AI Decision Governance Framework: Quebec's Law 25

Methodology | *Anastasia Bruce*

Purpose

Quebec's Law 25 is currently the strictest privacy law in North America, with penalties reaching twenty five million Canadian dollars or four percent of worldwide turnover. Section 12.1 of the law creates specific, binding obligations for any organisation that uses AI to make decisions about individuals. This framework was built to translate those statutory obligations into governance checkpoints that a regulated organisation can actually implement, rather than leaving compliance as an abstract legal reading.

Step 1: Establish the Statutory Basis

The framework was built by working directly from the text of Section 12.1 and the wider obligations of Law 25, including the requirement for Privacy Impact Assessments and a designated Privacy Officer. Working from the statutory text itself, rather than from secondary commentary, was a deliberate choice, since guidance from Quebec's Commission d'access a l'information on automated decision-making specifically remains limited at this stage.

Step 2: Determine Applicability

Section 12.1 applies only to decisions rendered exclusively through automated processing. If a human meaningfully participates in the decision, the requirement does not apply. The first checkpoint in the framework is therefore an applicability test. Each AI-driven decision an organisation makes is assessed against this threshold before any further obligation is considered, since applying the full requirement to a decision that does not legally require it would create unnecessary operational overhead.

Step 3: Map Each Statutory Requirement to a Control

Where a decision does meet the automated processing threshold, the law sets out four specific requirements. The individual must be informed that automated processing was used. They must be able to request an explanation of the factors and logic behind the decision. They must be able to access the personal data used to reach it. They must be able to submit observations to a staff member capable of reviewing the decision. Each of these four requirements was mapped to a concrete organisational control: a disclosure trigger embedded in the decision workflow, a documented explainability process, a data access procedure, and a designated human review pathway.

Step 4: Define Evidence and Review Cadence

For each control, the framework specifies the evidence an organisation would need to produce to demonstrate compliance during an audit or regulatory inquiry, along with a recommended review cadence appropriate for a regulated environment. This mirrors the checkpoint based structure used in the earlier Ghana Data Protection Act framework, adapted here to a jurisdiction with binding, enforceable automated decision-making requirements.

Approach to Ambiguity

Because regulatory guidance on automated decision-making under Law 25 remains limited, the framework was built to be conservative. Where the statutory language could reasonably support more than one interpretation, the framework treats the stricter reading as the operating assumption. This is intended to reduce compliance risk for any organisation adopting the framework, rather than optimising for the narrowest possible reading of the law.

Note: this methodology describes the framework's construction. The client and specific implementation details have been redacted to preserve confidentiality, while the underlying method remains fully transparent.