

Sound practices for financial institutions' responsible AI adoption

Overview of consultation responses

1. Introduction

The FSB published on 10 June 2026 a consultation report on the Sound Practices for financial institutions' responsible AI adoption¹ and invited public feedback until 22 July 2026. This note provides an overview of the consultation responses received and sets out the main changes to the report in light of the responses received.

Around 160 responses were received from a wide range of stakeholders that included banks, insurers, other nonbank financial institutions (e.g. asset managers, payment service providers, financial market infrastructures), industry associations, third-party technology and service providers, non-profit organisations, academics, public authorities and independent/individual experts. All responses not marked as confidential (i.e. public responses) have been published on the FSB website.²

2. Summary of feedback received

2.1. General comments and cross-cutting themes

Comments received

Overall, there was broad support across respondents for the FSB's report and its approach to identifying sound practices for AI adoption in the financial sector. Respondents generally welcomed the principles-based, proportionate and technology-neutral approach adopted in the report, noting that this approach supports responsible AI adoption by financial institutions with diverse asset sizes, business models, use cases and risk profiles. The FSB's emphasis on proportionality was particularly valued, facilitating financial institutions of different sizes and business profiles (in particular smaller and less-resourced financial institutions) to implement AI responsibly, without undue burden.

¹ See FSB (2026), *Sound Practices for Responsible Adoption of Artificial Intelligence (AI): Consultation report*, June

² Public responses are available on the [FSB website](#).

Most respondents highlighted that the non-prescriptive nature of the sound practices enables financial institutions to integrate the sound practices into their existing governance, operational resilience and risk management structures in a practical manner. This would in turn encourage innovation, support scalability and allow for adaptation as both AI technology and jurisdictional regulatory requirements evolve. A few respondents asked the FSB to consider more prescriptive sound practices.

The importance of international consistency and alignment with established standards and definitions (e.g. those put forth by the National Institution of Standards and Technology (NIST), International Organization for Standardization (ISO), Organisation for Economic Cooperation and Development (OECD)) was also widely noted, with respondents highlighting that consistency of practices as well as standards across jurisdictions helps avoid fragmentation, reduce compliance costs and supports cross-border activity and innovation.

Respondents welcomed the case studies in the report, noting that they help to illustrate how sound practices are implemented in real life to address risks and challenges. Some respondents, however, called for further details to be included in the case studies to aid implementation and the inclusion of a broader range of practical case studies, especially those relevant to smaller financial institutions as well as insurers and other nonbanks. Some also suggested incorporating negative or unsuccessful examples to provide a more balanced perspective.

Most respondents found the terms and definitions in the Glossary (Annex 4 of the report) to be broadly aligned with existing industry conventions. However, many respondents suggested enhancing the definitions of some key terms (e.g. agentic AI, AI (or AI system)) and adding more terms used in the report (e.g. human-in-the-loop, shadow AI) to improve the usability of the sound practices. A few respondents also suggested that the FSB considers establishing a standardised AI-related lexicon or including some terms in its Cyber Lexicon³ in the future so as to facilitate consistency in terminology and practices across jurisdictions.

Several respondents asked for greater clarity on how the sound practices apply when deploying third-party AI. Many respondents expressed concerns about their ability to manage risks from the use of third-party AI models, especially frontier models. They cited challenges with the lack of sufficient information from vendors to conduct proper due diligence for their own use of third-party AI models and expressed concerns with their dependencies on certain key service providers in the third-party AI ecosystem.

Several respondents proposed that the FSB expands the discussion of agentic AI, given the increased interest among financial institutions in adopting this technology and the unique challenges and associated controls it requires. Specific feedback suggested additional details with respect to:

- controls such as pre-execution authorisation, least-privilege access, immutable audit trails and runtime monitoring, with execution governance as a distinct category.
- human oversight, emphasising the need for clear accountability, traceability and real-time intervention capabilities, supported by assigning accountable owners, while

³ See FSB (2018), *Cyber Lexicon*, November

avoiding an excessively prescriptive approach, such as the need for human review for each action.

- business continuity planning and resilience testing tailored to agentic AI, recognising that manual fallback or vendor substitution may not be feasible during disruptions.
- traceability, focusing on actions, inputs and decisions due to the complexity and autonomy of such systems.

Although the focus of the report is on identifying financial institutions' sound practices for their adoption of AI, some respondents suggested that the report also discusses the macroprudential concerns and related financial stability implications from AI use, such as the correlated behaviour of market participants from using the same AI models or as a consequence of third-party provider concentration (including herding behaviour and amplification of procyclicality).

Some respondents emphasised the importance of consumer protection and market conduct risks in AI adoption. They highlighted risks related to explainability, fairness, bias and potential consumer harm, especially in customer-facing applications.

Finally, as the technology is rapidly evolving, many respondents suggested that the FSB continues its public-private sector collaboration, including participation from non-financial entities (e.g. vendors) to nurture shared understanding, increase overall transparency in the use of AI by financial institutions and facilitate greater consistency in approaches.

Changes in response to comments

- The final report retains the principles-based and technology-neutral approach for the sound practices and continues to emphasise that financial institutions can apply the relevant sound practices in a proportionate manner.
- Language was added to the section on risks to highlight that documentation gaps may hinder auditability, accountability and traceability of decision processes, which are important safeguards against consumer harm in consequential, customer-facing AI applications. Modifications were made to Sound Practice 3 (Incorporation of AI risks into risk management framework and effective documentation of AI use) to clarify that effective documentation should support accountability, traceability and the reconstruction of actions, particularly for agentic AI use cases.
- Additional case studies are included, showcasing practical applications from insurers and other nonbanks (e.g. asset managers, payment service providers) to provide a broader perspective. The level of detail in the case studies has been retained from the consultation report as the case studies are not intended to prescribe a specific approach to AI adoption. The case studies remain illustrative of possible real-life outcomes and are not intended to serve as benchmarks for any financial institution.
- Some changes have been made to the Glossary, with efforts directed toward explaining relevant terms used in the report to improve the usability of the sound practices. The FSB may consider including some terms in the context of any future revisions to the Cyber Lexicon.

- The report focuses on sound practices to support responsible AI adoption by individual financial institutions, and thus no changes have been made to address broader systemic risks or macroprudential concerns arising from AI use across the level of the financial sector. The FSB nevertheless recognises the importance of these issues⁴ and will continue to monitor AI-related systemic risks through its regular assessment of vulnerabilities.
- Modifications were made to Sound Practice 12 on third-party AI risk management (see details below), especially by adding language on how individual financial institutions could manage risks from third-party AI providers. However, the report's focus is on the use of third-party AI by individual financial institutions and not on broader topics related to macroprudential or systemic challenges associated with AI. Therefore, while these topics are important for the FSB and its members to consider, the final report does not address them.
- The final report continues to emphasise that the sound practices are applicable regardless of which AI technologies financial institutions choose to adopt. That said, in view of the feedback received, the final report has been enhanced in several areas relating to agentic AI, specifically the discussion of risks associated with agentic AI has been updated and the sound practices now include specific considerations for highly autonomous AI. These additional considerations on agentic AI are not fully consulted on but are distilled from the public consultation responses. The FSB notes that these sound practices specific to agentic AI and multi-agent systems are still emerging and may continue to evolve as understanding of these systems and their associated risks deepens. A more thorough treatment may be merited in the future. A new case study on agentic AI has also been added.

2.2. Benefits and Risks of AI Adoption

Comments received

Respondents generally agreed with the benefits and risks highlighted in the report. Some technical suggestions were provided, including considerations around the role of AI in the workforce (both benefits and risks) and additional detail on risks related to insufficient documentation, third-party solutions, generative AI (GenAI) and agentic AI.

Some respondents emphasised that the risks of not adopting AI, such as falling behind in cyber resilience or fraud detection, can be material, especially as malicious actors increasingly exploit advanced technologies.

⁴ For example, see FSB (2024), *The Financial Stability Implications of Artificial Intelligence*, November

Changes in response to comments

- The section on benefits (Section 1.1) now acknowledges AI as an augmentative technology and workforce multiplier and includes additional case studies showing how nonbank financial institutions have benefited from AI use.
- The section on risks and implementation challenges (Section 1.2) now provides more detail on workforce deskilling; challenges related to documentation, auditability and traceability; embedded AI in third-party solutions; and risks associated with GenAI and agentic AI.

2.3. Organisation-wide AI governance: Sound Practices 1-4

Comments received

Respondents generally welcomed Sound Practices 1-4 on organisation-wide AI governance. Some called for a clearer distinction between board oversight and senior management's operational responsibilities, with accountability proportionate to the scale, materiality and risk of each financial institution's AI adoption.

Respondents emphasised the importance of allowing for flexibility to embed AI risk management within existing governance structures and risk management frameworks, provided that AI risks are comprehensively identified, managed and monitored. They also highlighted the value of cross-functional collaboration to allow all relevant aspects of AI risks to be understood and appropriately addressed. Several respondents stressed the importance of integrating AI controls across all three lines of defence, an example being the hub-and-spoke governance models that combine central frameworks with business-line implementation.

Some respondents further stressed that financial institutions should regularly reassess their AI governance, risk management practices and capabilities as technology evolves. Continuous learning and organisational flexibility are essential to remain resilient to changes in the technological, regulatory and threat landscape.

Changes in response to comments

- The final report clarifies the respective roles of the board and senior management. It explains how small financial institutions can adapt the governance and risk management-related sound practices by leveraging existing roles and external expertise. It also provides further guidance on accountability, regular reassessment of AI strategies and the retention of human expertise and critical judgement as AI use expands.
- The final report further elaborates on the importance of cross-functional coordination. It highlights formal communication channels and regular forums as mechanisms for incorporating diverse perspectives into sound decision-making and risk management, particularly for highly autonomous AI systems.

2.4. AI lifecycle management: Sound Practices 5-12

Comments received

There was broad support for the approach taken towards sound practices across the AI lifecycle, where respondents welcomed the careful consideration of sound practices at different stages of the lifecycle. As noted earlier, several respondents raised the need to expand the discussion on agentic AI in the relevant stages. Additionally, respondents provided technical feedback on certain sound practices, particularly concerning third-party AI risk management.

Respondents generally supported robust, risk-based selection processes for AI systems and third-party service providers. Many respondents called for AI-specific considerations to be incorporated into existing third-party risk management frameworks, including the degree of transparency provided by and substitutability of service providers. They called for the development of standardised assurance and transparency documentation to streamline their due diligence and monitoring of third-party AI services and service providers. For example, recognised industry certifications and harmonised technical documentation, such as system cards. They also suggested the FSB consider greater guidance for smaller financial institutions with limited leverage against larger third-party service providers.

Respondents agreed with the importance of effective data management and governance. Some respondents, however, cautioned against creating AI-specific data management and governance practices, and suggested that data governance practices should remain proportionate to risk and materiality of the AI use case. Some respondents also highlighted that many financial institutions, particularly smaller ones, faced unique challenges in implementing data management and governance practices due to their reliance on third-party service providers that provided limited transparency and control over underlying data and data governance processes (e.g. enforcing data standards).

Respondents widely supported the report's risk-based approach to explainability, with some advocating that more material use cases, including customer-facing applications and regulatory reporting, warrant more explainable AI, while less explainable AI may be acceptable for low materiality and/or low-risk internal tools. Many also emphasised the importance of outcome-based monitoring, real-time oversight and continuous behavioural assurance for GenAI and agentic AI.

Respondents generally agreed that effective human oversight must involve genuine authority and operational feasibility that are evidenced, not just nominal review. For agentic AI, respondents emphasised the need for clear accountability, traceability and runtime controls, as well as business continuity that recognises that fallback or vendor substitution, although this may not always be feasible.

Respondents identified cyber and ICT risk management as a key priority, particularly as threats from advanced AI models evolve. Some suggested placing cyber risk management earlier in the sound practices and adopting practical, risk-based approaches to patching and vulnerability management.

Changes in response to comments

- A number of comments raised related to areas already addressed in the report. Where relevant, the final report has been updated to provide additional clarity or emphasis on these areas, which include materiality and risk assessment, selection, data governance and explainability. The ordering of the sound practices does not indicate prioritisation and some practices, also cut across multiple stages.
- The description of Sound Practice 8 (Explainability and transparency) already highlights that AI use with potentially higher impact on customers warrant greater explainability and has been refined to further emphasise the importance of transparency in higher-risk and higher-material use cases and the role transparency can play in a financial institution's engagement with supervisory authorities.
- The description of Sound Practice 9 (Performance management) has been expanded to include additional considerations on independent validation. It also highlights the role of subject matter experts in assessing GenAI performance and the importance of real-time monitoring, especially for AI capable of acting autonomously.
- The final report updates the description of Sound Practice 10 (Human oversight) with additional considerations on how financial institutions can overcome the challenges posed by GenAI and agentic AI.
- For agentic AI, the final report includes additional considerations as mentioned earlier.
- The final report broadens the description of Sound Practice 12 (Third-party AI risk management) to address risks associated with embedded AI in third-party products and sub-outsourced activities within the AI supply ecosystem. It also introduces considerations for managing concentration risk, such as obtaining more granular and timely assurances from third-party AI providers, ensuring transparency, and implementing robust business continuity and exit strategies for critical AI services. Additionally, the final report encourages sector-wide collaboration on standardising contractual terms and assurance tools, conducting joint scenario testing and enhancing oversight of supply chain dependencies.