


STRATEGY / OCTOBER 2026

Banking on AI with accountability

Scale the decision, the control and the economic case together



THE DECISION IN BRIEF

Banks should fund AI assistance and delegate decision authority as separate choices, each supported by its own customer, control and economic case.

- 01

A policy assistant, a customer explanation and a credit action have different consequences. Manage them as separate decision portfolios, even when they share technology.
- 02

Scale value after review, rework, integration and recovery costs. Useful assistance can remain permanently assistive; autonomy needs an additional case.
- 03

Connect every consequential decision to authoritative evidence, a responsible owner and a tested fallback. Apply regulatory requirements within their actual scope rather than treating one model framework as universal coverage.

THE MAYDEN VIEW


Fund assistance for the value it creates. Grant decision authority only on its own evidence.

INSIDE THIS PERSPECTIVE

The argument at a glance	3
01 / Separate value from authority	4
Exhibit 01 / The assistant prepares evidence; the bank owns the decision	5
02 / Connect controls to the decision	6-7
Exhibit 02 / Adoption is established in the survey; value is a separate question	8
03 / Price the complete operating role	9
04 / Earn the next scale decision	10
The leadership agenda	11
Method and references	12

ONE DECISION, FOUR CONNECTED QUESTIONS

Carry accountability with the decision



ACCOUNTABLE LEDGER / CONCEPTUAL ARTWORK

01 Separate value from authority

Separate useful assistance from decision authority

02 Connect controls to the decision

Apply model guidance within its actual boundary

03 Price the complete operating role

Test resilience across shared suppliers

04 Earn the next scale decision

Prepare the credit file without inventing the decision

Fund assistance for the value it creates. Grant decision authority only on its own evidence.

CHAPTER 01

01

Separate value from authority

Separate useful assistance from decision authority

A bank can save time assembling a credit file without giving AI authority to approve credit. It can improve a product explanation without permitting the system to invent an exception. These distinctions should drive the AI investment portfolio. The 2024 Bank of England and FCA survey found 75% of 118 responding UK financial-services firms used AI and another 10% planned adoption within three years. ^[1] That adoption evidence establishes neither bank profitability nor the case for autonomous decisions.

The more useful question is which decisions the bank can improve. An employee searching policy documents, an adviser explaining a product and a system changing a credit limit have different consequences. They may use related technology, yet the economic opportunity and accountability remain distinct. Treating them as one generative-AI programme encourages a single success metric, usually usage, and a single approval process, usually too heavy for some uses and too weak for others.

For each entry in the portfolio, record the decision owner, affected customer or asset, permitted authority, evidence threshold and value mechanism. mayden proposes this as the operating unit for investment. The shared platform supplies capability; it cannot settle the customer consequence or accountability on the bank's behalf.

Give three operating roles three distinct cases

Employee assistance can reduce the effort of finding, drafting and comparing information. Its initial value case should focus on the complete task, including review and correction. A faster first draft may still create more work if staff must verify every statement against inaccessible sources. Start with a bounded information domain, clear provenance and an employee who retains responsibility for the final action. Expanding access to sensitive records should be an explicit change, not a convenience added after launch.

Customer-facing interpretation has a different burden. An answer about fees, eligibility or contractual obligations can influence a customer's choice even when the system cannot execute a transaction. The bank needs authoritative content, controlled updates, tested language and a reliable escalation route. An apparently helpful answer that invents an exception to policy creates a commercial and conduct problem. Evaluate whether customers understand the answer and can reach a person when their situation does not fit the standard path.

Consequential execution requires a third case. Decisions affecting credit, payments, financial crime operations or customer access need bounded authority and a clear recovery design. US Regulation B, for example, requires specific principal reasons for covered adverse actions, supplied through the applicable notice or request procedure. ^[2] The broader management lesson is to design the decision and its explanation together. A generated narrative should not be allowed to invent reasons after the decision. Requirements in other jurisdictions must be mapped separately; the US example is not a global rule.

ILLUSTRATIVE DECISION TOOL

The assistant prepares evidence; the bank owns the decision

Illustrative credit-service authority chain

Stage	Permitted authority in the fictional case	Evidence and test
Extract and assemble	Assistant prepares a case; no credit action	Match identity, source version and disputed/missing facts
Recommend or decide	Only the authorized bank process decides within its scope	Actual factors, responsible authority and outcome
Communicate reasons	Draft from recorded decision basis; authorized release	Specific principal reasons for covered US actions, not generic policy prose
Correct or review	Established decision and remedy owners handle disputes	Preserve decision, correction and observed downstream change

THE MANAGEMENT IMPLICATION

Useful assistance can scale without giving the assistant power to decide credit or invent its reasons.

No applicant data, underwriting threshold or measured bank outcome. US Regulation B is a scoped example; other jurisdictions require separate analysis. Sources: [2], [6], [8]

CHAPTER 02

02

Connect controls to the decision

Apply model guidance within its actual boundary

US supervisors issued SR 26-2 in April 2026, replacing SR 11-7 and emphasizing a tailored approach to model risk. ^[3] The attached guidance is most relevant to institutions above \$30 billion in assets, with exceptions for smaller institutions carrying significant model risk. It sets no enforceable standards; violations of law or unsafe or unsound practices can still draw supervisory action. ^[8]

The scope boundary is explicit: footnote 3 excludes generative and agentic AI models, while covering traditional statistical and quantitative models and non-generative, non-agentic AI models. ^[8] A bank may adapt useful validation principles for a language assistant, but that is an internal governance choice, not evidence that SR 26-2 directly governs that assistant. Software, privacy, security, conduct and operational controls still need clear owners. Existing model oversight should connect to those functions where

appropriate, without pretending one framework covers every risk.

For the illustrative credit journey, mayden would allocate controls across three uses. For a conventional credit model, the credit-risk owner defines the decision policy and an independent validation function challenges performance, limitations and change evidence under the applicable model framework. For a generative adviser, credit operations owns source fidelity and the check that a communicated reason matches the actual decision; security owns access boundaries. For an action-taking agent, the transaction owner defines permitted actions, approval limits and reconciliation, while engineering tests execution and recovery. These are proposed internal allocations, not an extension of SR 26-2 to excluded models. Privacy, conduct and customer-remedy owners address the shared obligations. The business executive remains accountable for the complete workflow; independent challenge must be able to require a change when a handover leaves a control unowned.

CHAPTER 02

02

Connect controls to the decision

Repair the evidence chain before expanding discretion

For each banking use case, connect the data control to the decision it protects. A stale policy source can misdirect an adviser; an incorrect customer identifier can attach evidence to the wrong account; a supplier outage can interrupt an essential workflow. mayden proposes presenting those paths on one control map, with a business owner and fallback for each. The useful question is whether the bank can detect and contain the specific failure before it becomes a customer outcome.

The ECB's 2024 guide on effective risk data aggregation and reporting addresses governance, architecture, data quality and timely information in its supervisory context. ^[4] It is not a generative-AI deployment manual. It nevertheless provides a relevant reminder: the ability to assemble trustworthy information is an institutional capability, not a feature of the final analytical tool. mayden's inference is to invest in authoritative retrieval and data contracts

alongside models, with evidence that the selected sources are appropriate for the task.

A March 2026 FSI study examines AI data risks in financial services, with particular attention to generative AI. It identifies privacy, data quality and security challenges compounded by third-party dependencies and provider concentration. ^[9] The paper is supervisory analysis by its authors, not a new binding rule. It reinforces the need to assess the complete data chain rather than rely on a model validation label alone.

This creates a practical sequencing choice. Where records are reliable and permissions clear, assistance can scale quickly within a narrow domain. Where the bank cannot establish which record is authoritative, increasing autonomous action should wait. Use the early programme to identify and resolve those bottlenecks rather than conceal them behind a fluent interface. The economic benefit of repairing a shared data dependency may extend across several decisions; record that benefit separately to avoid claiming the same saving repeatedly in multiple AI business cases.

EMPIRICAL RESULT · RESPONDENT SURVEY

Adoption is established in the survey; value is a separate question

% of 118 survey respondents



- Already using AI
- Planning adoption
- Neither category

THE MANAGEMENT IMPLICATION

A descriptive adoption measure cannot establish profitability or deployment quality.

Bank of England/FCA 2024 survey; UK financial services, not banks only. 15% is calculated as 100-75-10. Planned adoption refers to the next three years from 2024. Sources: [1]

CHAPTER 03

03

Price the complete operating role

Test resilience across shared suppliers

A bank can reduce development time by buying a model service, but it also inherits a dependency on service availability, version changes, capacity and commercial terms. The FSB's 2024 assessment identifies concentration, correlated behavior, cyber risk and model/data concerns as potential financial-stability vulnerabilities. ^[5] These are channels of risk, not a forecast that AI will cause a particular crisis. At institution level, the immediate task is to identify which important operations rely on the same suppliers and assumptions.

The Basel Committee describes its operational-resilience principles in terms of continuity, testing, interdependencies and third-party management. ^[6] Apply that lens to AI-supported services: what continues if the model service is unavailable, degraded or withdrawn? A second vendor is not a complete fallback if it relies on the same cloud region, cannot use the same data lawfully or produces outputs that the bank has never validated. Redundancy must be demonstrated at the workflow level, including staff capacity to handle exceptions.

For Saudi-regulated institutions, SAMA's published Cyber Security Framework provides a jurisdiction-specific reference for safeguarding the financial sector. ^[7] It should not be generalized into an AI rule for every Gulf bank. Our portfolio proposal adds a dependency map that connects each decision to its model, hosting, data and recovery owner. That map helps the CEO see concentrations that individual project business cases can miss. The bank should negotiate change notice, incident cooperation and exit support in terms its operating teams can implement and test.

Measure value after review and recovery

The net-value equation is straightforward to state and difficult to measure: attributable benefit less technology, integration, review, exceptions, control and recovery costs. Time saved is not automatically a cash saving. It may become additional service capacity, faster response, better quality or reduced overtime. Each outcome can matter, but the finance team should specify the conversion mechanism. Avoid multiplying a survey estimate of minutes saved by every employee's salary and presenting the result as realized profit.

A useful evaluation compares complete workflows. For a customer-service assistant, measure resolution quality, repeat contacts, escalation and employee effort alongside handling time. For document analysis, examine corrections and downstream decisions, not just pages processed. For a credit-related use, evaluate performance over relevant customer segments and changing conditions, subject to the applicable law and validation requirements. A pilot should include difficult cases and a baseline process so that management can distinguish model assistance from other improvements introduced at the same time.

Fund a bounded trial for task quality and cost, then integration for the complete workflow. Expand authority only when its control and recovery case is demonstrated. A permanently assistive service can be the best investment. Value and autonomy are separate decisions, and the evidence may support one without the other.

CHAPTER 04

04

Earn the next scale decision

Prepare the credit file without inventing the decision

Consider a fictional bank using an assistant to organize evidence for a credit application. It may extract a document date and propose a case summary; it cannot choose an applicant outcome or change a credit limit. The authorized decision process separately records the factors it actually considered and the responsible authority. This separation lets the bank test useful clerical assistance without treating a good summary as evidence that autonomous credit decisions are justified.

Suppose a principal factor in a covered US adverse action was an existing delinquent obligation recorded in the decision file. A language assistant must not replace that factor with a generic statement about internal policy or invent a more persuasive explanation. Regulation B requires specific principal reasons, and its official interpretation links them to factors actually considered or scored. ^[2] The bank should compare proposed communications with the recorded decision basis, using the applicable notice or request procedure. This fictional example is not a global legal rule or a decision about a real applicant.

The test should include a mismatched identifier, disputed source information and a later correction. A reviewer confirms that evidence belongs to the correct application and routes disputes through the established process. The reasons record must remain connected to the authorized decision even when communication is revised. Missing evidence is a reason to investigate or use the appropriate process; it is not permission for a generative system to supply an assumed factor. Exhibit 01 records what each component may do and which evidence it must preserve.

Compare the assisted administrative workflow with the current process on the same cases. Include review, correction, repeat contact and escalation. Keep credit outcomes separate: losses and fairness emerge over a different period and depend on population and conditions. Lower preparation effort cannot establish those effects. Decide whether bounded assistance earns continued investment before considering any expansion of authority.

Scale the operating role that earns the investment

A decision portfolio should reduce ambiguity rather than add a second bureaucracy. Reuse platform, security and supplier evidence where it applies. Keep separate the customer consequence, permitted authority and owner of each use. A shared model can support several services without making their release decisions interchangeable.

Review one employee-assistance use, one customer-facing interpretation and one consequential workflow together. Establish the benefit owner, risk owner, baseline and shared dependencies. Test the fallback and ask an independent reviewer for the strongest reason to withhold scale. Retain human authority when that preserves the benefit at lower exposure.

The board's account should show which decisions improved, the benefit observed, the complete cost and the authority granted. Include failed or discontinued trials. Reconcile shared costs once across the portfolio. The next commitment should follow the evidence for a particular operating role - not the assumption that every successful assistant should become an autonomous decision maker.

FROM INSIGHT TO ACTION

Decisions to take forward

BUSINESS AND CREDIT-PROCESS OWNERS

Approve assistance separately from power to make consequential decisions.

Evidence to require: Case-level authority and actual-reasons record

INDEPENDENT RISK AND COMPLIANCE REVIEWERS

Demonstrate the relevant duty and control for the actual activity.

Evidence to require: Identity, disputed-source, reasons and correction cases

FINANCE AND OPERATIONS

Scale on complete-work value; keep immature credit outcomes and expected benefits separate.

Evidence to require: Review, correction, repeat-contact and exception account

mayden connects senior judgement, research and engineering to turn strategic choices into working capabilities.

EVIDENCE YOU CAN INSPECT

Research notes

Method and boundaries

Structured synthesis of nine opened primary regulator, supervisory and institutional sources; cutoff 10 October 2026. The adoption chart reproduces a descriptive 2024 survey with one transparent residual calculation.

Sources cover adoption, credit explanation, current US model guidance, European data governance, global resilience and stability, and a Saudi cybersecurity reference.

No bank-level financial analysis, portfolio backtest, client data or independent operating audit. Jurisdiction-specific guidance is not treated as globally binding. Proposed decision tracks and economics require local measurement.

Primary sources and author research

[1] Artificial intelligence in UK financial services 2024

Bank of England and FCA / 2024-11-21. Sections 1.2 and 2.1; 118 responses, 75% using, 10% planning.

[Read the source](#)

[2] Regulation B, §1002.9 Notifications

CFPB / Current text accessed 2026-10-10. Section 1002.9(b)(2); official interpretation 9(b)(2)-2 through -7.

[Read the source](#)

[3] SR 26-2: Revised Guidance on Model Risk Management

Federal Reserve/OCC/FDIC / 2026-04-17. Applicability and supersession; attached guidance sections II-III.

[Read the source](#)

[4] Guide on effective risk data aggregation and risk reporting

ECB Banking Supervision / 2024-05. Governance, architecture and data-quality sections.

[Read the source](#)

[5] FSB assesses the financial stability implications of artificial intelligence

Financial Stability Board / 2024-11-14. Principal vulnerability channels.

[Read the source](#)

[6] Basel Committee issues principles for operational resilience and risk

Basel Committee on Banking Supervision / 2021-03-31. Final paragraph: continuity/testing, interdependencies, third-party and incident management.

[Read the source](#)

[7] Cyber Security Framework: Introduction

Saudi Central Bank, SAMA Rulebook / 2017-05-24; status in-force at access 2026-10-10. Section 1.1 and status header.

[Read the source](#)

[8] Revised Guidance on Model Risk Management: attachment to SR 26-2

Federal Reserve, OCC and FDIC / 2026-04-17. PDF pages 2-3, Purpose and Scope; footnotes 1-3.

[Read the source](#)

[9] In data we trust? Emerging policy and supervisory approaches to AI data use in financial services

Crisanto, Currat, Ehrentraud and Wu; BIS Financial Stability Institute, FSI Insights 73 / 2026-03-26. Executive summary; Section 2 data challenges; official publication abstract. Authors' views, not necessarily BIS or standard-setter positions; not binding guidance or evidence of business returns.

[Read the source](#)