

Responsible AI Governance in Banking: Integrating Model Risk, Consumer Protection, and Operational Resilience Across the United States, European Union, United Kingdom, Singapore, and Hong Kong

Rajeew Vishvakarma
Project Manager, Infosys

Abstract

Banks are embedding artificial intelligence (AI) into credit underwriting, fraud monitoring, trading, collections, customer service, and anti-financial-crime operations. The resulting risk profile is no longer captured adequately by a single governance lens. A bank AI system can simultaneously be a model-risk object, a consumer-impact mechanism, and a technology dependency whose failure may disrupt critical services. This paper develops an integrated Responsible AI governance framework for banks through a comparative documentary analysis of primary regulatory and supervisory materials across five influential jurisdictions: the United States, the European Union, the United Kingdom, Singapore, and Hong Kong. The analysis shows a convergence of supervisory expectations around three control families: lifecycle model governance, customer protection and explainability, and operational resilience with third-party accountability. Building on this convergence, the paper proposes an integrated control architecture covering governance bodies, AI inventory and use-case classification, risk tiering, validation, explainability, fairness review, deployment controls, incident response, vendor oversight, and ongoing monitoring. Illustrative case studies show that fragmented governance often misses failure modes even where no single regulation has yet been formally breached. The paper concludes with a phased implementation roadmap and argues that banks should govern AI as a cross-cutting enterprise capability rather than as a narrow model-development topic.

Keywords: Responsible AI; banking; model risk management; consumer protection; operational resilience; AI governance; AI Act; DORA; SR 11-7; FEAT; HKMA

1. Introduction

Artificial intelligence has moved from experimentation into production in core banking functions. Financial institutions now use AI in credit decisioning, transaction monitoring, anti-money laundering workflows, fraud detection, personalization, collections, treasury support, document processing, and customer service. These applications promise efficiency, scale, and speed, but they also reshape the bank risk surface. AI can amplify traditional model risk through data drift, poor generalisation, feedback loops, weak validation, and opaque decision logic. It can create consumer-protection risk where automated systems affect pricing, access to credit, treatment in servicing, or the adequacy of explanations given to customers. It can also create operational-resilience risk because many AI capabilities depend on external data sources, cloud infrastructure, foundation models, application programming interfaces, and third-party vendors.

Historically, banks governed these domains in separate silos. Model risk management focused on conceptual soundness, validation, and performance monitoring. Compliance and conduct teams concentrated on fair lending, disclosure, complaints, and customer outcomes. Technology and operational-risk functions addressed continuity, cyber resilience, outsourcing, and incident management. That separation is increasingly untenable. The same credit model may require independent validation under supervisory guidance, explanation and adverse-action controls under consumer law, and contingency planning if it depends on external infrastructure. Likewise, a

customer-service chatbot may be treated as a digital channel, yet it can also create legal and reputational risk if it obstructs required servicing or generates misleading information.

Recent supervisory developments sharpen this convergence. In the European Union, the AI Act explicitly treats certain banking-related uses, including creditworthiness assessment and credit scoring of natural persons, as high-risk. The same institutions also face the Digital Operational Resilience Act (DORA), which places management-body responsibility on ICT risk governance, incident reporting, and oversight of ICT third-party dependencies. In the United States, long-standing model-risk guidance remains highly relevant, while the Consumer Financial Protection Bureau has made clear that creditors using complex algorithms remain fully responsible for providing specific adverse-action reasons and that automation cannot be used to evade customer-service obligations. In the United Kingdom, the PRA and FCA shape the field through outsourcing, operational resilience, and consumer-outcome requirements. Singapore and Hong Kong increasingly articulate practical expectations for fairness, accountability, testing, and governance in financial services.

This paper addresses a practical and policy-relevant question: how should banks integrate model risk management, consumer protection, and operational resilience into a unified Responsible AI governance framework? The paper makes three contributions. First, it synthesises primary legal and supervisory sources across five jurisdictions and identifies the common control logic beneath different regulatory styles. Second, it translates those expectations into a bank operating model with clear ownership, risk-tiering criteria, lifecycle controls, monitoring, and assurance. Third, it demonstrates through public case illustrations that fragmented governance leaves material gaps at the seams between fairness, explainability, resilience, and accountability. The central argument advanced here is that responsible AI governance in banking should not be implemented as an isolated ethics overlay or as a technical add-on to validation. It should instead be embedded into enterprise governance, risk management, compliance, and operational resilience architecture.

2. Research and Supervisory Context

A growing body of scholarship examines the benefits and risks of AI in finance, especially in credit, fraud, and customer interaction contexts. Research consistently shows that machine-learning systems can improve predictive performance, but may also reduce interpretability, complicate human review, and generate disparate impacts or distributional shifts. Work on explainability in credit decisions, fairness-aware learning, and responsible AI therefore has direct relevance for banking environments where decisions are consequential and regulatory scrutiny is high.

The supervisory literature increasingly reflects a convergence around trustworthy and accountable deployment. In the United States, model-risk governance remains anchored in guidance that emphasises governance, validation, and controls proportionate to the scale and complexity of model use. In the European context, prudential supervisors have explored the use of machine learning in internal-ratings-based models while highlighting interactions with legal frameworks on AI and data governance. Standards bodies have reinforced this direction. The NIST AI Risk Management Framework frames trustworthy AI around governance, mapping, measurement, and management, while international bodies such as the Financial Stability Board emphasise concentration, third-party dependency, model risk, and supervisory capability as systemic concerns.

What remains underdeveloped, however, is an integrated banking-specific framework that connects these threads. Much of the existing literature treats fairness, explainability, validation, and resilience as adjacent but separate themes. Yet banking use cases demand integration because a single system can touch all of them simultaneously. A credit underwriting model requires performance validation and monitoring, but it must also support customer-facing explanations and operate inside a resilient technology stack. A fraud model may not fall within every AI-specific statute, yet it still requires governance over false positives, customer treatment, escalation, and

operational dependency. Generative AI adds further complexity because governance must now cover unstructured outputs, prompt-driven behaviour, hallucination risk, vendor updates, and retrieval-quality failures.

The most useful conceptual baseline is therefore not a narrow ethics checklist, but a layered control model. At the first layer, governance establishes accountability, oversight, and risk appetite. At the second layer, lifecycle controls govern design, data, validation, deployment, monitoring, change, and retirement. At the third layer, impact controls address customer outcomes, explainability, fairness, complaints, and human review. At the fourth layer, resilience controls address continuity, incident response, cyber and ICT risk, and third-party oversight. This paper adopts that integrated lens and applies it to cross-jurisdictional banking regulation.

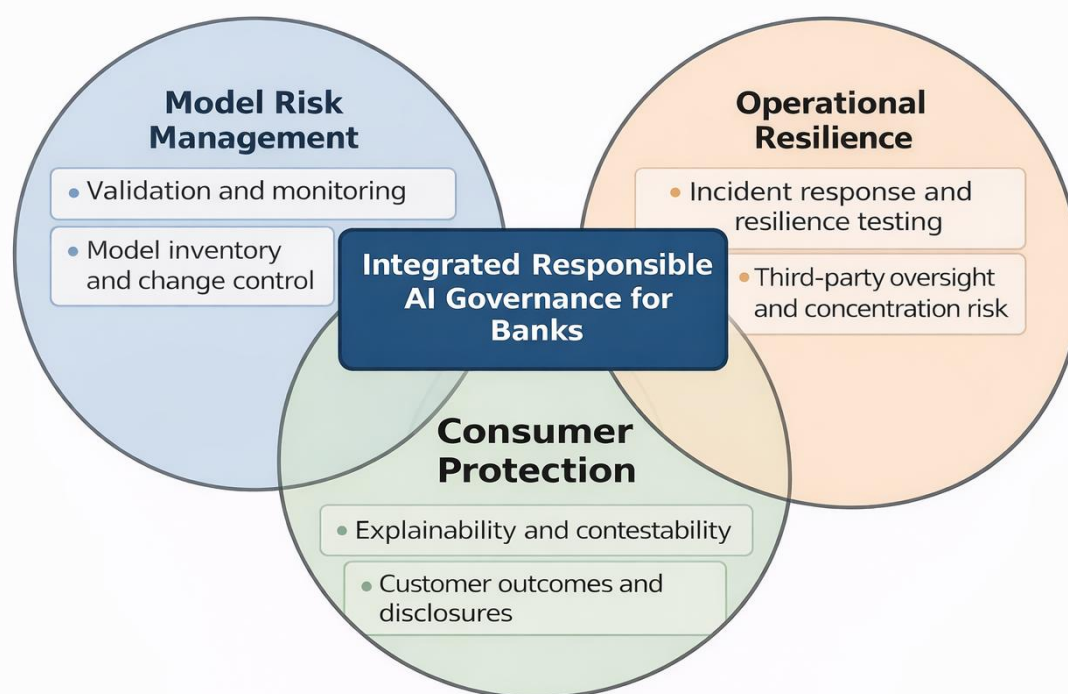


Figure 1. *Convergence of the three supervisory control families that underpin integrated AI governance in banking.*

3. Methodology

The paper uses a qualitative documentary analysis methodology. The primary materials reviewed are publicly available laws, supervisory statements, policy documents, information papers, and official reports relating to AI, model risk, consumer protection, outsourcing, and operational resilience in banking. The core jurisdiction set comprises the United States, the European Union, the United Kingdom, Singapore, and Hong Kong because together they capture a range of regulatory styles and have influenced supervisory practice beyond their domestic markets.

The analysis proceeded in three steps. First, the documents were coded for explicit requirements or expectations relevant to bank use of AI, including governance and board accountability, documentation, data governance, validation, explainability, human oversight, customer treatment, incident reporting, outsourcing controls, and third-party due diligence. Second, those requirements were mapped to recurring bank control mechanisms such as

model inventories, governance committees, validation standards, change management, vendor assessment, and incident response procedures. Third, public case materials and supervisory reports were used to test whether fragmented governance structures would likely miss the types of failures or weaknesses highlighted by regulators. This is a conceptual and normative study rather than an empirical measurement exercise. Its aim is not to estimate model performance, but to derive a regulator-grounded governance design that banks can operationalise.

4. Comparative Regulatory Analysis

Across the five jurisdictions, regulatory language differs, but the underlying control logic is increasingly consistent. Supervisors expect banks to know what AI systems they use, classify risk, document and test them, assign accountable owners, control customer impacts, and manage the operational dependencies on which those systems rely. The most important differences concern the degree of prescriptiveness and the pathway through which expectations are imposed. The United States remains largely framework-based and relies on prudential, fair-lending, consumer-protection, and third-party rules. The European Union now combines horizontal AI regulation with financial-sector resilience law. The United Kingdom governs through sectoral supervision, especially outsourcing, resilience, and conduct. Singapore and Hong Kong remain more principles-oriented, but both have become influential in translating responsible AI ideas into operational guidance for financial institutions.

Table 1. *Cross-jurisdiction summary of core AI governance expectations in banking*

| Control area | United States | European Union | United Kingdom | Singapore | Hong Kong |
|--------------------------------------------|------------------------------------------------------------|----------------------------------------------------------------------------|---------------------------------------------------------|-------------------------------------------------------|-------------------------------------------------------------------|
| Primary posture | MRM + consumer law + vendor oversight | AI Act + DORA + prudential overlays | Prudential + conduct + resilience | Principles + thematic guidance | Principles + supervised experimentation |
| Inventory and accountability | Expected through MRM and compliance | Required for high-risk AI lifecycle evidence | Expected through governance and outsourcing oversight | Strong emphasis in FEAT / Veritas practice | Expected through governance and sandbox-style oversight |
| Explainability / customer treatment | Specific adverse-action reasons; service quality oversight | Transparency, human oversight, FRIA in some contexts | Consumer Duty drives outcome-based transparency | Fairness and transparency are central | Customer disclosure and use-case governance increasingly stressed |
| Validation and testing | Independent challenge, validation, monitoring | Risk management, data governance, logging, testing, post-market monitoring | Risk-based validation embedded in governance | Structured assessment and model governance encouraged | Risk assessment and bounded deployment expected |
| Operational resilience | Safe-and-sound operations and vendor oversight | Board ICT accountability, incident reporting, ICT third-party risk | Important business services, outsourcing, exit planning | Governance and control expectations growing | Supervisory focus on risk-managed innovation |

4.1 United States

In the United States, the most durable starting point remains model risk management. Federal Reserve SR 11-7 makes clear that model governance, validation, and controls apply across the lifecycle of model use and should be proportionate to the size, nature, complexity, and sophistication of the institution model environment. For banks using AI in credit, fraud, or risk management, this means that machine-learning systems are not outside the model-risk perimeter merely because their architecture differs from traditional scorecards or econometric models. Documentation, independent challenge, outcome analysis, benchmarking, limitations analysis, and change control remain essential.

However, U.S. banking AI governance cannot be reduced to SR 11-7 alone. Consumer law imposes separate and sometimes sharper obligations. The CFPB has clarified that adverse-action notice requirements apply equally to complex or so-called black-box models, and that technological opacity is not an excuse for failing to provide specific reasons for a credit denial or other adverse decision. This matters because some AI techniques may produce probabilistic or latent features that are not easily translated into customer-facing explanations. Governance must therefore ensure that design and deployment processes include explanation pathways, reason-code governance, quality assurance, and escalation routes for disputed decisions.

The U.S. approach also highlights customer-service risk. The CFPB work on chatbots underscores that banks cannot use automation as a substitute for legally required servicing quality. A chatbot that misdirects customers, blocks access to hardship options, or creates confusion during complaints and disputes can expose the institution to consumer-harm findings even if the technology performs efficiently from an operating-cost perspective. Operationally, U.S. banks must also account for cloud vendors, AI tool providers, data suppliers, and concentration risk through safe-and-sound operations, business continuity, and third-party oversight. The practical implication is clear: U.S. banks need enterprise AI governance that links model-risk, compliance, operations, and technology functions rather than assuming that existing silos will coordinate spontaneously.

4.2 European Union

The European Union is now the clearest example of regulatory convergence around banking AI. The AI Act classifies AI systems used to evaluate the creditworthiness of natural persons or to establish their credit score as high-risk, except where used solely for fraud detection. This classification activates a set of lifecycle obligations around risk management, data governance, record-keeping, transparency, instructions for use, human oversight, accuracy, robustness, cybersecurity, and post-market monitoring. For banks, these are not only technical requirements. They imply institutional structures: identifiable ownership, a documented intended purpose, provider-deployer role clarity, human review mechanisms, and change-management controls. Article 27 adds a further layer through fundamental-rights impact assessment requirements in specified contexts, reinforcing that deployment decisions must consider effects on individuals as well as model performance.

The second major EU pillar is DORA. Whereas the AI Act governs high-risk AI as a product and lifecycle matter, DORA governs the digital environment in which banking AI operates. It requires an internal governance and control framework for ICT risk, places ultimate responsibility on the management body for ICT arrangements, requires reporting of major ICT-related incidents, and treats ICT third-party risk as an integral component of the framework. This is especially important for AI because many banking use cases depend on vendors, cloud platforms, model providers, data pipelines, and outsourced software components.

The combined effect is that a bank cannot responsibly govern AI in the EU merely by validating model performance. It must also evidence resilience, logging, oversight, change management, due diligence on third parties, incident reporting capability, and alignment between AI controls and enterprise ICT governance. The EBA work on machine learning for internal-ratings-based models reinforces this direction by identifying explainability, overfitting, management understanding, and skills constraints as supervisory concerns. The EU therefore offers the most explicit legal foundation for integrated AI governance in banking today.

4.3 United Kingdom

The United Kingdom has not adopted a single cross-sector banking AI statute comparable to the EU AI Act. Instead, UK banks face a combination of prudential, outsourcing, operational-resilience, and conduct obligations that together shape AI governance. PRA Supervisory Statement SS2/21 sets expectations for outsourcing and third-party risk management, which is directly relevant to AI because many banks procure AI capabilities from external vendors or depend on third-party infrastructure, data, and software components. A bank remains

accountable for regulated outcomes even where technology is sourced externally. That principle pushes firms toward stronger due diligence, contractual control, contingency planning, concentration-risk analysis, and exit-strategy design for AI providers.

The FCA Consumer Duty sets higher and clearer standards for consumer protection and requires firms to put customers needs first. Although technology-neutral, the Duty has obvious implications for AI-enabled products and channels. If an algorithmic product design, communication, or servicing workflow leads to poor customer outcomes, opacity will not excuse weak governance. For AI uses in retail banking, governance must therefore address explainability, suitability of communications, challenge and appeal pathways, oversight of vulnerable-customer impacts, and evidence that customer outcomes are being monitored.

Operational resilience provides the third UK pillar. Where AI supports important business services, banks should determine whether disruption, erroneous outputs, or vendor failure could cause intolerable harm to customers or market operations. This makes resilience testing, fallback procedures, and escalation design especially important for high-dependency AI use cases. The UK model therefore encourages banks to govern AI through enterprise risk disciplines rather than through a separate ethics-only track.

4.4 Singapore

Singapore has been a leading jurisdiction in articulating operationally useful responsible AI guidance for finance. The Monetary Authority of Singapore FEAT Principles - fairness, ethics, accountability, and transparency - provide a foundational vocabulary for AI and data analytics in financial services. Subsequent industry work, including the Veritas initiative and toolkit, has sought to translate those principles into assessment methodologies and implementation practice. MAS has also supported sector-focused work on generative AI and published an information paper on AI model risk management based on thematic review work with banks.

The Singapore approach is significant for two reasons. First, it recognises that responsible AI in financial institutions must be operationalised through governance, testing, and documentation rather than left at the level of values statements. Second, it treats fairness and accountability as management topics rather than purely technical properties of models. For banks, the practical implication is that AI governance should include explicit policy principles, accountability mapping, use-case review, explainability standards proportionate to use, and structured oversight of generative AI. Singapore therefore offers a strong implementation reference point for banks designing governance before binding requirements become highly prescriptive.

4.5 Hong Kong

Hong Kong has also taken a pragmatic, finance-specific approach to AI governance. The HKMA work on AI application in banking helped frame sector expectations around risk, accountability, transparency, and adoption barriers. More recent HKMA activity on generative AI signals a move from exploratory interest toward more operational governance. Through research papers, circulars, and sandbox arrangements, the HKMA has emphasised the need for risk assessment, data management, customer protection, and control over customer-facing deployment.

The Hong Kong approach is relevant for institutions balancing innovation ambition with supervisory caution. It does not rely solely on formal rulemaking. Instead, it uses white papers, principles, and supervised experimentation to steer governance practice. For banks, that means AI governance should be sufficiently mature to demonstrate use-case traceability, accountability, testing, customer disclosures where appropriate, and controls over data use and vendor components. Taken together, Singapore and Hong Kong show how Asian financial centres have operationalised responsible AI through a mix of principles, thematic guidance, and controlled innovation.

5. An Integrated Responsible AI Governance Framework for Banks

The comparative analysis points to a central conclusion: banking AI governance must be built as an integrated system rather than as a stack of disconnected reviews. A workable framework should combine nine components. First, governance and accountability. Banks should establish clear executive sponsorship and a cross-functional AI governance forum that coordinates model risk management, compliance, legal, technology, operational risk, cybersecurity, data governance, internal audit, and relevant business lines. Second, enterprise inventory and use-case taxonomy. Every AI use case should be registered with minimum metadata including purpose, business owner, model owner, deployment channel, affected products and customers, data sources, third-party dependencies, and applicable jurisdictions. Third, risk tiering and impact assessment. Each use case should be classified using a framework that considers regulatory criticality, consumer impact, model complexity, autonomy, operational dependency, and legal sensitivity.

Fourth, lifecycle model governance. Design, development, validation, and change management should follow rigorous standards including intended purpose, data documentation, performance limits, challenge testing, back-testing, benchmarking, monitoring thresholds, and change logs. For generative AI, additional controls should cover prompt management, output constraints, prohibited-behaviour testing, retrieval quality, and deployment boundaries. Fifth, explainability, fairness, and customer-impact controls. High-impact AI should undergo fairness review and explanation-readiness assessment before deployment. Governance should confirm that adverse-action reasons, customer communications, escalation routes, and human-review pathways are adequate. Sixth, deployment controls and human oversight. Production release should require evidence that monitoring, logging, fallback arrangements, owner sign-offs, and access controls are in place. The degree of human review should depend on use-case criticality, model confidence, and legal impact.

Seventh, operational resilience and incident response. For important AI-enabled services, banks should identify failure scenarios and determine whether they threaten important business services or regulatory obligations. Incident taxonomies should include erroneous outputs, drift, unavailable model services, degraded vendor performance, data-feed corruption, prompt-injection events, and customer-harm incidents caused by automated interactions. Eighth, third-party and concentration-risk management. Vendor-provided AI creates layered dependency risk, so due diligence should cover technical capability, governance standards, subcontracting, data handling, audit rights, model-update notifications, resilience posture, and exit options. Ninth, monitoring, assurance, and audit. AI governance should generate management information that reaches senior leaders, including inventory completeness, validation timeliness, override rates, fairness alerts, drift alerts, incident trends, vendor concentration, and remediation aging. Internal audit should assess not only whether policy exists, but whether the governance system works in practice.

AI use-case intake, risk tiering, and governance workflow

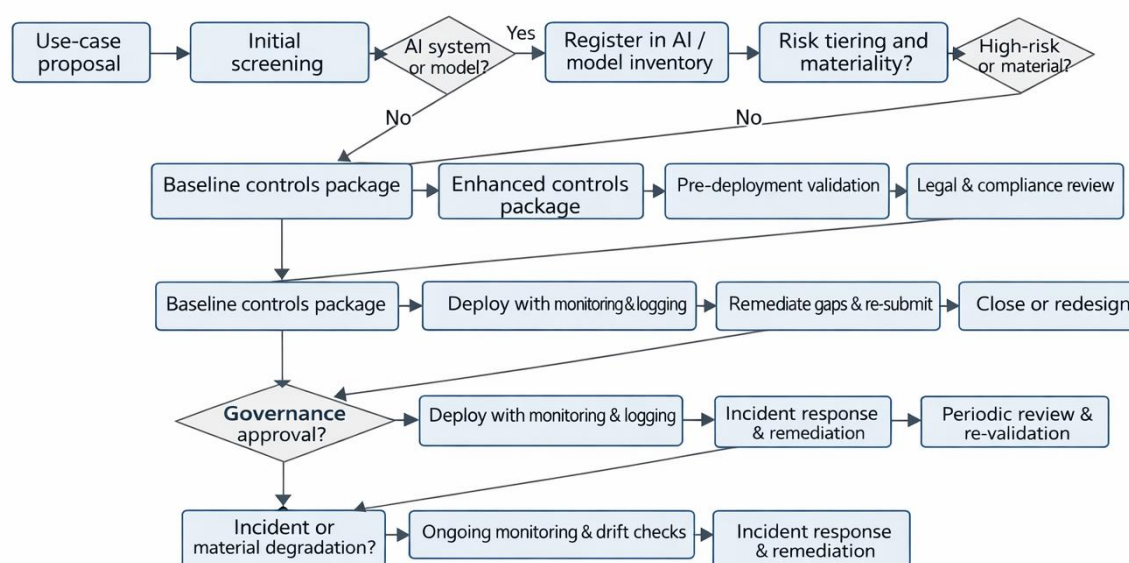


Figure 2. AI use-case intake, risk tiering, validation, approval, and post-deployment governance workflow.

Table 2. Core framework components, leading owners, and illustrative artefacts

| Framework component | Lead owner(s) | Illustrative artefacts / outputs |
|----------------------------------------------|--------------------------------------------------------|----------------------------------------------------------------------|
| AI policy and accountability | Board / executive sponsor / AI committee | Policy, risk-appetite statement, committee charter |
| Enterprise inventory and taxonomy | Business owners + model risk + enterprise architecture | AI register, use-case metadata, ownership map |
| Risk tiering and impact assessment | Model risk + compliance + legal | Tiering rubric, impact assessment, jurisdictional applicability memo |
| Validation and lifecycle control | Model risk management | Validation reports, monitoring thresholds, change-control log |
| Fairness and explanation governance | Compliance + fair-lending / conduct teams | Reason-code mapping, fairness review record, escalation workflow |
| Operational resilience and incident response | Operational risk + technology + cyber | Critical-service mapping, playbooks, incident taxonomy |
| Third-party AI oversight | Procurement + legal + technology risk | Due-diligence file, contractual clauses, exit plan |
| Monitoring, assurance, and audit | First line + second line + internal audit | Dashboard, issue log, audit plan, remediation tracker |

6. Illustrative Case Studies and Governance Lessons

The practical value of integrated governance becomes clearer when viewed through public case material. The first illustration is the Apple Card investigation by the New York Department of Financial Services. The investigation did not find evidence of unlawful discrimination in the underwriting data reviewed, but it showed how algorithmic decisioning can create concern where explanations, escalation pathways, and customer understanding are weak. Governance takeaway: even where discriminatory intent is absent, banks need documented rationale, meaningful customer communication, and structured human review for disputed automated outcomes.

The second illustration concerns chatbots in consumer finance. The CFPB has highlighted the rapid adoption of chatbots by major banks and the risks created when automated service channels substitute for effective customer assistance. A chatbot may appear operationally efficient while still generating legal and conduct risk if it misstates rights, loops customers endlessly, or obstructs access to hardship or dispute processes. Governance takeaway: customer-service automation should be treated as a regulated decision environment, not merely as a digital-experience feature.

The third illustration is the EBA follow-up work on machine learning for internal-ratings-based models. The report shows that the prudential use of machine learning cannot be considered in isolation from broader legal and governance concerns. Even where a model supports credit-risk estimation rather than a direct consumer-facing decision, questions of explainability, governance, and interaction with AI regulation arise. Governance takeaway: prudential models require governance broad enough to anticipate cross-regime obligations. The fourth illustration comes from MAS thematic work on AI model risk management. Even where guidance is expressed as soft law, supervisory information papers are revealing because they show what banks are expected to operationalise, including inventories, governance, testing, and structured controls. Governance takeaway: banks should treat strong supervisory signals as early governance requirements rather than waiting for highly prescriptive binding rules.

7. Implementation Roadmap

A bank seeking to operationalise the framework proposed here can proceed in three phases over roughly twelve to fifteen months. Phase one is foundation. The institution should appoint an executive sponsor, approve an enterprise Responsible AI policy, define the scope of the AI inventory, establish a cross-functional governance forum, and agree on a risk-tiering methodology. This phase should also define common terminology because many institutions use inconsistent labels for models, analytics, automation, machine learning, and generative AI.

Phase two is control build-out. The bank should complete the inventory, classify use cases, and align lifecycle controls across model-risk, compliance, and technology teams. Validation standards should be refreshed to address AI-specific issues such as data representativeness, feature drift, explanation quality, prompt robustness, and retrieval quality where applicable. Compliance and conduct teams should build standard review templates for fairness, reason-code governance, customer communication, and complaint feedback. Technology and operational-risk teams should identify AI-supported important business services, key dependencies, and failure scenarios.

Phase three is assurance and scaling. The institution should pilot the framework on selected high-impact use cases, generate management reporting, and embed escalation and incident pathways. Internal audit should review both policy design and implementation effectiveness. Vendor contracts should be remediated where rights, transparency, update controls, or exit planning are inadequate. Once the framework is stable, it can be extended to lower-risk use cases using proportional controls. Proportionality is essential throughout: smaller banks do not need the same committee density or specialist tooling as global systemically important banks, but they still need clear ownership, inventory discipline, risk tiering, and escalation design.

AI incident response for banking systems

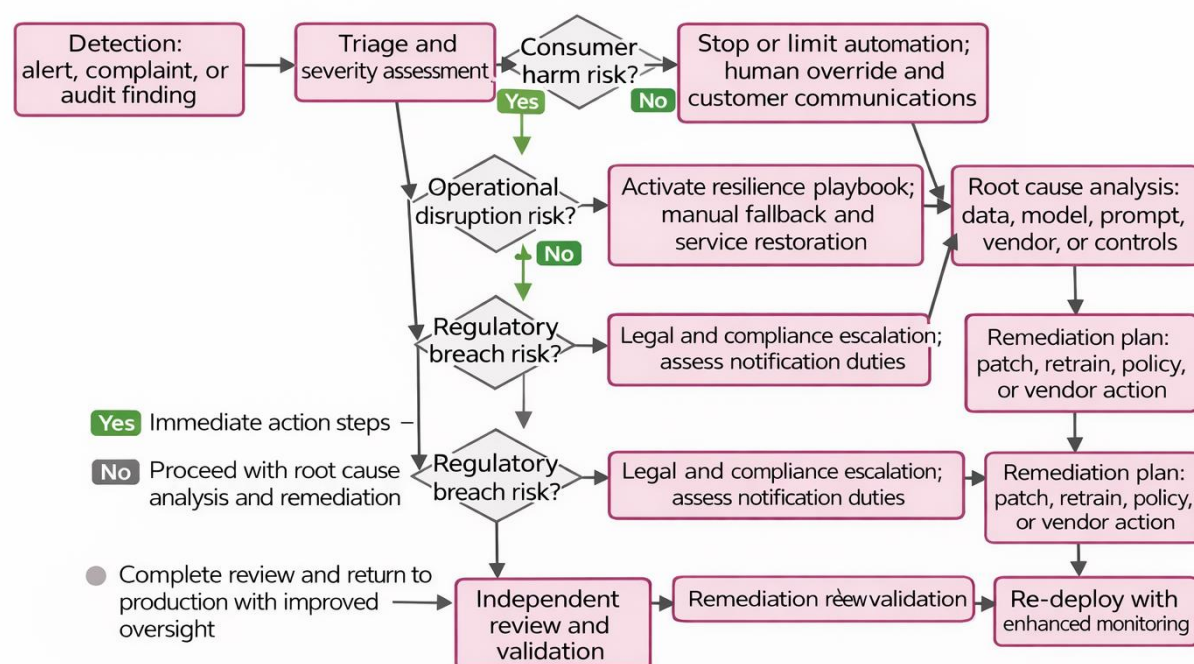


Figure 3. AI incident response logic linking customer harm, operational disruption, and regulatory escalation.

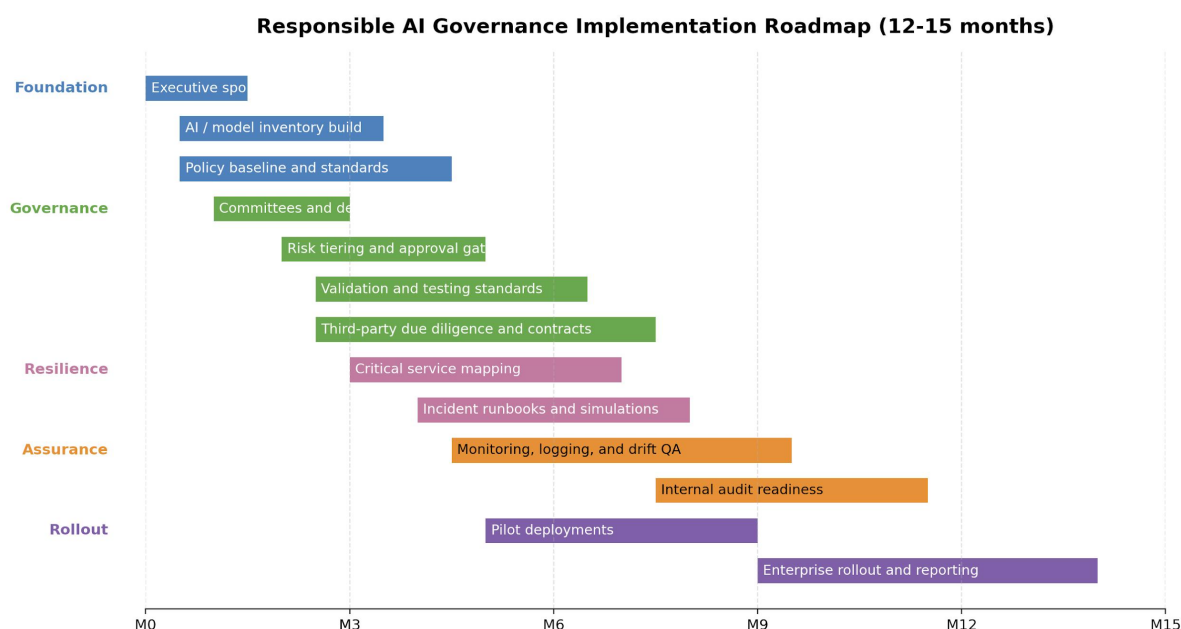


Figure 4. Indicative implementation roadmap for a bank-scale Responsible AI governance programme.

8. Discussion and Limitations

The framework proposed in this paper is intentionally normative. It argues for a governance design broader than many current bank operating models. Some institutions may object that integrating model risk, consumer protection, and resilience creates process friction or duplicates existing controls. That concern is understandable, but it underestimates the cost of fragmented governance. AI-related failures often surface precisely because no

single function sees the whole picture. A technically strong model can still cause consumer harm if explanation and complaint-handling pathways are weak. A compliant use case can still create major operational disruption if vendor concentration or continuity planning is neglected.

A second limitation is that regulation remains dynamic. The AI Act is still moving through phased implementation, supervisory practice under DORA will continue to mature, and Asian jurisdictions are actively evolving their guidance for generative AI. The framework should therefore be viewed as a living governance architecture rather than a static checklist. Its value lies in the durability of the core control families: accountability, inventory, risk classification, lifecycle control, customer-impact control, resilience, and assurance are unlikely to become less important as AI adoption expands.

A third limitation is that this paper relies on public sources. Confidential examination findings could reveal additional nuance about how supervisors prioritise controls in practice. Future work could add expert interviews, comparative policy scoring, or bank-policy audits to test implementation maturity. Even so, the documentary evidence already supports the main claim that banking AI governance is converging toward integrated oversight across prudential, conduct, and resilience domains.

9. Conclusion

AI governance in banking is no longer a niche question for data-science teams. It is an enterprise governance challenge with direct implications for safety and soundness, customer treatment, and continuity of critical services. The comparative analysis in this paper shows that regulators across the United States, European Union, United Kingdom, Singapore, and Hong Kong are approaching the problem through different legal pathways but with increasingly similar expectations. Banks are expected to know their AI use cases, classify risk, document and validate systems, maintain meaningful oversight, protect customers, manage third-party dependencies, and respond effectively when technology fails or behaves unexpectedly.

The central argument of this paper is that these expectations are best met through an integrated Responsible AI governance framework. Such a framework should connect model risk management, consumer-protection controls, and operational resilience rather than leaving them in parallel silos. Doing so does not merely improve compliance posture. It creates better management information, clearer accountability, more resilient deployment, and greater institutional capacity to innovate responsibly. The institutions most likely to succeed will be those that govern AI as a cross-functional capability with board-level visibility, lifecycle discipline, customer-impact awareness, and operational resilience built in from the start.

References

- Board of Governors of the Federal Reserve System. (2011). SR 11-7: Guidance on Model Risk Management.
- Consumer Financial Protection Bureau. (2022). Consumer Financial Protection Circular 2022-03: Adverse action notification requirements in connection with credit decisions based on complex algorithms.
- Consumer Financial Protection Bureau. (2023a). Chatbots in consumer finance.
- Department of Financial Services, State of New York. (2021). DFS issues findings on the Apple Card and its underwriter Goldman Sachs Bank.
- European Banking Authority. (2023). Follow-up report on machine learning for IRB models.
- European Union. (2022). Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (DORA).
- European Union. (2024a). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act).
- Financial Conduct Authority. (2021). PS21/3: Building operational resilience.
- Financial Conduct Authority. (2022). PS22/9: A new Consumer Duty.
- Financial Stability Board. (2024). The financial stability implications of artificial intelligence.
- Hong Kong Monetary Authority. (2019). Report on artificial intelligence application in banking.

- Hong Kong Monetary Authority. (2024). Generative artificial intelligence in the financial services space and sandbox arrangements.
- Monetary Authority of Singapore. (2018). Principles to promote fairness, ethics, accountability and transparency (FEAT) in the use of AI and data analytics.
- Monetary Authority of Singapore. (2023a). MAS-led industry consortium releases toolkit for responsible use of AI in the financial sector.
- Monetary Authority of Singapore. (2023b). MAS partners industry to develop generative AI risk framework for the financial sector.
- Monetary Authority of Singapore. (2024). Artificial intelligence model risk management.
- National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0).
- Prudential Regulation Authority. (2021/2024). SS2/21 - Outsourcing and third party risk management.