FSB FINANCIAL STABILITY BOARD

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

Response to Consultation

Bank Ready AI

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, the introduction section is broadly thorough on business strategy, business model benefits, and interestingly mentions specific AI-prone processes like credit risk management. While benefits are broadly described, the risks are broadly described too. However, the coverage of recursive governance is non-existent. There is coverage of AI overseeing AI (sound practice 10), but this isn't the same as recursive governance, where management governance expands because an aspect of governance was outsourced to technology.

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

   It is comprehensive and clear, however there could be more around what evidence outcomes should be produced to support AI governance work completed by institutions (the scale to which can tie into the proportionality argument).

   Also, there could be more discussion of where the 12 practices apply when AI is vendor-provided vs built/fine-tuned in-house. Where AI is vendor-provided, for example at the community institution level, sound practices 5-10 are implemented through sound practice 12 mechanisms (due diligence, contractual disclosure, and compensating controls). So, this level of direction could be made clear.

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 report differentiates across traditional ML, GenAI, and agentic AI appropriately. Section 3.6 and Annexes 2 and 3 are thorough. Per our response to Q2, we believe there is an opportunity to better serve the AI adoption modes: vendor-provided vs. in-house built/fine-tuned.

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

What would make the practices more flexible would be to ground them in operational control points. For example, rather than describing features of the technology, the sound practices could describe technology features from the perspectives of autonomy (or the autonomy level); how deterministic the technology's output is (the output determinism); how much access the technology has (the access scope); and how reversible the technology's actions are (reversibility of actions).

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.**

We don't see institutional coverage for community financial institutions, yet section 1.3's own risks – service provider concentration, and homogeneous behaviors and correlated outcomes – are directly in play at this institutional level (e.g., thousands of small institutions relying on the same handful of foundation models and vendors). Here are two community finance AI adoption case studies we authored: https://www.a2z.financial/research.

On the request for nonbank case studies, we defer to other respondents.

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

Partly. The report's case studies across sections 1-3 include descriptive practical AI use-case examples. But they don't provide actionable insights on what documentation or regulatory evidence documentation institutions should maintain. These case studies could be held to the same standard as those in Annex 2 and 3. The annex 2 and 3 worked examples actually discuss the governance reasoning applied to each AI use-case.

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

Three suggestions:

1. Recommend replacing the current ‘open-source models’ entry with 3 separate entries for open-source, open-weight (permissively licensed), and open-weight (restricted commercial license): the report’s own third-party case study in section 3.9 mentions there’s an open source “spectrum”, but the glossary doesn’t expand on this. The glossary additions could be based on due-diligence and license terms: permissive licenses include those released under Apache 2.0 and MIT.

2. Add AI use-case: since the report is based around different use-cases, make it clear that this refers to a business application of the technology.

3. Add proportionality and materiality, not just criticality: since one of the key suggestions in the report is to use proportionality when applying the practices, and materiality is a big driver, both can have glossary items, like “criticality.”

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

It's a good report. Very detailed. A few suggestions that we believe make the report more robust are included in responses above.