



FINANCIAL
STABILITY
BOARD

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

Response to Consultation

Sia Partners

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?

The benefits and risks of AI adoption presented in Sections 1.2 and 1.3, respectively, are comprehensive and reflect current industry trends and practices. We suggest that the report could also consider expanding its discussion on:

a) systemic risks associated with reliance on small number of AI foundation model developers and cloud infrastructure providers. While concentration risk is identified in Section 1.3, we suggest that the report address how these correlate. Where many institutions depend on the same providers/models, a single provider outage or shared model could affect multiple firms at the same time;

b) workforce transition risks, including deskilling or loss on institutional expertise as AI takes over routine analysis, the operational risks that can arise on large-scale role changes and;

c) benefit-realisation risk. The benefits of AI adoption for large financial institutions may take years to materialise, while the cost of maintaining the underlying foundation, data pipeline, model monitoring are recurring than one-off. There is a risk that financial Institutions will build business cases on overly optimistic assumptions, resulting in misallocation of investments or under provisioning of ongoing compliance costs.

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

The sound practices are clearly structured and broadly consistent with practices being adopted by leading financial institutions, and the concept of proportionality is particularly valuable in supporting adoption across financial institutions of varying size and complexity.

The report could be further enhanced by providing prioritisation and practical guidance, particularly for extensive illustrative lists such as documentation requirements by distinguishing core expectations from supplementary practices. A tiered implementation guide outlining baseline and enhanced practices for higher-risk AI use cases would also support more consistent and proportionate adoption across financial institutions. FSB could

also provide clarity in the areas of risk-tiering for AI (for example by customer impact, autonomy, materiality etc.), minimum control expectations, and board and senior management report. FSB could also provide mapping or cross-reference to international recognised AI governance framework such as NIST AI Risk Management Framework and ISO/IEC 42001 to facilitate implementation.

With regards to third-party risk management, FSB could provide more clarity on the management of secondary AI features and dependencies. Examples include risk management for embedded AI features incorporated at a later date in an already existing/approved vendor tool or subcontracting of AI services. While the sound practices address data access and quality, greater emphasis could be placed on data usage rights, specifically regarding how financial data from FIs can be used in third-party AI tools.

For greater clarity in the recommended cybersecurity measures, we recommend considering AI through two distinct lenses: (i) internally developed AI tools, which should be governed by security-by-design principles, and (ii) externally sourced AI tools, which should fall within third-party risk management, security hardening and monitoring frameworks.

While the FSB provides a non-exhaustive list of AI-powered cyber and ICT risk management tools [page 46], broader adoption would benefit from a pragmatic use-case-led approach. Financial institutions should prioritize high-value use cases with clear risk reduction and operational benefits, aligned with their cyber and AI maturity, capabilities and risk appetite.

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?

The FSB has adopted a balanced and forward-looking approach, particularly through its explicit consideration of emerging agentic AI risks, including reward hacking and memory poisoning.

While the report recognises that performance management guidance for agentic AI remains a developing area of research, further supplementary guidance in this area would be valuable given the pace of institutional adoption. The proposed human oversight measures, including kill-switches and agent identifiers are practical safeguards that could be highlighted as near-term implementation priorities for financial institutions. Additionally, FSB could distinguish more clearly between the security baseline practices applicable to all AI model, as well as additional controls for high-autonomy, and high-impact AI systems – such as agentic AI.

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

Yes, it does provide a sound basis for accommodating newer types of AI, as it focuses on outcomes and risks rather than specific technologies. We suggest the following enhancements to ensure this flexibility over time:

a) FSB could consider establishing a defined review cadence (e.g., every 18–24 months), particularly given the rapid evolution of topics such as agentic AI;

b) Initiatives such as Project Mind Forge and the Bank of England and /FCA AI Consortium demonstrate good practices and could be further developed or formalized into an ongoing mechanism for monitoring emerging developments and informing future updates.

- 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 included generally provide valuable practical insights. The report already includes a non-bank example through the insurer case study. However, the overall representation of non-bank financial institutions remains limited relative to banks. Given their growing role in the financial system and increasing adoption of AI, additional standalone case studies covering asset managers, payment firms, and financial market infrastructures would broaden the report's relevance across the sector.

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

The outcomes presented in several case studies, including the early identification of non-performing SME loans and fraud prevention metric are valuable and practical insights. The case studies could be further enhanced by providing more details on the governance and control mechanisms that support these outcomes. For example, the relationship management case study [on page 10] highlights efficiency gains but provides limited detail on the oversight processes applied. A more consistent case study format linking outcomes to the relevant sound practices would enhance their value as practical implementation references for financial institutions.

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

The glossary is clear and generally well aligned with current industry usage, for example the detailed treatment of "explainability." For consistency, we would suggest the glossary also capture emerging agentic AI concepts that are used in the body of the report but not currently defined, such as reward hacking and memory poisoning.

Memory poisoning in particular could sit alongside the existing "data/model poisoning" entry, given the report treats agentic memory poisoning as distinct.

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

While the report provides a comprehensive assessment of the evolving role of AI across the financial sector, further insights on areas such as data privacy, consumer protection, and cross-border supervisory coordination could help strengthen the overall analysis and support more consistent implementation across jurisdictions.