

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

Response to Consultation

Commencement Bank

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 generally agree with the benefits and risks described in the report, as they accurately reflect the opportunities and challenges associated with AI adoption. The report appropriately identifies efficiency gains, improved risk management, and enhanced customer service as key benefits, all of which are relevant in a community bank environment where resource constraints make productivity gains particularly valuable. The risks identified, including data quality, explainability, cybersecurity, and third-party dependencies, are also well aligned with our experience. Additional considerations include the importance of AI as a workforce multiplier and the reliance on vendor-delivered AI capabilities, along with risks related to vendor opacity, shadow AI usage, and evolving regulatory expectations.

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

The sound practices are comprehensive and provide a strong foundation for responsible AI adoption. They cover governance, lifecycle management, and key control areas such as data governance, explainability, human oversight, and third-party risk. The structure aligns well with existing risk management frameworks; however, while the principles are clear, they would benefit from additional practical guidance on implementation. Smaller institutions would gain value from examples of proportional controls and simplified approaches that meet regulatory expectations without requiring enterprise-level complexity.

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 report strikes an appropriate balance between managing risks across all AI types and addressing emerging risks associated with generative and agentic AI. It recognizes that many traditional risks are amplified by greater complexity and autonomy while also highlighting new concerns such as hallucinated outputs and autonomous decision-making. This balanced approach is appropriate for community banks that primarily consume AI through vendors. Additional clarification differentiating internal AI development and vendor-provided AI would further enhance applicability.

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

The sound practices are sufficiently flexible to accommodate evolving AI technologies. The non-prescriptive design and emphasis on proportionality and continuous reassessment allow institutions to adapt governance practices as their use of AI matures. This flexibility is particularly important for community banks, where adoption is incremental and driven by vendor capabilities. The framework supports long-term adaptability while maintaining alignment with regulatory expectations.

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 provide useful examples of AI benefits but are primarily focused on large institutions, limiting their applicability to smaller banks. Additional case studies reflecting community bank use cases, such as AI-assisted document processing, compliance support, and vendor-based fraud detection, would improve relevance. Including nonbank examples, such as fintech applications in onboarding and fraud prevention, would further broaden the perspective.

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

The case studies provide helpful conceptual insights and demonstrate business value and governance considerations; however, they are less actionable from an implementation standpoint. They do not provide sufficient detail on how governance structures are established or how controls can be scaled appropriately. Additional practical guidance, such as implementation examples, templates, and minimum control expectations, would enhance their usefulness for smaller institutions.

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

The glossary definitions are clear and aligned with industry standards, including emerging AI concepts such as generative and agentic AI. They provide a strong technical foundation for understanding the framework. Minor enhancements could include simplified explanations for non-technical stakeholders and examples relevant to financial services to improve usability.

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

None