

Just in Time

Sound Practices for Responsible Adoption of Artificial Intelligence (AI)

October 2026



Executive Summary

- The Financial Stability Board (FSB) report, [*Sound Practices for Responsible Adoption of Artificial Intelligence \(AI\)*](#), published on 15 October 2026, sets out **12 practices for financial institutions**, organised around **organisation-wide governance and lifecycle management**. The framework responds to the rapid expansion of machine learning, generative AI (GenAI) and agentic systems across credit, fraud prevention, investment management, compliance, cybersecurity, customer service and operations. AI can improve efficiency, analytical capacity, resilience and customer outcomes, but it also creates risks relating to **opacity, unreliable outputs, weak data governance, cyber threats, third-party concentration, automated decisions and autonomous action**. The FSB therefore advocates a **proportionate, technology-neutral approach** aligned with each institution's strategy, risk appetite, business model and systemic relevance. **Boards and senior management remain accountable**; existing risk frameworks should be enhanced rather than replaced; and controls should cover materiality assessment, documentation, model selection, data, explainability, performance, human oversight, cyber risk and third parties throughout the AI lifecycle. Particular attention is given to **agentic AI**, whose ability to plan, use tools and execute multi-step tasks can increase both value and operational risk. The report's central message is that **responsible adoption requires coordinated action** across business, risk, compliance, technology, security, legal and audit functions, together with continuous learning as technologies and regulatory expectations evolve.



At a Glance

01	<u>Introduction: Objectives, Scope and Regulatory Positioning</u>	4
02	<u>An Overview of the Use of AI in the Financial System</u>	6
03	<u>Organisation-wide AI Governance</u>	9
04	<u>AI Lifecycle Management</u>	11
05	<u>Conclusions & Key Takeaways</u>	15

Keywords: AI, Financial Regulation



01

Introduction: Objectives, Scope and Regulatory Positioning



Introduction: Objectives, Scope and Regulatory Positioning

Just in Time

The report builds on the FSB's 2017 and 2024 work on AI and financial stability.

Recent advances in foundation models, large language models and autonomous agents have shifted AI from supporting analysis and individual decisions toward systems capable of planning tasks, interacting with databases and software, and executing actions.



This distinction between **AI-assisted decision-making** and **AI-enabled autonomous execution** is important: autonomy can increase efficiency, but it can also produce actions outside intended parameters that are difficult to monitor or reverse.



The practices are broadly applicable across technologies, while controls should reflect specific characteristics such as **autonomy, delegated authority and intervention capability**. They are also deliberately evolutionary because understanding of agentic and multi-agent risks remains immature. Implementation should be **proportionate to the institution, function and use case**: an internal productivity tool should not face the same controls as an autonomous system involved in payments or consequential customer decisions. The case studies illustrate reported benefits and governance approaches rather than universal benchmarks.



Following the June 2026 consultation, the final report clarifies board and senior management responsibilities, strengthens documentation and transparency expectations, expands third-party considerations and adds guidance on agentic AI. Overall, the FSB positions **responsible AI as an enterprise governance issue** requiring coordinated ownership across the three lines of defence, not as a technology project managed in isolation.

02

An Overview of the Use of AI in the Financial System



An Overview of the Use of AI in the Financial System 1/2

Just in Time

AI adoption in the financial sector

AI is becoming a major driver of change across banking, insurance, capital markets, asset management, payments and financial market infrastructures. Adoption is supported by greater computing capacity, cloud infrastructure, data availability and foundation models, while institutions seek **efficiency, better analytics, stronger customer engagement and competitive advantage**. Traditional machine learning mainly supports prediction, classification and recommendations; large language models process language; GenAI creates content; and **agentic AI can plan and execute complex activities with varying autonomy**. Adoption nevertheless remains uneven. Banks use AI in credit risk, fraud detection, controls and document processing, but tend to be cautious in consequential front-office decisions. Insurers apply it to underwriting and claims; asset managers to research, portfolio analysis and forecasting; payment providers to monitoring and onboarding; and infrastructures to surveillance, settlement, reconciliation and anomaly detection. The pace depends on **institutional capability, regulatory context, technological maturity and risk tolerance**.

Use cases and benefits of AI in the financial system

The report illustrates **benefits across credit, investment management, compliance, customer service and operations**. In credit, AI can combine financial, behavioural and alternative data to improve risk differentiation, early-warning systems and preliminary underwriting. One regional bank reportedly identified **more than 95% of non-performing SME loans** at least three months before repayment difficulties, while another reduced preparation of preliminary credit proposals from over a day to about **15 minutes**. In investment management, AI supports research, sentiment analysis, liquidity forecasting and portfolio optimisation; however, evidence of fully autonomous GenAI trading remains limited, and case studies retain **final accountability with human risk managers**. In compliance and financial crime, machine learning can detect unusual transactions and relationships, while GenAI can structure investigations. One bank's agentic fraud system proposed detection rules subject to analyst approval and reportedly helped reduce fraud losses by **more than 20%**.

An Overview of the Use of AI in the Financial System 2/2

Just in Time

Use cases and benefits of AI in the financial system

Customer applications include chatbots, relationship-management support, complaints and claims processing; reported gains include halved preparation time and substantially faster claims handling. GenAI is also widely used for coding, document preparation, retrieval and workflow automation. One insurer reported removing about **400,000 manual processing instances** and reducing a preliminary underwriting step to roughly **45 seconds**. AI can strengthen operational resilience through faster cyber monitoring and continuity tools; a payment network, for example, uses more than 200 variables to support authorisation during issuer outages. These outcomes are institution-specific, and **benefits must be weighed against customer protection, staff deskilling and control effectiveness**.

Risks and implementation challenges

The FSB highlights **four vulnerabilities with potential financial-stability effects: third-party concentration, market correlations, cyber risk, and model, data and governance weaknesses**. Poorly defined accountability and fragmented implementation can produce inconsistent controls, while **shadow AI reduces visibility**. Low-quality or unrepresentative data can generate errors and bias; limited explainability can obstruct validation, compliance and customer redress; and performance can deteriorate through hallucinations, drift or changed conditions. Human reviewers may succumb to **automation bias** or lose the expertise needed to challenge outputs. AI systems are also exposed to **data poisoning, prompt injection, jailbreaking, evasion and model extraction**, while malicious actors can use AI for phishing, deepfakes and code generation. Dependence on cloud, model, data and platform providers can create **common points of failure**. Customer and conduct risks include unsuitable recommendations, discrimination, inadequate disclosure and weak challenge processes. **Agentic systems intensify these concerns** because they can execute transactions, access data, use tools and interact with other agents before human intervention is possible; multi-agent delegation and sub-agent creation may further reduce traceability and control.

03

Organisation-wide AI Governance



Organisation-wide AI Governance

Strategic direction and oversight

Sound Practice 1

Boards and senior management should align AI adoption with the institution's business model, strategy and risk appetite, assessing both adoption risks and the strategic cost of failing to use technologies that may improve competitiveness, fraud prevention or cyber defence. The board provides oversight and promotes accountability; senior management implements the approach and defines expected outcomes. Management information should cover **adoption, performance, incidents, control deficiencies, audit findings, customer outcomes and remediation**. Boundaries should specify **acceptable and prohibited uses**, especially for autonomous or consequential applications, taking account of explainability, customer impact, regulation and dependencies. Institutions should allocate sufficient skills, infrastructure and data capability, and **periodically reassess strategy** as technology, maturity, risks and requirements change.

Governance and accountability

Sound Practice 2

Responsibilities should be explicit and integrated into effective organisational structures. Under the **three lines model**, business functions own AI applications, assess materiality and implement controls; the second line provides independent challenge across model, compliance, conduct, data, cyber and operational risks; and internal audit provides assurance. Control personnel need **independence, expertise and authority**, while material uses may warrant a named accountable executive. The FSB does not prescribe a single structure: centralised models support consistency, decentralised models support specialist needs, and hybrids combine enterprise standards with local flexibility. A case study shows a systemic bank embedding AI into existing risk disciplines and adding a **coordinating AI governance committee** rather than creating a separate framework.

AI risk management framework and effective documentation

Sound Practice 3

Institutions should **incorporate AI into enterprise risk management**, with targeted enhancements for AI-specific characteristics. Frameworks should **identify applications, assess materiality and risk, assign controls and maintain accurate documentation**. Records should cover purpose, design, intended and prohibited uses, data and models, limitations, ownership, approvals, testing, monitoring, incidents and material changes. Existing model inventories, application registers and third-party registries can provide visibility, including over AI embedded in external services. Requirements should be proportionate, but **complex or autonomous systems need sufficient traceability** to reconstruct outcomes and support validation, assurance and audit.

Organisational adaptability

Sound Practice 4

Governance cannot remain static. Institutions should regularly reassess controls, capabilities and oversight as technologies, use cases and threats evolve. Staff involved in development, deployment, use, monitoring and assurance need **role-appropriate AI literacy** and should be encouraged to challenge outputs, report unexpected behaviour and escalate concerns. Learning should incorporate **incidents, near misses, validation findings and drift**, while external monitoring should cover innovation, cyber threats, industry practices and regulation. Engagement with peers, providers, academics, supervisors, sandboxes and public-private initiatives can improve shared understanding, particularly for rapidly evolving agentic AI.

04

AI Lifecycle Management




AI Lifecycle Management 1/3

Just in Time

Lifecycle management applies risk controls from planning through retirement, rather than concentrating them at initial approval. Controls should respond to **changes in model behaviour, data, operating conditions, technology and third-party dependencies**.

Stages of AI development and deployment



AI Lifecycle Management 2/3

Materiality and risk assessment

Sound Practice 5

Institutions should assess **both materiality and risk at inception and after material change**. Materiality reflects importance to operations, customers and legal obligations; risk reflects the likelihood and severity of adverse outcomes. Factors include **scale, affected customers, financial consequences, reversibility, model limits, data, explainability, autonomy, bias, cyber exposure and third parties**. Results should determine the intensity of **validation, monitoring and human oversight**. Reassessment should follow model or use changes, new dependencies, deteriorating performance or threats.

Selection process

Sound Practice 6

Model selection should start from the business objective and use-case risk, not technological novelty. Institutions should evaluate **suitability, accuracy, robustness, explainability, scalability, cyber resilience, implementation complexity, data, skills and monitoring**. A simpler model may be preferable where it achieves the objective with greater controllability. Choices among internal, open-source and third-party solutions create different governance and dependency risks, and selected systems should remain subject to **documented change management**.

Data governance

Sound Practice 7

Data must be **accurate, complete, consistent, secure, representative and fit for purpose**. Governance should cover **provenance, quality, access, privacy, confidentiality and intellectual property**, as well as changes in the relationship between training data and the operating environment. Diverse structured, unstructured and synthetic sources create additional challenges. For proprietary third-party models, institutions may need assurance over providers' controls and compensating measures where direct access is unavailable. **Data governance is a continuous control supporting reliability and accountability**.

Explainability and transparency

Sound Practice 8

Explainability varies by technology and should be considered in selection, development and use. More interpretable models may be preferable where performance is comparable. Post-hoc techniques can support analysis but provide partial, assumption-dependent insight; triangulating methods may improve confidence. Where adequate explanation is infeasible, institutions should apply **compensating controls such as stronger testing, restricted use, enhanced monitoring and meaningful human oversight**. Transparency should be proportionate to the consequences of the use case and support **accountability, assurance, compliance and customer redress**.

AI Lifecycle Management 3/3

Just in Time

Performance management

Sound Practice 9

Performance evaluation should reflect the technology and use case. Predictive models can often be tested against outcomes; GenAI may require qualitative criteria and expert judgement because several answers can be acceptable; and **agentic AI must be assessed across sequences of actions as well as final outputs**. Testing should cover **accuracy, reliability, robustness, benchmarking and suitability**, with stricter tolerances for critical uses. Independent assessment and post-deployment monitoring should detect **hallucinations, omissions, drift and inappropriate actions**, with changes triggering reassessment and remediation.

Human oversight

Sound Practice 10

Human oversight must be meaningful: reviewers need knowledge, authority, information and incentives to challenge outputs and intervene. Controls should address **automation bias and deskilling** and should vary with materiality and autonomy, from approval before consequential decisions to monitoring within predefined boundaries. Institutions should record overrides and test whether oversight actually works. For agentic AI, continuous review may be impractical, so systems need **behavioural limits, escalation, monitoring and effective interruption**. Oversight requirements should also inform third-party selection and due diligence.

Cyber and ICT risk management

Sound Practice 11

AI can strengthen cyber defence while creating new attack surfaces. Existing secure-development, access, network, vulnerability, monitoring and incident-response controls should be adapted for **data poisoning, prompt injection, safeguard bypass, model extraction and AI-generated code**. AI may also accelerate phishing, vulnerability discovery and exploitation. The FSB favours **defence-in-depth, limited exposure, segmentation, strong authentication and secure-by-design principles**. Shadow AI should be controlled. Agentic systems require tightly scoped permissions, isolation, monitoring and response capabilities because compromised credentials or instructions can propagate actions across connected systems.

Third-party AI risk management

Sound Practice 12

Third-party risk extends across cloud, models, platforms, data, libraries and embedded AI, not only direct counterparties. Institutions should identify dependencies and assess **provider governance, security, performance, data management, resilience and supply chains**. Contracts should address data use and retention, permitted uses, service levels, model or version changes, drift, discontinuity and responsibility for autonomous actions. Where proprietary restrictions limit transparency, institutions may need independent assurance and compensating controls. **Concentration in common infrastructure, models or datasets** can create correlated behaviour and common failures, potentially amplifying herding, procyclicality and disruption across the financial system.

05

Conclusions & Key Takeaways



Conclusions & Key Takeaways 1/2

The FSB's final report identifies 12 sound practices designed to support financial institutions in capturing the benefits of AI while managing emerging risks. The following are the principal regulatory and operational takeaways.

AI governance is a board-level responsibility

Boards and senior management should ensure that AI adoption is aligned with business strategy, risk appetite and institutional objectives, supported by clear accountability and effective oversight.

A risk-based and proportionate approach is essential

AI controls should reflect the materiality, complexity, autonomy and potential impact of individual use cases. Higher-risk applications, particularly those supporting critical financial functions, warrant more robust governance and monitoring.

Existing risk management frameworks should be enhanced, not necessarily replaced

AI controls should reflect the materiality, complexity, autonomy and potential impact of individual use cases. Higher-risk applications, particularly those supporting critical financial functions, warrant more robust governance and monitoring.

AI transparency, explainability and documentation are critical

AI controls should reflect the materiality, complexity, autonomy and potential impact of individual use cases. Higher-risk applications, particularly those supporting critical financial functions, warrant more robust governance and monitoring.

Human oversight must remain effective

Human involvement should provide meaningful challenge and intervention capabilities. Particular attention is required to address automation bias, overreliance on AI-generated outputs and the potential erosion of staff expertise.

Conclusions & Key Takeaways 2/2

Agentic AI introduces a new dimension of operational risk

Autonomous AI agents may execute unauthorised actions, interact unpredictably with connected systems or expose sensitive information. Their deployment requires appropriate boundaries, monitoring and intervention mechanisms.

AI risk management must cover the entire lifecycle

Risk assessment, testing, validation and performance monitoring should continue from initial development through deployment, operation and retirement, accounting for model drift and changing operating conditions.

Data governance remains fundamental to reliable AI

Data quality, representativeness, security and appropriate usage are essential to limiting inaccurate outputs, unacceptable bias and potential regulatory breaches.

Cybersecurity and third-party concentration are major financial stability concerns

AI can strengthen cyber resilience but also enable sophisticated attacks. Dependencies on a limited number of cloud, infrastructure and foundation model providers may amplify operational disruptions and systemic vulnerabilities.

Responsible AI adoption requires continuous adaptation

Financial institutions should invest in AI literacy, technical capabilities and organisational learning, regularly updating governance arrangements as technologies, risks and regulatory expectations evolve.

ESSENTIAL SERVICES FOR FINANCIAL INSTITUTIONS

iason is an international consulting firm that has been supporting both financial institutions and regulators in topics related to Risk Management, Finance and ICT since 2008

Strategy

Strategic advisory on the **design** of **advanced frameworks** and **solutions** to fulfil both **business** and **regulatory needs** in Risk Management and IT departments

Methodology & Governance

Implementation of the designed **solutions** in bank departments **Methodological support** to both **systemically important financial institutions** and **supervisory entities**

Solution

Advanced **software solutions** for **modelling, forecasting, calculating** metrics and **integrating** risks, all on cloud and distributed in Software-as-a-Service (**SaaS**)

Company Profile

iason is an international firm that consults Financial Institutions on Risk Management.
iason integrates deep industry knowledge with specialised expertise in Market, Liquidity, Funding, Credit and Counterparty Risk, in Organisational Set-Up and in Strategic Planning.

Dario Esposito



Lorena Corna



This is an **iason creation**.

The ideas and the model frameworks described in this presentation are the fruit of the intellectual efforts and of the skills of the people working in iason. You may not reproduce or transmit any part of this document in any form or by any means, electronic or mechanical, including photocopying and recording, for any purpose without the express written permission of iason.

www.iasonltd.com