalent mechanism, cannot generate.

---

[11] UK Government Digital Service, *Algorithmic Transparency Recording Standard: Guidance for Public Sector Bodies* (GOV.UK, first published January 2023, mandatory scope confirmed February 2024); see also *Algorithmic Transparency Recording Standard: Mandatory Scope and Exemptions Policy* (GOV.UK, December 2024).

[12] *Code des relations entre le public et l'administration* (France), Art L 311-3-1 (created by Loi no 2016-1321 du 7 octobre 2016, art 4); implementing Decree no 2017-330 du 14 mars 2017, Art R 311-3-1-2.

[13] City of Amsterdam, *Algorithm Register* (algorithmregister.amsterdam.nl, operational since 2020); City of Helsinki, *Artificial Intelligence Register* (ai.hel.fi, operational since 2020).

> **Recommendation:** Adopt the ADS transparency provisions with the explanation right extended to the inference step that materially influences a consequential decision; broaden the human review right to require systemic fairness assessment with mandatory referral obligations where patterns of group-level error are detected; and specify that the post-decision review right supplements rather than replaces pre-decision oversight requirements.

Proposals 6, 7, and 10: Data Collection

Proposal 6: Privacy Notice Requirements, Including for Automated Decision Systems

The proposal to strengthen privacy notice requirements addresses deficiencies: notices are currently hard to find, written in technical language, not reliably provided when ADS are used, and not connected to any centralized repository. Plain language requirements, central registry publication, and notification when ADS process personal data should be adopted. One element requires tightening.

The proposal provides that, for decisions made or supported by ADS, notice should be given 'before or as soon as reasonably possible after the system starts using their personal data.' The breadth of the post-processing alternative risks eliminating the notice's accountability function for consequential decisions. If notice may follow the processing and the outcome it produces, it serves only a historical record function: the individual learns, after the fact, that their data was processed in a system they could not observe or contest at the relevant time. For decisions affecting access to benefits, immigration status, legal rights, or public service eligibility, pre-processing notice should be required, not an available alternative. Post-processing notice should be confined to operational contexts where pre-processing notice is impracticable, a category that should be defined narrowly in the Act and subject to justification on request to the Commissioner.

The content of ADS notices should be specified in the Act rather than left entirely to policy elaboration. A notice that provides 'a general explanation so the person can understand how the ADS handled their personal data' is consistent with the Act's current transparency deficiencies: in practice, general explanations of automated processes routinely tell individuals that their data was used in a computerized decision without providing the specific information needed to assess or challenge the outcome. The Act should require that ADS notices specify: the categories of personal data used; whether the decision was fully automated or substantially AI-assisted; the primary factors influencing the outcome; and the individual's rights to request a full explanation and human review. Quebec's Law 25 provides the closest domestic comparator, granting individuals the right to know the personal information used and the reasons and factors that contributed to the decision.[14] The federal *Privacy Act* should extend equivalent specificity to the public sector.

[14] *An Act Respecting the Protection of Personal Information in the Private Sector*, CQLR c P-39.1, as amended by Law 25 (SQ 2021, c 25), ss 12.1, 12.2.

> **Recommendation:** Adopt the notice requirements with mandatory pre-processing notice for decisions affecting legal rights and benefits; confine post-processing notice to narrowly defined impracticable contexts subject to justification; and specify ADS notice content requirements in the Act to include data categories used, automation level, primary decision factors, and the individual's rights to explanation and review.

Proposal 7: Spectrum of Data Sensitivity and Identifiability

The proposal's recognition of a spectrum of identifiability, distinguishing anonymized data, de-identified data, and identifiable personal data and attaching differentiated safeguard requirements to each, should be adopted in full. The Act's current treatment of de-identified data as simply outside its scope is untenable in the AI context. 99.98% of Americans remain uniquely identifiable in any dataset using fifteen demographic attributes, with the study finding that even heavily sampled anonymized datasets are unlikely to satisfy modern anonymization standards, challenging the assumption that de-identification and sampling provide adequate protection.[15] Treating de-identified data as outside the Act's coverage while organizations use it to make probabilistic inferences about individuals on the basis of group membership is a gap that this proposal correctly closes.

The recognition of sensitive categories, including racial or ethnic origin, political opinions, religious beliefs, genetic data, biometric data, health data, and sexual orientation, should be enacted. Treating breaches of sensitive data as automatically material, and requiring the highest safeguards for their collection and retention, provides a baseline protection consistent with GDPR Article 9's special categories framework and with modern privacy statutes across comparable jurisdictions.

The sensitivity spectrum, as drafted, has one limitation: the enumerated categories are defined by reference to the data at the point of collection. AI systems routinely derive sensitive information from data that does not fall within any sensitive category when collected. Location data yields inferences about religious practice and political participation. Purchase patterns yield inferences about sexual orientation and health conditions. Facial imagery, collected to verify liveness for fraud prevention, yields racial classification. In each case, the input data is non-sensitive under the proposal's definition; the output inference is highly sensitive. A sensitivity framework that attaches only to data as collected provides no constraint on inference from non-sensitive inputs to sensitive outputs.

The inference step is how AI systems generate privacy harms at scale. The Article 29 Working Party identified the problem in its 2017 guidelines: it is not the information collected in itself that is sensitive, but rather the inferences drawn from it, that most frequently generate rights violations.[16] The OPC's own consultation paper for the reform of private-sector privacy legislation concluded that inferences can be as harmful to individuals and groups as collected

[15]Luc Rocher, Julien Hendrickx & Yves-Alexandre de Montjoye, 'Estimating the Success of Re-Identifications in Incomplete Datasets Using Generative Models' (2019) 10 Nature Communications 3069.

[16]Article 29 Data Protection Working Party, *Guidelines on Automated Individual Decision-Making and Profiling for the Purposes of Regulation 2016/679* (WP 251rev.01, 6 February 2018) at 17-18.

information and recommended that personal information definitions encompass inferences.[17] A sensitivity classification applied at collection addresses the data; it does not address what the data becomes.

The Act should include an inference-based sensitivity trigger alongside the collection-based sensitivity categories. Data should be treated as subject to the highest safeguards, and its breach as automatically material, when it is processed in a way that is reasonably likely to generate outputs within a sensitive category, regardless of whether the input data itself falls within that category. The Privacy Impact Assessment process under Proposal 3 provides the vehicle for this determination: institutions should be required, as a component of the PIA, to identify whether their AI system can generate sensitive-category outputs and to apply enhanced safeguards accordingly. The OPC should be empowered to issue binding guidance calibrating the 'reasonably likely' standard to specific types of systems and input combinations.

> **Recommendation:** Adopt the identifiability spectrum in full; supplement the enumerated sensitivity categories with an inference-based trigger, operationalized through the mandatory PIA, covering AI systems reasonably likely to generate sensitive-category outputs from non-sensitive inputs; and empower the OPC to issue binding guidance on the standard.

Proposal 10: Necessity Test for the Collection Threshold

The proposal to add a statutory necessity test to section 4 should be enacted in full. The four-part test, comprising legal authority, necessity, effectiveness, and minimal intrusiveness, corrects a documented gap between the Act's current standard and the constitutional standard articulated in *R v Spencer*.[18] The current 'directly related' standard is a threshold of proximity, not necessity: an institution that can show a substantial connection between personal data and its broad program mandate satisfies the standard without demonstrating that collection is the least privacy-intrusive means of achieving the program's objectives. In an environment where the marginal cost of collecting additional data approaches zero and machine learning systems extract value from virtually any data point, the proximity standard provides no meaningful constraint on systematic over-collection. The OPC recommended replacing it with a necessity and proportionality standard in 2006 and again in 2016; Parliament has not acted.[19]

Two extensions are necessary to make the necessity test adequate for the AI context.

The first extension addresses inference generation. The proposal targets over-collection: it requires that collection be limited to what is 'reasonably required to achieve a clearly defined

[17]Ignacio Cofone, *Policy Proposals for PIPEDA Reform to Address Artificial Intelligence* (OPC, November 2020) at 8.

[18]*R v Spencer*, 2014 SCC 43 at para 69 (Cromwell J), establishing a 'reasonably necessary' standard for state information demands under s 8 of the *Charter*.

[19]OPC, *Privacy Act Reform in an Era of Change and Transparency* (Submission to ETHI Committee, 2016) at 6; David H Flaherty, 'Reflections on Reform of the Federal Privacy Act' (2008) 21 Can J Admin L & Prac 272 at 285.

purpose.' It does not address the subsequent processing step in which existing data is run through an inference model to generate new personal information. A government institution that collects only what is reasonably necessary to administer a program but then feeds that data into a predictive model to generate estimates of recidivism likelihood, immigration flight risk, or benefit fraud probability has satisfied the necessity test on the proposal's current terms while generating the privacy harms that the test is designed to prevent. The Act should specify that the necessity standard applies when institutions generate personal information from existing data through algorithmic processing, not only when they collect from external sources. The proposed definition of 'creation of personal data' in Proposal 14, which treats inference as a form of collection, must be aligned with the necessity test so that the same standard governs both steps.

The second extension addresses the proposal's 'effectiveness' criterion: that 'the data collected is likely to achieve that purpose.' For AI systems, this criterion should include a requirement that the institution has validated, before deployment, that the system achieves its stated purpose without producing disparate outcomes across demographic groups defined by protected characteristics. An algorithmic risk assessment tool that predicts recidivism with reasonable overall accuracy but misclassifies individuals from a particular racial group at substantially higher rates does not achieve 'accurate risk assessment' consistently with the Act's quasi-constitutional status. Demographic validation, meaning assessment of whether the system's error rates are consistent across groups, is a pre-condition for claiming that collection satisfies the effectiveness criterion. This requirement connects the necessity test to the PIA obligation proposed under Proposal 3 on a common factual footing.

The objection that a statutory necessity test imposes unfeasible operational requirements is not supported by comparable regimes. The *Directive on Privacy Practices* already requires institutions to limit the collection of personal data to what is 'demonstrably necessary' for their programs; the four-part test in this proposal is substantively similar. Legislating the necessity standard does not change what institutions are already required to do by policy; it provides the enforcement that the policy requirement lacks. The test requires institutions to specify what data they need, why they need it, and why less intrusive alternatives are inadequate. For AI systems it adds one requirement: pre-deployment validation that the system works as claimed across the relevant population. Institutions with clear mandates and well-designed systems will satisfy it; institutions that cannot do so will have identified the problem the test exists to detect.

> **Recommendation:** Adopt the four-part necessity test; extend it explicitly to inference generation as defined in Proposal 14; and interpret the effectiveness criterion to require pre-deployment demographic accuracy validation confirming that the system does not produce disparate outcomes across protected groups.

Proposal 12, 13, and 14: Fundamental Rights and Principles

Proposal 12: Privacy as a Fundamental Right

The proposal to amend the Act's purpose clause and preamble to recognize privacy as a fundamental right should be adopted. The Supreme Court's characterization of the Act as 'quasi-constitutional' in *Lavigne v Canada* and confirmed in *Dagg v Canada* reflects the normative position that privacy is not merely an administrative interest but a foundational value that constrains state conduct.[20] Codifying that characterization in the Act's own statement of purpose would align the Act with comparable statutes internationally, including the GDPR's explicit recognition of data protection as a fundamental right under the EU Charter of Fundamental Rights, and would provide interpretive guidance that the Act's current purpose clause fails to supply.

Two supplements to the proposal are necessary.

The first concerns the adequacy of rights-based framing alone. Recognizing privacy as a fundamental right is necessary but not sufficient to address the power asymmetry that AI creates. Individual rights, however generously interpreted, cannot remedy distributional harms that are invisible to any individual complainant. An institution that deploys a recidivism algorithm producing disparate error rates across racial groups has violated the privacy and equality interests of a population of affected individuals, none of whom can identify the harm without access to the system's aggregate performance data. A right-bearer who cannot detect that a right has been violated cannot vindicate it through individual complaint procedures. The purpose clause should therefore recognize systemic institutional accountability, not merely individual rights, as a co-equal objective of the Act. Rights without institutional accountability mechanisms are claims without effective remedies.

The second concerns the inclusion of 'enabling service delivery' as an additional statutory objective. This is legitimate in principle: the Act currently provides no positive basis for responsible data reuse in the public interest, and service integration objectives should be expressly recognized. But a purpose clause that treats service efficiency as a co-equal object of legislation alongside a fundamental right invites interpretive balancing that treats privacy as a constraint to be traded off against service objectives rather than as a foundational value within which service delivery must operate. The purpose clause should be explicit that service delivery operates within the constraints of the fundamental right to privacy, not as a qualification of it. The GDPR, which recognizes the free flow of personal data as an ancillary objective that must be consistent with a high level of protection, provides the appropriate model.[21]

> **Recommendation:** Adopt the fundamental rights amendment; add systemic institutional accountability as a co-equal statutory objective alongside individual rights; and draft the service delivery objective as explicitly subordinate to privacy protection rather than co-equal with it.

[20] *Lavigne v Canada (Office of the Commissioner of Official Languages)*, 2002 SCC 53 at para 25; *Dagg v Canada (Minister of Finance)*, [1997] 2 SCR 403 at 426 (La Forest J).

[21] Regulation (EU) 2016/679 (GDPR), Art 1(3) (free flow of personal data as an objective that must be consistent with a high level of protection of natural persons).

Proposal 13: Privacy Principles in the Privacy Act

The proposal to incorporate privacy principles into the Act's operative provisions should be adopted. The existing regime places principles in Schedule 1 in non-legal language that mixes obligations with best practices; moving principles into the Act's operative provisions, as rights and obligations rather than aspirations, aligns the Act with PIPEDA, with the approach the *Consumer Privacy Protection Act* would take in the private sector, and with the OECD 2013 Privacy Guidelines that have served as the international baseline for four decades.

Privacy by design requires elaboration. It is a Canadian conceptual innovation, developed by Ann Cavoukian as Information and Privacy Commissioner of Ontario beginning in the 1990s, subsequently codified in GDPR Article 25, and recently incorporated in Australia's privacy reforms following the Robodebt Royal Commission.[22] Canada's federal public-sector statute lacks it despite its presence in comparable international regimes. Codifying privacy by design should require that institutions implementing AI systems embed data minimization, necessity, and de-identification from the design stage. The Robodebt Royal Commission's findings confirm the value of this sequencing: the scheme's systematic errors were detectable at the design stage and were not detected because no design-stage accountability norm existed.[23]

Privacy by design should include, as an explicit component, the obligation to verify that third-party data suppliers whose data feeds an AI system collected that data in compliance with the Act. This requirement codifies the OPC's implicit section 4 obligation developed in its *Clearview AI* and *Project Wide Awake* investigations, an obligation the RCMP has contested because its implicit, interpretive character leaves the legal position uncertain.[24] Giving this obligation statutory form through the privacy-by-design principle resolves that dispute.

Principles, to be effective, must work as legal constraints rather than solely as interpretive resources. The Act should empower the OPC to issue binding guidance specifying how each principle applies in the AI context, and that guidance should be subject to periodic review. Without this, principles function as contested reference points: the history of PIPEDA demonstrates that principles-based legislation without binding regulatory elaboration produces litigation over interpretation rather than clear compliance incentives. The *Consumer Privacy Protection Act* addressed this by including order-making power and administrative guidance authority; the reformed *Privacy Act* should do the same.

The consultation document proposes that principles would be outlined at the outset of the Act and integrated into its provisions. That sequencing is right. The question is how. A principles-based statute is not a statute with a list of principles followed by substantially unchanged operative provisions; it is a statute whose operative provisions are themselves

[22]Ann Cavoukian, *Privacy by Design: The 7 Foundational Principles* (Toronto: Information and Privacy Commissioner of Ontario, 2009); GDPR, Art 25; Privacy and Other Legislation Amendment Act 2024 (Cth) (Australia).

[23]Robodebt Royal Commission Report, above n 1, at Recommendations 17.1–17.2 (calling for a legislative framework for automated decision-making and an independent audit body).

[24]OPC, Clearview Report, above n 4; OPC, Wide Awake Report, above n 4.

organized around and expressed through those principles. The *Consumer Privacy Protection Act* provides a model: it incorporates principles into operative provisions, so that accountability, necessity, and accuracy appear as statutory obligations generating enforceable rights. The reformed *Privacy Act* should follow this structure. Each principle should receive operative expression in the Act's body: necessity as a collection threshold in section 4, privacy by design as a pre-deployment obligation, accountability as a statutory duty with institutional consequences. Principles that appear only in a prefatory statement and are then applied by reference to independently drafted operative provisions function as interpretive guides rather than legal obligations, informing the construction of provisions without themselves generating rights or duties. This distinction matters for enforcement: an OPC investigation grounded in a statutory obligation to design for privacy produces different and more determinate outcomes than one grounded in the interpretive weight of a prefatory principle.

> **Recommendation:** Adopt the principles with OPC binding guidance authority; incorporate each principle as an operative obligation generating enforceable rights in the Act's body, following the CPPA model, rather than as a prefatory statement applied to independently drafted provisions.

Proposal 14: Key Definitions

The proposed definitions collectively address a drafting deficit: the Act's 1983 definitions were adequate for a paper-based information environment and are inadequate for one in which inference, de-identification, and automated processing are routine. Three definitions are consequential for the AI context: 'creation of personal data,' 'automated decision system,' and 'sensitive data.'

The proposed definition of 'creation of personal data,' treating derivation or inference of personal data from existing personal data as itself a form of collection, should be adopted.[25] The gap this definition closes is real: the Act's collection limitation in section 4 does not, on its plain terms, apply to internal processing that generates new personal information from existing data. A government institution that collects ordinary administrative information, including address, employment status, and prior government interactions, and feeds it through a predictive model to generate estimates of health risk, financial reliability, or political inclination has not, on a straightforward reading of the existing section 4, 'collected' those estimates. The OPC has attempted to address this gap through interpretation of section 4, but the RCMP's continued resistance to the OPC's Clearview finding confirms that interpretation cannot substitute for statutory text. Defining inference as collection, and subjecting it to the same necessity and proportionality constraints as direct collection, resolves the gap.

This definition must be read together with the necessity test in Proposal 10. Defining inference as collection is effective only if the necessity standard that governs collection also

[25]OPC, Clearview Report, above n 4. The RCMP accepted the OPC's recommendations while maintaining disagreement with the finding that s 4 was contravened, which is the characteristic failure mode of the ombudsman model against powerful institutional actors.

governs the inference step. The Act should make explicit that the statutory necessity standard applies to inference generation as defined, and the OPC's binding guidance should specify what necessity and proportionality require at the inference stage in practice.

The proposed definition of 'automated decision system,' comprising any technology that assists or replaces human judgment through rule-based systems, regression analysis, predictive analytics, machine learning, deep learning, or neural networks, is appropriate and consistent with Bill C-27's formulation.[26] It should be read to encompass systems that generate inferences or predictions that materially influence decisions, not only systems that make final determinations. Without this reading, the definition is susceptible to the 'human in the loop' evasion: nominal human review of algorithmic recommendations does not displace the legal significance of the inference step that precedes it.

The proposed definition of 'sensitive data' raises the inference concern developed under Proposal 7 and should be supplemented accordingly. The definition as proposed identifies categories determined by the data at collection. The Act should provide that data is also treated as sensitive, for purposes of safeguard requirements, breach notification thresholds, and the enforcement regime, when it is processed in a way that is reasonably likely to generate outputs falling within a sensitive category, regardless of whether the input data itself falls within any sensitive category. The Privacy Impact Assessment process provides the vehicle for this assessment: institutions should document whether their system can generate sensitive-category outputs and apply enhanced safeguards when it can. The removal of the requirement that personal information be 'recorded' should be adopted. The Act should specify separately that the right of access under section 12 continues to apply only to recorded personal data, since access to an unrecorded processing event is impracticable by definition.

> **Recommendation:** Adopt all proposed definitions; align 'creation of personal data' explicitly with the necessity test in Proposal 10; supplement 'sensitive data' with an inference-based trigger operationalized through the PIA process; and read 'automated decision system' to encompass systems generating inferences that materially influence decisions, not only systems making final determinations.

Proposals 19, 21, and 22: Updating the Compliance Framework

The consultation's enforcement proposals are a necessary component of any adequate reform. Each requires strengthening.

Proposal 19: Binding Order-Making Powers

The proposed corrective action plan advances the current ombudsman model by requiring institutions to develop and publish remediation plans and by authorizing Federal Court

[26] *Consumer Privacy Protection Act* (Bill C-27, Part 1, lapsed on prorogation of Parliament, January 2025; not enacted), s 2 (definition of 'automated decision system'); Treasury Board of Canada Secretariat, *Directive on Automated Decision-Making* (2019, amended 2023), s 4.

referral for non-compliance, mechanisms the current regime, which permits institutions to accept a finding while disputing its legal basis, does not provide. The CAP mechanism should be adopted.

It is not sufficient for the AI context. The mechanism operates in three steps: a finding, a corrective action plan, and a court referral if the institution fails to comply. This produces a minimum of two additional procedural stages between a finding of non-compliance and any binding legal consequence. AI systems that process millions of records and generate thousands of consequential decisions per hour cannot wait for a corrective action plan to be developed, published, implemented, and then litigated before compliance is achieved. Slow enforcement is inadequate for a fast harm-generating technology. The proposed reform mitigates this mismatch without resolving it.

The Commissioner should be granted direct binding order-making power to require institutions to modify, suspend, or discontinue AI systems found to violate the Act, without an intermediate corrective action plan requirement for urgent or systemic violations. This authority has clear precedent in provincial regimes: the Information and Privacy Commissioners of Ontario, British Columbia, and Alberta exercise binding order-making powers, demonstrate high institutional compliance rates, and maintain cooperative relationships with the entities they regulate.[27] The experience of those regimes confirms that it is the availability of binding orders, not their frequent exercise, that disciplines institutional behaviour. Binding orders are rarely issued, but their prospect changes how institutions respond to investigations and recommendations.

Financial penalties are essential and are not contemplated by the proposal. Without a financial penalty mechanism, the CAP regime produces accountability through reputational pressure alone, a mechanism that is weaker against government departments than against private companies because departmental reputations are not measured in market terms. Financial penalties calibrated to institutional size and violation severity create deterrence that reputational pressure cannot, and they generate upstream incentives for investment in system quality and bias mitigation before deployment. The asymmetry argument for financial penalties is strong: if the *Consumer Privacy Protection Act* is enacted in its current form, the federal government will hold private companies handling Canadians' data to a more rigorous enforcement regime, including administrative monetary penalties, than it applies to its own institutions.[28] A statute the Supreme Court has recognized as quasi-constitutional cannot be less protective than ordinary commercial legislation.

A constitutional objection to this might be that an officer of Parliament exercising adjudicative functions raises separation of powers concerns. This objection has been tested and resolved in the provincial context. The Information and Privacy Commissioners of Ontario, British Columbia, and Alberta have exercised binding order-making powers for decades without successful constitutional challenge. The *Information Commissioner* exercises

[27]David H Flaherty, 'Reflections on Reform of the Federal Privacy Act' (2008) 21 Can J Admin L & Prac 272 at 297.

[28]*Consumer Privacy Protection Act* (Bill C-27), ss 93–95 (administrative monetary penalties); s 103 (private right of action).

binding order-making power under the *Access to Information Act* following the reforms introduced by Bill C-58 (reforms that the Privacy Commissioner argued should have been extended to the OPC as well). The constitutional path for a parliamentary officer to hold such powers is established, despite the OPC's continued omission from that model, and the argument that it cannot be replicated for the *Privacy Act*'s ombudsman function does not reflect existing law.

Another objection might be that the CAP mechanism is sufficient because the prospect of Federal Court referral provides adequate deterrence. It does not. The deterrent effect of a mechanism requiring a finding, a remediation plan, and further proceedings before any legal consequence attaches deters less than one that can impose a binding order directly. The provincial evidence confirms this: institutions subject to direct order-making power respond to investigations and recommendations differently from institutions subject only to recommendations. The change in institutional behaviour occurs at the investigation stage, not at the order-issuance stage, because the prospect of a binding order changes how institutions weigh the costs of non-compliance in a way that the CAP mechanism alone cannot replicate.

> **Recommendation:** Adopt the CAP mechanism; add direct binding order-making power for urgent and systemic violations, without an intermediate CAP requirement; and introduce financial penalties calibrated to institutional size and violation severity, consistent with the enforcement model the CPPA will apply to the private sector.

Proposal 21: Re-Identification Offences

The proposal to create offences for intentional re-identification of de-identified data should be adopted. The re-identification research makes clear that the privacy value of de-identification depends entirely on legal and technical backstops against re-identification: de-identified data retains identifiability risk, and without sanctions against deliberate re-identification, the flexibility afforded by the de-identification regime in Proposal 7 creates a risk that the protections are nominal. Bill 64 in Quebec, the UK Data Protection Act, and privacy statutes in Australia, South Korea, and Japan all include re-identification prohibitions.[29] Canada's federal public-sector statute should do the same.

The proposal's exceptions for investigations and national security are legitimate but must be narrowly drafted and subject to judicial authorization. An unqualified exception for law enforcement risks swallowing the rule: identification of individuals from de-identified datasets is the objective that motivates many law enforcement applications of personal data. The exception should require prior judicial authorization except in exigent circumstances, and all uses should be subject to oversight reporting requirements to the Commissioner.

The Act should also require organizations that de-identify data for the purpose of invoking the flexibility measures under Proposal 7 to implement technical safeguards designed to

[29]Bill 64 (Quebec), s 90; UK Data Protection Act 2018, s 171; Privacy Act 1988 (Cth) (Australia), s 13G; Act on the Protection of Personal Information (Japan), as amended 2023.

prevent third parties from re-identifying it. A re-identification offence without a corresponding obligation to make re-identification technically difficult places the entire deterrent burden on the sanctions regime and ignores the prevention mechanism that technical safeguards provide. Data protection authorities internationally recommend pairing the two: technical safeguards as a pre-condition for de-identification flexibility, and legal sanctions for circumventing those safeguards.

> **Recommendation:** Adopt re-identification offences; require prior judicial authorization for law enforcement and national security exceptions; and condition de-identification flexibility on implementation of technical safeguards against re-identification.

Proposal 22: Expanded Judicial Remedies

The proposal to expand Federal Court authority to cover privacy violations beyond the access-denial context should be adopted. Section 41's current restriction to access denials means that violations of the fair information obligations in sections 4 through 9, including unlawful collection, unauthorized use, and improper disclosure, carry no route to individual compensation outside the Commissioner's investigative process. For AI-era harms, including erroneous benefit denials, misidentification by facial recognition systems, and discriminatory credit or immigration outcomes, this is a gap incompatible with the rights-based framework the Act is intended to provide.

The minimum reform should grant individuals the right to seek compensation in Federal Court for violations of sections 4 through 9 following an OPC finding, consistent with the *Consumer Privacy Protection Act*'s sequential model. This is the legislative floor, not the ceiling.

A more adequate reform would allow direct judicial action, without requiring prior OPC proceedings, for AI decisions affecting *Charter*-protected interests: equality under section 15, liberty and security under section 7. The GDPR model provides an independent right to judicial remedy without conditioning judicial access on prior regulatory proceedings.[30] The Supreme Court's decision in *Douez v Facebook*, holding that a contractual forum selection clause was unenforceable against statutory privacy rights, signals judicial willingness to protect access to domestic courts for privacy claims against powerful actors. That foundation could support explicit class proceeding recognition for AI-era distributional harms.[31]

Class proceedings are essential for AI distributional harms and should be recognized. Where an AI system produces disparate error rates across groups defined by protected characteristics, the harm is collective in character: it cannot be adequately addressed by individual complaint procedures, and the evidence necessary to establish the pattern is available only across a population of cases. A class action vehicle allows individual and group harms to be adjudicated together, provides the aggregation that distributional injuries require, and creates reputational and financial incentives for institutions to investigate and correct

[30]GDPR, Arts 79–82 (independent right to judicial remedy and compensation without prior regulatory proceedings).

[31]*Douez v Facebook, Inc*, 2017 SCC 33.

systematic bias that individual complaint procedures cannot generate. The Act should provide that class proceedings may be certified for violations arising from AI systems operating in systematic or distributional patterns. Alternatively, amendments to the *Federal Courts Act* should confirm that privacy class actions in the public-sector context may be certified, a question that current law leaves uncertain in a way that favours powerful institutions over the individuals and communities they harm.

> **Recommendation:** Extend Federal Court access to all violations of sections 4–9 following an OPC finding as the minimum reform; allow direct judicial action without prior OPC proceedings for AI decisions affecting Charter-protected interests; and recognize class proceedings for distributional AI harms.

Conclusion

The Act's central reform task is to address a power asymmetry that AI generates between data-rich government institutions and the citizens they govern: institutions hold inferential profiles of individuals, make consequential decisions through systems whose logic cannot readily be examined or challenged, and currently face no meaningful accountability for the accuracy or fairness of those decisions. Individuals have neither the information nor the tools to detect harms that may not be occurring to any identifiable individual but are occurring across populations. The amendments proposed here aim to operationalize what each proposal's own rationale requires under current (and future) technology.

The three sets of reforms addressed in this submission are mutually reinforcing. Transparency without contestability provides information but no recourse. Contestability without systemic obligations corrects individual errors while leaving systemic issues such as bias unaddressed. Enforcement without transparency lacks the evidentiary foundation for action. None of these strengthening amendments alter the direction of the proposals. Enacted together and strengthened as this submission proposes, they would give Canada a public-sector privacy statute capable of constraining the power asymmetry that AI creates. The *Privacy Act*'s quasi-constitutional status demands no less.

Executive Summary

This submission supports Proposals 3, 5, 6, 7, 10, 12, 13, 14, 19, 21, and 22, subject in each case to specific strengthening amendments identified below. The amendments close gaps between each proposal's stated rationale and its proposed mechanism, particularly where the mechanism reflects a collection-centred drafting model that technology has transcended.

Proposal 3 (Privacy Impact Assessments). The PIA trigger should cover any AI system processing personal information at scale, not only programs making individual decisions. PIAs should mandate demographic accuracy validation, third-party data supplier verification, and publication of contractor-held AI documentation.

Proposal 5 (ADS Transparency). The explanation right should attach to the inference step that generates a consequential recommendation, not only to the final decision. The human

review right should require assessment of systemic fairness and include referral obligations where patterns of group-level error are detected.

Proposal 6 (Notice Requirements). Pre-processing notice should be mandatory for decisions affecting legal rights and benefits; post-processing notice should be limited to impracticable contexts. ADS notices should specify data categories used, automation level, and the individual's rights to explanation and human review.

Proposal 7 (Sensitivity and Identifiability Spectrum). The identifiability spectrum should be adopted in full. The sensitivity categories should be supplemented with an inference-based trigger: data processed in a way reasonably likely to generate sensitive-category outputs should be subject to the highest safeguards regardless of input sensitivity.

Proposal 10 (Necessity Test). The necessity test should apply to inference generation, not only to collection from external sources. The effectiveness criterion should include pre-deployment demographic accuracy validation.

Proposal 12 (Privacy as Fundamental Right). The purpose clause should also recognize systemic institutional accountability as a co-equal objective alongside individual rights. The service delivery objective should be explicitly subordinate to the fundamental right.

Proposal 13 (Principles). The OPC should be empowered to issue binding guidance specifying how each principle applies in the AI context. Each principle should receive operative expression as an enforceable obligation in the Act's body, following the *Consumer Privacy Protection Act* model, rather than appearing only as a prefatory statement.

Proposal 14 (Key Definitions). The definition of 'creation of personal data' should be adopted and aligned with the necessity test in Proposal 10. The 'sensitive data' definition should be supplemented with an inference-based trigger consistent with the amendment to Proposal 7.

Proposal 19 (Order-Making Powers). The Commissioner should receive direct binding order-making power for urgent and systemic violations, without an intermediate corrective action plan requirement. Financial penalties calibrated to institutional size and violation severity are essential and should be introduced.

Proposal 21 (Re-Identification Offences). Exceptions for law enforcement and national security should require prior judicial authorization. De-identification flexibility should be conditioned on implementing technical safeguards against re-identification.

Proposal 22 (Judicial Remedies). Federal Court access should be extended to all violations of sections 4–9 following an OPC finding. Class proceedings should be recognized for AI systems generating distributed harms.