

Accountability and Fair Lending Evidence in Outsourced Credit Agents

A Bank Decision on Explainable Credit Delivery

Phil Hatch | Akholi

October 23, 2026

Bank leaders should approve outsourced credit assistance only when the bank can reconstruct each consequential decision, test relevant outcomes, and provide the explanations and review required by applicable law.

Require Evidence of the Credit Decision

The bank's head of lending and chief risk officer should require a decision evidence pack before approving an outsourced credit agent. The pack must connect information received, policy applied, recommendation produced, human judgment, final action, and customer communication. A provider's model score does not establish that the bank's complete credit process is lawful or explainable.

In US lending, Regulation B requires specific principal reasons for adverse action in the circumstances it governs. Its official interpretations require those reasons to reflect factors considered or scored. A fluent explanation generated after the event is insufficient if it describes a reason the decision did not use.¹

The executive approval should define the permitted assistance and retained bank decisions. Document extraction, preparation, recommendation, pricing, approval, and notice generation

¹ Consumer Financial Protection Bureau, "Regulation B," Consumer Financial Protection Bureau, current version accessed October 23, 2026, 12 CFR 1002.9(b)(2) and official interpretations 2–7, <https://www.consumerfinance.gov/rules-policy/regulations/1002/9/>.

carry different consequences. Management should not describe all of them as decision support when the runtime can change the customer's outcome without a qualified bank decision-maker.

Trace Inputs and Overrides Through the Two Firms

A hypothetical agent could summarize an applicant's income, retrieve policy, recommend a decline, and draft a notice. An error in the source document, policy version, or summary can affect the result even if the underlying score operates as designed. The bank should preserve the evidence necessary to identify where the consequential change occurred.

The study's manual-alteration accounts support scrutiny of human corrections between input and delivery. They do not establish unfair lending or defective credit decisions. Their relevance is that the bank must know when a provider changed information, process, or output before relying on the resulting recommendation.²

Require a recorded reason and authority for overrides. A human reviewer should have enough information, time, and standing to challenge the agent. An approval click does not establish substantive review. Management should test whether reviewers can reject, correct, and escalate a recommendation without being forced back into the same automated outcome.

Build a Credit Decision Evidence Pack

Figure 1 proposes four connected records. The pack should retain the relevant input and provenance, policy and system versions, decision path and overrides, and the notice or explanation delivered. Access should be controlled, with retention set by applicable duties rather than a general desire to keep every prompt indefinitely.

Outcome monitoring should examine relevant groups, product segments, exceptions, and changes in the service. Compliance specialists should define lawful testing methods and appropriate comparisons. A difference between groups is a signal to investigate; it does not by itself establish illegal discrimination or explain the mechanism producing it.

The bank should be able to rerun selected cases against the recorded configuration or otherwise reconstruct the decision faithfully. If an upstream model cannot be reproduced exactly, management should state that limitation and preserve the evidence needed for a reliable explanation. It should not imply perfect reproducibility from a version label alone.

² Phil Hatch, *Agentic AI Outsourcing at Tier 1 Financial Institutions: Full Study Report*, version 1.8.2 (Akholi, forthcoming manuscript dated November 1, 2026), Part 5, "Five Manual Alteration Methods" and "Disclosure of Manual Alteration," table 5.8, <https://doi.org/10.2139/ssrn.7164559>.

Figure 1. Connect the recommendation to the customer outcome

Record	Required connection	Executive assurance
Input and provenance	Information actually used	Correct case and source
Policy and configuration	Rules and versions applied	Approved decision basis
Judgment and override	Recommendation to final action	Effective bank challenge
Outcome and explanation	Actual factors to customer notice	Faithful explanation and monitoring

Author's proposed operating framework. Adapt to the service and applicable jurisdiction; not an empirically validated standard.

Source: Phil Hatch, proposed framework developed for this paper.

Keep EU UK and Swiss Conclusions Distinct

Swiss FADP article 21 addresses decisions based exclusively on automated processing that have legal consequences or considerable adverse effects, including information and requested human review, subject to stated exceptions. The bank should assess whether the actual workflow meets that definition rather than infer the answer from the presence of a human somewhere in delivery.³

EU and UK lending require separate analysis of credit, equality, privacy, and relevant AI requirements. The UK's legislative changes should not be treated as identical to EU automated-decision rules. Counsel should identify the effective provisions for each product and entity, including the actual role of any human decision-maker.

FINMA's AI guidance includes observations on bias, explainability, documentation, and independent review. Those observations help frame Swiss assurance questions but do not

³ Federal Data Protection and Information Commissioner, "Duty to Provide Information," Federal Data Protection and Information Commissioner, accessed October 23, 2026, discussion of automated individual decisions under article 21 FADP, <https://www.edoeb.admin.ch/en/duty-to-provide-information>.

prescribe a universal fair-lending test. The common operational objective is a decision path that the responsible bank can explain and challenge under its own obligations.⁴

Make Provider Change a Credit Risk Decision

A new model, extraction method, reference source, policy retrieval rule, or notice template can change credit outcomes. The provider should identify affected decisions and supply comparison evidence before material change. The bank should retain approval over changes that alter the permitted authority or the basis of customer treatment.

An explanation test should begin with an actual recorded decision and compare the delivered reason with the factors used at that time. A generated reason that sounds plausible but cannot be tied to the decision should fail the test. Where a human override changes the outcome, the record should show the override's actual basis. This protects the bank against using an articulate after-the-event narrative as a substitute for evidence of its credit judgment.

Mayer Brown's agentic governance analysis identifies biased or unfair decisions and consequential action boundaries as assessment issues. The bank should translate those issues into credit-specific acceptance criteria rather than rely on a generic AI impact questionnaire. The criteria should include decision accuracy, explanation fidelity, effective challenge, and exception handling.⁵

The release test should include incomplete applications, contradictory documents, unavailable information, and requests requiring escalation. The system should preserve uncertainty instead of inventing facts or defaulting silently to a consequential outcome. Failures should be assigned to owners and retested before wider use.

Approve Only the Authority the Evidence Supports

The board should receive a concise statement of scope, retained bank authority, material outcome findings, unresolved explanations, and remediation. Management should restrict the affected use when it cannot reconstruct a consequential decision or meet an applicable explanation duty. A provider indemnity cannot supply missing decision evidence.

⁴ Swiss Financial Market Supervisory Authority, *Governance and Risk Management When Using Artificial Intelligence* (Swiss Financial Market Supervisory Authority, December 18, 2024), FINMA Guidance 08/2024, secs. 2.1–2.7, https://www.finma.ch/~media/finma/dokumente/dokumentencenter/myfinma/4dokumentation/finma-aufsichtsmittelungen/20241218-finma-aufsichtsmittelung-08-2024.pdf?hash=60E2745D7AC1D7F4E06CE43474539F68&sc_lang=en.

⁵ Arsen Kourinian, "Governance of Agentic Artificial Intelligence Systems," Mayer Brown, February 24, 2026, "Mitigation Measures," <https://www.mayerbrown.com/en/insights/publications/2026/02/governance-of-agentic-artificial-intelligence-systems>.

The approval question is whether the bank can remain accountable for the outsourced credit process in practice. An evidence pack linking the decision to the information, policy, interventions, and communication gives executives a defensible basis for bounded use and for refusing an expansion that the current controls cannot support.

Evidence and Source Note

The study is a nonrepresentative cross-section of 850 vendor delivery-team respondents on 150 engagements serving Tier 1 financial institutions, collected during January–June 2026. Neither banks nor vendors corroborated the accounts. It establishes reported experience, not causation, legal breach, or market prevalence. The author-supplied version 1.8.2 is dated November 1, 2026 and is cited as a forthcoming manuscript consulted October 23, 2026.⁶

This is an analytical decision paper based on the supplied study and selected primary public sources checked through October 23, 2026. The original contribution is a two-firm credit decision evidence pack separating input provenance, configuration, judgment, and customer outcome. The proposed method has not been independently validated. Legal requirements depend on the entity, activity, jurisdiction, and effective provisions; supervisory observations, research, and professional commentary have distinct authority.

Bibliography

Consumer Financial Protection Bureau. “Regulation B.” Consumer Financial Protection Bureau. Current version accessed October 23, 2026. <https://www.consumerfinance.gov/rules-policy/regulations/1002/9/>.

Federal Data Protection and Information Commissioner. “Duty to Provide Information.” Federal Data Protection and Information Commissioner. Accessed October 23, 2026. <https://www.edoeb.admin.ch/en/duty-to-provide-information>.

Hatch, Phil. *Agentic AI Outsourcing at Tier 1 Financial Institutions: Full Study Report*. Version 1.8.2. Akholi. Forthcoming manuscript dated November 1, 2026. Author-supplied manuscript consulted October 23, 2026. <https://doi.org/10.2139/ssrn.7164559>.

Kourinian, Arsen. “Governance of Agentic Artificial Intelligence Systems.” Mayer Brown. February 24, 2026. <https://www.mayerbrown.com/en/insights/publications/2026/02/governance-of-agentic-artificial-intelligence-systems>.

Swiss Financial Market Supervisory Authority. *Governance and Risk Management When Using Artificial Intelligence*. Swiss Financial Market Supervisory Authority, December 18, 2024. https://www.finma.ch/~media/finma/dokumente/dokumentencenter/myfinma/4dokumentation/finma-aufsichtsmittelungen/20241218-finma-aufsichtsmittelung-08-2024.pdf?hash=60E2745D7AC1D7F4E06CE43474539F68&sc_lang=en.

⁶ Hatch, *Agentic AI Outsourcing*, Introduction, “Executive Summary” and “Evidentiary Basis and Its Limits.”