FSB FINANCIAL STABILITY BOARD

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

Response to Consultation

Asociación Española de Banca (AEB)

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?**

   We broadly agree with the benefits and risks of AI adoption described in the report. We particularly welcome the recognition that not adopting AI may itself create risks, as this provides a more balanced perspective on AI adoption. We do not identify any substantive benefits or risks that are missing from the report.

   However, the final report should continue to preserve a principles-based, proportionate and technology-neutral approach. In particular, risk descriptions should not be transposed into prescriptive supervisory expectations. For example, references to explainability should not be read as a general constraint on the use of less explainable models where performance, risk and business objectives justify their use and appropriate controls are in place. Similarly, discussions on cyber and ICT risk information-sharing should recognise that cross-border sharing of AI incidents, near misses and emerging threats depends on applicable national legal and regulatory frameworks.

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

   We consider the sound practices to be broadly comprehensive and clear. However, some parts could be interpreted as more prescriptive than intended. We encourage the FSB to preserve the report's principles-based, proportionate and technology-neutral approach, particularly in areas such as documentation, data governance, testing and human oversight.

   For example, detailed documentation expectations under Sound Practice 3, or specific cyber frameworks such as OWASP ML Security Top 10, MITRE ATLAS and the NIST adversarial machine learning framework under Sound Practice 11, should be clearly framed as examples rather than expected approaches. Similarly, case studies should be presented as illustrative examples, rather than preferred implementation models.

   This would be consistent with the FSB's stated intention to provide a non-binding menu of sound practices, rather than a prescriptive international standard.

Sound practices could also benefit from more non-binding practical examples on the application of proportionality, materiality assessment, board reporting, GenAI testing, human oversight and third-party contractual expectations. Such examples would help support consistent interpretation across jurisdictions, while preserving the report's principles-based and technology-neutral approach.

For example, it would be useful to exemplify concepts such as "materiality", "high materiality" and "high risk" - which are mentioned on section 1.1- in a manner that is aligned with the report's risk-based and proportionate approach.

Similarly, references in sound practice 1 to board and senior management oversight should be understood proportionately. They should not be read as implying granular board-level monitoring of all AI applications, or as creating fixed expectations on management information, reporting frequency or escalation to the board or relevant committees. These arrangements should remain proportionate to the materiality, complexity and risk profile of the relevant AI use cases and should be determined by institutions within their broader governance and risk management frameworks.

The same principle applies to AI literacy expectations for boards and senior management. Appropriate awareness of AI-related opportunities, risks and controls is important. However, these expectations should not be interpreted as requiring all board members, or a specific board committee, to have professional AI technical expertise. Institutions should retain flexibility to allocate expertise, support and challenge functions within their governance structures.

It would also be helpful to further clarify the paragraph regarding "preventing automation bias (i.e. overreliance) by creating incentive alignment for thorough, thoughtful review, such as by incorporating performance metrics that reward quality of oversight decisions".

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?

We consider that the sound practices generally strike an appropriate balance. They remain broadly applicable to all forms of AI while recognising that certain risks and controls may be particularly relevant for GenAI and agentic AI.

Given the relative novelty of agentic AI, we agree that this area will require further monitoring and may evolve over time. The final report should therefore preserve a flexible and technology-neutral approach, recognising that implementation approaches will continue to evolve as industry practices mature.

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

We consider that the sound practices are broadly flexible enough to accommodate newer types of AI and responsible adoption over time. However, this flexibility should be reinforced in the final report by ensuring that examples, case studies and lists of controls are clearly framed as illustrative and optional.

In particular, wording that may be interpreted as endorsing specific implementation approaches could be softened. For example, expressions such as "typically includes" could be replaced with more flexible language such as "could include actions such as". Similarly, examples relating to documentation, data governance, testing, or human oversight should not be interpreted as preferred implementation models, but rather as possible approaches depending on the institution's size, complexity and AI maturity.

Finally, the sound practices should also acknowledge that financial institutions' ability to implement certain governance and risk management practices may depend on the information, safeguards and support made available by third-party AI providers, particularly in the case of general-purpose AI models and systems.

Stronger recognition of the importance of provider transparency and cooperation could help establish a common industry expectation around the information that providers should make available to support more effective third-party AI risk management and more consistent industry practices.

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 generally useful in illustrating how different financial institutions can benefit from responsible AI adoption. However, they should be clearly framed as illustrative examples rather than preferred or expected implementation approaches.

We would also encourage the FSB to consider including additional examples reflecting practical approaches already being developed by financial institutions, including responsible use of AI for internal knowledge management, process automation, customer interfaces, fraud detection, and model monitoring and validation. Examples could also illustrate the application of AI in governance and risk management, such as risk-based tiering of AI use cases, ongoing performance monitoring, independent model validation, and governance of third-party AI services. These examples would further demonstrate how responsible AI adoption can be implemented proportionately across institutions with different business models, AI maturity and risk profiles.

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

The case studies provide useful and actionable insights for financial institutions. However, the final report should make clear that they are illustrative examples rather than preferred or expected implementation approaches.

This is important to preserve proportionality and avoid creating de facto supervisory expectations around specific tools, governance models or control arrangements. For example, case studies describing dedicated AI inventories, specific documentation practices or governance structures may be appropriate for some institutions but should not be interpreted as the FSB's preferred implementation model.

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

   As the consultation paper refers to high risk use cases in some parts, a definition of what is understood as a "high-risk use case" should be provided.

   We suggest the following definition:

   "High risk AI use case: An AI use case that meets the criteria for high-risk classification under the applicable national or regional AI governance and risk management framework. This may include AI use cases that could have a significant impact on the health, safety, or citizens fundamental rights.

   Other concepts commonly used by financial institutions in their AI governance frameworks could also be included, such as "AI use case", "AI model" and "AI inventory". In particular, the report could refer to AI inventories as centralised records of AI use cases and associated AI models, systems or components, supported by internal documentation explaining how the inventory should be completed consistently across the organisation.

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