

# Canon Admissibility Failure Memo (CAF M v2)

Structural assessment of an AI-assisted consumer credit decision pipeline (UK).

**Method:** Canon (non-interpretive structural admissibility diagnostic)

**Scope:** UK consumer lending decisioning (regulatory anchors provided for provenance only)

**Acronym note:** CAF M here means **Canon Admissibility Failure Memo** (not Computer-Aided Facility Management).

## Preface (read first)

This memo is a **non-interpretive structural diagnostic**. It does not assess legality, adequacy, fairness, correctness, or outcomes. It operates on publicly stated constraints and commonly deployed pipeline structures. Where an undeclared hand-off is described, it is a structural dependency (reliance on silent reconstruction), not an allegation of wrongdoing.

## Regulatory anchors (provenance)

- FCA Handbook — CONC 5.2A (Creditworthiness assessment):  
<https://handbook.fca.org.uk/handbook/CONC/5/2A.html>
- FCA Policy Statement PS18/19 (Assessing creditworthiness in consumer credit):  
<https://www.fca.org.uk/publication/policy/ps18-19.pdf>
- FCA Handbook — PRIN 2A.4 (Consumer Duty: price and value outcome):  
<https://handbook.fca.org.uk/handbook/PRIN/2A/4.html>
- ICO guidance — Automated decision-making and profiling: <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/individual-rights/automated-decision-making-and-profiling/>
- Legislation.gov.uk — GDPR Article 22: <https://www.legislation.gov.uk/eur/2016/679/article/22>

## 1. Decision pipeline under consideration (concrete)

- 1 Consumer submits a credit application (often including credit bureau data; sometimes Open Banking transaction data).
- 2 Credit risk model executes (score / classification used for approve, decline, or pricing).
- 3 Explanation artefact generated (reason codes / rationale used in governance and customer communications).
- 4 Human oversight varies (threshold automation; manual review for borderline cases; compliance review/audit).
- 5 Decision issued (approve/decline/price), carrying significant effects for the customer.

## 2. Claimed safeguards and authority (accepted as stated)

Pipelines of this type commonly rely on claims such as: responsible lending / creditworthiness assessment, explainability, human-in-the-loop oversight, governance controls, and (where relevant) automated decision-making safeguards. These claims are not evaluated here; they are treated as the declared basis for authority.

## 3. Canon structural analysis

### 3.1 Load-bearing abstraction

The pipeline relies on the abstraction that **explanation + human oversight** is sufficient to render an automated output accountable as a credit decision.

### 3.2 Undeclared hand-off

At the point of review, interpretive responsibility is transferred to humans under uncertainty and partial information. What is frequently not explicitly declared includes: what level of understanding is required, how uncertainty should be handled, how conflicts are arbitrated, and where responsibility sits when an explanation is technically correct but materially incomplete for accountability.

### 3.3 Structural consequence

Responsibility is operationally real but structurally underrepresented. Accountability is therefore reconstructed after the fact during complaints, audits, or supervisory scrutiny, rather than being explicitly represented in the pipeline at the point authority is exercised.

## 4. Admissibility result

Under Canon, this pipeline is **structurally incomplete for the authority it claims** where accountability depends on undeclared reconstruction. This does not imply non-compliance, harm, or fault. It records a missing representation requirement.

## 5. The non-optional question

**If this pipeline is deployed, who explicitly owns responsibility for interpreting the model output under uncertainty, and where is that responsibility declared?**

### Addendum: why this matters commercially (non-promissory)

- **Consumer Duty context:** outcome evidencing increases the cost of silent reconstruction and post-hoc narrative assembly.
- **Upstream control:** Canon operates before downstream GRC artefacts by making the boundary between stated evidence and assumptions explicit.
- **Remediation cost:** incomplete files drive manual review and reconstruction; Canon makes absence visible early and consistently.
- **Integration:** deploy as a deterministic validation service (API-first) that returns stated/not stated plus a machine-readable audit log.