FSB FINANCIAL STABILITY BOARD

Sound Practices for Responsible Adoption of Artificial Intelligence (AI): Consultation report

Response to Consultation

Li L (independent response)

1. Do you agree with the benefits and risks of AI adoption by financial institutions described in this report? Are there any substantive benefits or risks not covered?

Yes, broadly. The report provides a balanced and practical description of the benefits and risks of AI adoption by financial institutions. I agree that AI can improve operational efficiency, fraud detection, credit assessment, customer service, compliance monitoring, cyber defence, and internal knowledge management. I also agree that AI can amplify model risk, data-governance risk, explainability challenges, cyber risk, conduct risk, and third-party dependency.

I would suggest that the final report give more explicit attention to one additional risk channel: infrastructure-level dependency and financing-related risk arising from the AI infrastructure that supports financial institutions' AI adoption.

AI adoption by financial institutions does not occur only at the model or application layer. It relies on a broader stack of critical inputs, including cloud infrastructure, data centres, GPU and accelerator capacity, high-speed networks, foundation models, data vendors, energy supply, and deployment platforms. A financial institution may have sound governance over its internal AI use case, but still be exposed to concentrated external dependencies at the infrastructure level. If many financial institutions rely on the same cloud providers, foundation models, chip supply chains, or AI infrastructure platforms, the resulting concentration could become a common vulnerability across institutions and jurisdictions.

A related issue is that AI infrastructure is increasingly financed through banks, leases, special-purpose vehicles, private credit funds, infrastructure funds, insurers, and other institutional capital. From a banking risk-governance perspective, this creates a second-order but important financial-stability channel. The physical data-centre shell may have a long life, while the compute layer, such as GPUs and AI servers, may lose financing value much more quickly due to hardware generation shifts, energy-efficiency gaps, or weaker secondary-market liquidity. In such cases, the asset may remain operational while its refinancing value weakens before any payment default occurs.

Therefore, I would recommend that the FSB consider adding "AI infrastructure dependency and financing risk" as a substantive risk category, especially under third-party risk, ICT risk, and financial-stability monitoring. This would not discourage AI adoption. Rather, it would

help institutions understand that responsible AI adoption requires not only model governance, but also a look-through assessment of the infrastructure, contracts, funding structures, and concentration risks that enable AI at scale.

2. **Are the sound practices sufficiently comprehensive and clear to enable financial institutions' responsible AI adoption?**

Yes, the sound practices are comprehensive and clear at a high level. In particular, the distinction between organisation-wide AI governance, AI lifecycle management, cyber and ICT risk, and third-party risk is useful. The emphasis on board and senior management oversight is also important, because AI adoption is no longer only a technology-management issue; it is becoming part of business strategy, risk appetite, resilience planning, and financial-stability governance.

I would suggest three clarifications to make the sound practices more operational for financial institutions.

First, Sound Practices 1 to 4 could more explicitly require boards and senior management to consider AI dependency in setting risk appetite. Risk appetite should not only define which AI use cases are permitted or prohibited. It should also define the institution's tolerance for reliance on external AI infrastructure, foundation models, cloud platforms, data providers, and highly concentrated technology vendors. For systemically important institutions or institutions using AI in critical functions, this dependency should be visible to the board.

Second, Sound Practice 3 on AI risk-management frameworks could encourage institutions to maintain an "AI dependency map". This map would identify material AI use cases, the models used, critical data sources, third-party providers, hosting arrangements, cloud regions, fallback arrangements, and interconnections with core banking systems. Such a map would help institutions understand whether AI risks are isolated within a use case or embedded in critical operations.

Third, Sound Practice 12 on third-party risk management could further distinguish between ordinary vendor risk and critical AI infrastructure risk. In some cases, a third-party AI service provider is not merely an outsourcing vendor. It may be part of a concentrated AI supply chain involving cloud computing, specialised hardware, foundation models, power supply, and platform governance. Due diligence should therefore include not only service performance and data protection, but also capacity constraints, business continuity, vendor model changes, portability, exit strategies, and concentration across the financial sector.

These clarifications would make the practices more useful for banks, insurers, asset managers, payment firms, and nonbanks that increasingly depend on external AI infrastructure.

3. **Do the sound practices strike an appropriate balance between managing risks relating to all forms of AI, and addressing some of the risks relating to emerging and new complex forms of AI, such as GenAI and agentic AI?**

Yes, the practices strike a broadly appropriate balance. A technology-neutral and risk-based approach is preferable because financial institutions use a wide range of AI techniques, from

traditional machine-learning models to GenAI tools and agentic AI systems. The same governance principles should apply across technologies, but controls should become more robust as autonomy, complexity, opacity, and materiality increase.

For GenAI and agentic AI, I suggest that the final report place more emphasis on operational boundaries. The distinctive risk of agentic AI is not only that outputs may be inaccurate. It is that an AI agent may interact with databases, APIs, workflow systems, trading tools, customer-service platforms, fraud-monitoring systems, credit workflows, or cybersecurity tools. This creates the possibility of unauthorised actions, excessive automation, cascading errors, or rapid execution of harmful decisions.

Therefore, human oversight should be complemented by technical and procedural boundaries. These may include permission limits, transaction limits, approval thresholds, independent logging, audit trails, rollback arrangements, kill switches, real-time exception monitoring, and pre-defined escalation procedures. For high-materiality use cases, the relevant question should not simply be whether a human is “in the loop”, but whether the human can meaningfully understand, intervene, override, and remediate the AI system before harm spreads.

This approach would preserve the flexibility of the FSB’s practices while making them more practical for emerging AI systems with higher autonomy.

4. Are the sound practices sufficiently flexible to accommodate and address newer types of AI and responsible adoption over time?

Yes, the sound practices are sufficiently flexible in principle, especially because they are framed as a menu of practices rather than prescriptive international standards. However, flexibility over time depends on whether financial institutions have clear triggers for reassessment. AI risks evolve quickly, and a use case that appears low-risk at inception may become higher-risk as the model, vendor, data, infrastructure, or business process changes.

I recommend that the final report encourage financial institutions to define specific reassessment triggers. These may include: major model-version changes; migration from an in-house model to a third-party model; changes in cloud hosting or data-centre location; deterioration in vendor transparency; new AI features released by third-party vendors; increased autonomy of an AI system; data drift or model drift; cyber incidents; material changes in laws or supervisory expectations; loss of portability or exit options; concentration of workloads in a small number of providers; and disruptions in energy or compute availability.

For financial institutions that finance AI infrastructure, reassessment triggers should also include hardware generation shifts, accelerated technological depreciation, contract-renewal pressure, secondary-market liquidity deterioration, energy-cost shocks, and refinancing-window compression. These risks may arise before default metrics or accounting indicators show deterioration.

In this sense, responsible AI adoption should be treated as a dynamic governance process rather than a one-time model approval. A practical approach would be to combine periodic

review with event-driven reassessment. The frequency and depth of reassessment should depend on use-case materiality, autonomy, explainability, third-party reliance, and systemic relevance.

5. **Do the case studies in this report sufficiently highlight how different types of financial institutions can benefit from responsible AI adoption? Are there additional case studies for inclusion in the report? If so, please provide such case studies, particularly for nonbanks.**

The case studies are helpful, especially in showing how banks and other financial institutions may use AI in credit risk management, fraud detection, relationship management, operational efficiency, and cyber resilience. However, I believe the final report would benefit from an additional case study on nonbank participation in AI infrastructure finance and the related bank-nonbank risk channel.

A possible case study could be titled: “Risk migration in AI infrastructure finance: from bank credit to nonbank capital.”

In this case study, an AI data-centre or compute-infrastructure project is financed through a layered structure. A commercial bank may originate, arrange, or provide a warehouse facility for the transaction. The project assets may be held by a special-purpose vehicle. The long-lived physical layer includes land, buildings, grid connections, power systems, and cooling infrastructure. The shorter-lived compute layer includes GPUs, AI servers, networking equipment, storage, and specialised configurations. The cash-flow layer may depend on a long-term compute-service contract, capacity offtake agreement, lease, or anchor customer arrangement. The final risk may be held by private credit funds, leasing platforms, infrastructure funds, insurers, pension funds, or other institutional investors.

Such a structure can be beneficial. It may diversify funding sources, reduce pressure on bank balance sheets, and match certain long-duration claims with institutional investors that have a longer investment horizon. Nonbank participation can therefore support productive AI infrastructure investment.

The risk is that risk migration may be mistaken for risk elimination. The underlying collateral remains technology-sensitive. A GPU cluster or AI server pool may remain operational while losing financing value after a hardware generation shift, a change in energy efficiency, weaker workload demand, or a decline in secondary-market liquidity. The project may still pay interest, but the next refinancing may become more difficult. The first signs of stress may appear not as a bank non-performing loan, but as lower residual-value assumptions, weaker fund marks, tighter leasing terms, reduced warehouse capacity, higher refinancing spreads, or greater reliance on sponsor guarantees.

This case study would be useful because it connects three FSB workstreams: responsible AI adoption, third-party dependencies, and nonbank financial intermediation. It would also show that AI adoption by financial institutions creates not only model risk and operational risk, but also infrastructure, funding, and interconnection risks.

The case study could recommend several monitoring indicators: technology-adjusted collateral value, technology-adjusted LTV, refinancing-window ratio, tenant or anchor-

customer concentration, contract duration versus hardware economic life, energy-cost sensitivity, power-availability risk, residual-value assumptions, private-credit exposure, bank liquidity lines to SPVs, bank guarantees, servicing obligations, and reputational support channels.

The central lesson would be: nonbank participation is stabilising when the risk is transparent, well matched to investor horizon, and supported by robust disclosure. It becomes more fragile when technological depreciation, maturity mismatch, and bank-nonbank interconnections are opaque.

6. Do the case studies in this report provide actionable insights for financial institutions in their responsible AI adoption?

Yes, the case studies provide useful practical insights. They show that responsible AI adoption can improve risk detection, operational efficiency, fraud prevention, underwriting consistency, customer service, and cyber resilience. They also illustrate the importance of proportionality and governance.

To make the case studies more actionable, I suggest that each case study include a short "implementation checklist" and "risk-control checklist". For example, a case study on AI credit underwriting could identify: the business objective; the AI model or system used; whether it is developed internally or by a third party; the data sources; the materiality classification; the human-oversight design; performance metrics; explainability tools; drift monitoring; complaint or appeal mechanism; fallback process; and periodic reassessment triggers.

For third-party AI use, an actionable case study could include: due-diligence questions for vendors; required disclosures; model-change notification clauses; data-use restrictions; incident-reporting obligations; exit plans; portability arrangements; vendor concentration limits; and contingency testing. It would also be useful to distinguish between ordinary vendor management and critical AI infrastructure dependency.

For AI infrastructure financing and nonbank participation, a case study could include a transaction-monitoring template. This template may include conventional credit indicators, such as DSCR and LTV, but also technology-specific indicators, such as technology-adjusted LTV, refinancing-window ratio, contract-renewal risk, hardware age, energy-cost sensitivity, secondary-market liquidity, sponsor-support assumptions, and bank-nonbank linkages.

Such checklists would help financial institutions translate broad principles into actual governance documents, credit policies, model-risk procedures, third-party-risk procedures, and board reporting.

7. Are the definitions in the glossary clear and aligned with industry sound practices, including recent developments in AI?

The glossary is broadly clear and aligned with current industry usage. I suggest that the final report consider adding or further clarifying several terms that are increasingly relevant to responsible AI adoption by financial institutions.

First, “agentic AI” could be defined not only as AI with autonomy, but also by reference to its ability to plan, use tools, interact with APIs, access systems, and execute multi-step tasks. This would help institutions distinguish ordinary GenAI output generation from AI systems that can take actions in operational environments.

Second, “human oversight” could be clarified to include meaningful ability to understand, challenge, intervene, override, and remediate AI decisions or actions. In high-autonomy use cases, nominal human review may not be sufficient if the human cannot realistically detect or stop harmful behaviour in time.

Third, I suggest adding “AI infrastructure dependency”. This term could refer to reliance on cloud services, data centres, specialised hardware, foundation models, data providers, deployment platforms, and energy or network resources that are necessary for the operation of material AI use cases.

Fourth, I suggest adding “AI third-party concentration” or “foundation-model concentration”. These terms would help institutions and supervisors monitor common reliance on a small number of providers or models.

Fifth, for financial institutions that finance AI infrastructure, terms such as “technology-adjusted collateral value”, “technology-adjusted LTV”, and “refinancing-window risk” could be useful. They would help translate technological obsolescence into credit-risk monitoring language.

8. Are there any comments you would like to make on this report? Please fill in.

I appreciate the FSB’s work on this important consultation. The proposed sound practices provide a timely and practical framework for financial institutions to adopt AI responsibly.

I submit these comments in my personal capacity as a banking practitioner and researcher in credit risk governance, AI infrastructure finance, and bank-nonbank risk migration. The views expressed here are my personal professional views and do not represent the official position of my employer.

My main recommendation is that responsible AI adoption should be understood as a three-layer governance issue.

The first layer is model-level governance: data quality, explainability, performance, bias, validation, human oversight, and monitoring.

The second layer is institution-level governance: board oversight, senior management accountability, risk appetite, documentation, model inventory, policies, controls, and organisational adaptability.

The third layer is ecosystem-level resilience: third-party dependency, cloud and foundation-model concentration, AI infrastructure availability, cyber and ICT resilience, energy constraints, and bank-nonbank financial interconnections.

The third layer is becoming increasingly important as financial institutions rely on external AI infrastructure and as AI infrastructure itself becomes a major object of bank and nonbank

finance. A financial institution may be an AI user, an AI operator, an AI infrastructure lender, a liquidity provider to an SPV, a counterparty to a private credit fund, or an investor in AI-related instruments. These roles should not be considered in isolation.

The objective should not be to slow responsible AI adoption. Rather, it should be to ensure that AI adoption is supported by transparent governance, resilient infrastructure, effective third-party risk management, and early-warning tools that can detect vulnerabilities before they appear in default statistics or operational failures.

I therefore encourage the FSB to consider adding more explicit language on AI infrastructure dependency, nonbank risk migration, and technology-sensitive refinancing risk in the final report.