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

Response to Consultation

Cornerstone Quality Partners, 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?**

   The FSB's 12 sound practices constitute the most operationally grounded international guidance on AI governance that financial institutions have received to date. Practices 1 through 4 correctly place governance accountability at the board and senior management level before any AI system reaches deployment. Practices 5 through 10 establish a lifecycle logic that maps to how AI risk actually accumulates: at selection, at data intake, at deployment, and during ongoing monitoring. The inclusion of specific human oversight modalities in Practice 10, ranging from human-in-the-loop to kill switch mechanisms, reflects direct engagement with how financial institutions are structuring their AI operations today.

   The practices are also appropriately non-prescriptive. Proportionality is built into the structure. A community bank running a single ML-based fraud detection model does not require the same governance architecture as a globally systemic institution deploying agentic AI across trading, compliance, and customer service simultaneously. The framework respects that distinction.

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

   Three gaps between what the practices state and what implementation requires merit specific attention in the final report. Each is observable in the current US market, where the Federal Reserve, OCC, and FDIC issued updated model risk management guidance in April 2026 (SR 26-2) that explicitly excluded generative AI and agentic AI from scope. The result is a documented governance vacuum: 72 percent of US banks lack formal AI model kill-switch protocols and failure reporting procedures (American Banker, 2026), and only 18 percent of banking leaders express full confidence that their AI controls would pass an independent review within 90 days (Grant Thornton AI Impact Survey, 2026). These are not hypothetical risks. They are current conditions in a major financial market that the FSB's practices are intended to address.

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?**

Gap 1: Practice 9 (Performance Management) Requires Continuous Behavioral Assurance, Not Point-in-Time Testing. The current draft addresses developmental testing and ongoing monitoring but does not adequately distinguish the evidence standard required for deterministic AI systems from the evidence standard required for probabilistic AI systems. A language model deployed via a third-party cloud provider can change its output distribution when the provider updates the underlying model, without any code change, without any notification, and without any version increment that would trigger a standard re-validation cycle. Practice 9 should explicitly require that financial institutions operating GenAI or agentic AI implement continuous behavioral assurance processes, not as an optional compensating control, but as the default evidence standard for this class of system.

Gap 2: Practice 10 (Human Oversight) Does Not Define Who Is Accountable When Third-Party Models Update. Kill switch mechanisms are named and defined. What is not addressed is the accountability question that arises in the most common enterprise AI deployment pattern: the institution is a downstream consumer of a third-party model, the third-party provider deploys model updates on a schedule outside the institution's control, and no one inside the institution has been designated responsible for assessing whether post-update behavior remains within approved parameters. This is the precise governance gap that banks in the United States are currently managing by disabling AI features entirely rather than governing them. Practice 10 should specify that accountability for post-update behavioral review must be assigned to a named individual or function within the institution, regardless of where the model is hosted.

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

   Each of the gaps above shares a common root: the distance between a governance document and a governed system. The practices are designed to be documented by institutions. What is not yet specified is the independent evidence standard that confirms whether documentation reflects operational reality.

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.**

   Yes.

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

   The Practices Do Not Address the Regulatory Sequencing Problem for GenAI and Agentic AI. Practices 1 through 4 correctly require AI risk management frameworks and documentation. What they do not address is the situation where no supervisory playbook exists for the AI system a financial institution is currently operating. When examiners arrive at a US bank running GenAI tools in AML or lending workflows under SR 26-2's exclusion, the examiners are applying judgment calls, not written rules. Institutions without self-defined governance frameworks are the most exposed. The FSB's final report would benefit from explicit guidance on how financial institutions should document their governance rationale when no jurisdiction-specific regulatory standard yet covers their AI use case.

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

Yes.

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

An AI Readiness Assessment conducted by an independent reviewer, structured against the FSB's 12 practices, closes this distance. The assessment does not replace internal governance. It produces the external evidence that internal governance is functioning as documented: that risk assessments are current, that kill-switch protocols are tested, that behavioral drift is monitored, and that human oversight assignments are maintained and not nominal. The FSB's final report would strengthen institutional uptake of the practices by recommending that independent readiness assessments be conducted proportionate to the materiality of an institution's AI use cases.

The financial services compliance community will follow what the FSB recommends. The gap between the current consultation draft and the operational reality of AI governance in 2026 is one the FSB has the authority to close. The 16-day window before this consultation closes is the opportunity to do so.