FINANCIAL STABILITY BOARD

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

Response to Consultation

Hamlaw Investment Concepts LLC

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?**

I fundamentally agree with the FSB's identification of AI benefits and risks for financial institutions. The report accurately captures operational efficiency gains, enhanced risk modeling, and improved customer experience as key benefits.

Nevertheless, I would highlight two underexplored risk areas:

First, the compounding risk of AI opacity at the institutional level. While the report addresses explainability at the use-case level (Sound Practice 8- Explainability and Transparency), it does not sufficiently address the systemic risk created when multiple institutions simultaneously deploy similarly opaque AI models trained on overlapping datasets. This creates correlated, system-wide blind spots that no single institution can detect independently.

Second, workforce transition risk. The rapid displacement of roles traditionally filled by human experts such as compliance officers, credit analysts and fraud investigators — introduces institutional knowledge loss that may not surface until a crisis reveals gaps that AI systems cannot navigate without human judgment.

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

All 12 FSB sound practices form a strong foundational framework. FSB Sound Practices 1-4 on organization-wide governance structures are particularly well-structured and replicate real-world best practices.

My recommendation: FSB Sound Practice 6 (Selection) would benefit from explicit guidance on the governance of foundation model selection — specifically when institutions choose to build on third-party large language models (LLMs). The current framing was designed with traditional ML models in mind, and the rapid shift to generative AI and agentic AI systems requires updated specificity here.

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?**

   This could be the most important question in the consultation. To be frank, we aren't there yet, but the framework is showing a lot of potential for the future.

   Agentic AI — systems that take autonomous, multi-step actions with real-world consequences — introduces qualitatively different risks than predictive AI. These include goal misalignment, cascading autonomous decisions, and the challenge of defining meaningful "human oversight" when the system performs much faster than a human review is possible.

   I recommend the FSB consider a dedicated supplementary guidance document specifically addressing agentic AI governance, including:

   - Clear definitions distinguishing agentic AI from automated decision-making systems

   - Minimum latency requirements for human review checkpoints in high-stakes autonomous workflows

   - Mandatory incident reporting protocols for autonomous AI actions that materially deviate from intended outcomes

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

   The principle-based approach (rather than prescriptive rules) is inherently the right design choice for a fast-moving technology landscape. Its flexibility is a core strength.

   A structural improvement I recommend is for FSB to establish a formal annual review cycle for the sound practices, with public stakeholders' input, to ensure the framework keeps pace with AI capability growth. The strategic rationale is clear, what constitutes "sound practice" for AI governance in 2026 will look materially different by 2028.

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 useful, although they skewed towards large, well-resourced institutions. Community banks, credit unions, and smaller non-bank financial intermediaries face meaningfully different resource constraints and risk profiles.

   I recommend the FSB solicit and publish a dedicated set of case studies from smaller and mid-sized institutions, particularly from emerging markets where AI adoption in finance is growing rapidly with less regulatory infrastructure in place.

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

While the case studies are illustrative, they could benefit from being more prescriptive. To maximize utility, practitioners, require deeper insight not just into the action taken, but into the specific governance decisions made, initial points of failure and subsequent institutional adaptations. Consequently, I recommend future case studies adapt standardized, structured format comprising five key pillars: the strategic objective, the tactical approach, the early implementation/challenges, governance adjustments, and outcome metrix.

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

Overall, the glossary is generally clear and consistent with industry standards. However, I have flagged two specific definitions for reconsideration to ensure complete alignment and precision.

1. AI systems ---the current definition of the AI System, may not clearly/explicitly encompass the multi-model pipelines or orchestration layers common in modern enterprise AI deployments. Clarifying this definition is critical to prevent potential regulatory arbitrage, wherein an institution might claim that individual components of a system, when viewed in isolation, fall outside the scope of governance.

2. Human oversight— the term Human oversight requires greater operational specificity. Oversight that is technically available but practically impossible due to system speed or complexity, does not constitute meaningful control. To address this, the Financial Stability board (FSB) should establish minimum standards defining what constitutes effective human oversight across different AI risk tiers.

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

The FSB's initiative to establish global sound practices for AI adoption in finance is timely and vital. Responsible AI governance is not a constraint on innovation; rather, it is the foundation that makes sustainable innovation possible. I highly commend the FSB for leading this critical global conversation and look forward to seeing the final guidance reflect the rapidly evolving realities of AI deployment.