FSB FINANCIAL STABILITY BOARD

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

Response to Consultation

Yeh and Johnson Consulting, 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?**

   Yeh and Johnson Consulting LLC agrees with the benefits and risks of AI adoption described in the report. The FSB's identification of operational risk, third-party dependency, model opacity, and the governance challenges posed by agentic AI reflects the landscape as practitioners encounter it.

   We note one risk category that is present in the report's framing but not named explicitly as a distinct risk: the risk of an absent governance record.

   Financial institutions may have governance policies, risk frameworks, and human oversight processes in place and still be unable to produce, after the fact, a contemporaneous record of who authorized a specific AI deployment, under what authority, with what controls active, and at what date and time. This is not a failure of policy. It is a failure of record — and it creates a distinct class of risk.

   When an AI-assisted decision is challenged — by a regulator, in litigation, or in an internal audit — the institution's exposure is not only whether it had a governance framework. It is whether it can demonstrate that the framework was followed at the moment the consequential decision was made or authorized. Where no execution-time record exists, governance decisions become subject to re-litigation: the absence of evidence creates the conditions for contested accountability.

   We recommend the FSB name this explicitly as a distinct risk — the execution-time governance record gap — alongside the operational, model, and third-party risks already identified. Naming it as a risk category would strengthen the motivation for Sound Practice 3's documentation requirements and make explicit what institutions should be able to demonstrate, not only what policies they should maintain.

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

   The proposed sound practices are well-structured and the three-pillar architecture — organisation-wide governance, lifecycle management, and third-party risk — reflects the

actual governance challenge as it presents in practice. The proportionality principle is appropriately calibrated, and the attention to agentic AI throughout is both timely and substantive.

We offer one observation about a gap that is present across multiple practices but not addressed explicitly: the difference between having a governance framework and having a governance record.

The proposed practices describe what institutions should do — maintain accountability, document AI use cases, implement human oversight, conduct testing. They are appropriately non-prescriptive about how. However, the practices do not address the question of what evidence of governance looks like at execution time — what an institution should be able to produce, to whom, and under what circumstances, to demonstrate that its governance practices were actually followed at the moment a consequential AI decision or deployment was authorized.

This gap has practical consequences. In our practitioner experience, financial institutions that have governance policies in place frequently cannot produce an execution-time record of who authorized a specific AI deployment, under what authority, with what controls active, and at what date and time. The policy exists. The governance act — the moment at which a human exercised authority over a consequential AI decision — cannot be demonstrated after the fact.

Sound Practice 3 states, in Footnote 37, that "proactive record-keeping from inception promotes effective AI risk management." This is correct. We suggest the FSB consider making explicit that the documentation requirements of Sound Practice 3 are intended to produce artifacts capable of demonstrating, at the level of individual decisions and deployments, that governance was exercised — not merely that governance policies exist. The distinction between a governance policy and a governance record is the difference between an institution that has a framework and an institution that can prove it followed one.

We recommend the FSB consider explicitly naming the execution-time governance record as a distinct governance output — not a new practice, but a clarification of what Sound Practice 3's documentation requirements are intended to produce in practice.

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 proposed practices address agentic AI thoughtfully, particularly in Sound Practices 10, 11, and 12. The requirements in SP10 — assigning individual identifiers to AI agents, documenting prohibited actions and required human approval checkpoints, maintaining audit trails of agent transactions, monitoring and documenting the process of AI agents including key intermediate steps — reflect the actual governance challenge that agentic systems present in production environments.

One observation: the practices describe oversight mechanisms — human-in-the-loop, human-on-the-loop, kill switches, contestability — but do not address the documentation of authority that precedes those mechanisms. Specifically: who designated which oversight

form for a given AI deployment, under what authority, and what record exists of that designation?

As agentic AI systems become more prevalent, the question of who authorized an agent to act — not only whether a human was technically on the loop — will become the central governance and liability question. An institution may have a human-on-the-loop process in place. Whether that process was applied, by whom, and under what mandate, is a separate question that the practices do not yet require to be recorded.

We recommend the FSB consider adding explicit guidance that governance documentation for agentic AI include a record of:

- who authorized the agent's permitted action scope;
- under what authority that authorization was granted;
- what conditions and controls were active at the time of authorization; and
- what record exists to demonstrate all three at the time of any subsequent challenge, audit, or review.

This is not a new requirement — it is implicit in SP3's documentation list and SP10's oversight requirements. Making it explicit for agentic AI would materially strengthen the practices and address the accountability gap that SP12 correctly identifies: "financial institutions remain accountable for AI outcomes even when the AI was not built in-house." That principle applies with equal force to agentic systems operating within an institution's own infrastructure.

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

The principles-based design of the sound practices supports flexibility. By focusing on what institutions should achieve — accountability, documentation, human oversight, third-party risk management — rather than prescribing how, the practices can accommodate AI forms that do not yet exist at scale. Sound Practice 4's emphasis on organisational adaptability and periodic governance review reinforces this.

We offer one observation about where flexibility has a structural limit: the governance record.

As AI types evolve, the governance practices can adapt. The governance record — the execution-time artifact documenting that governance was exercised — must adapt with them. The documentation requirements of Sound Practice 3 were designed primarily with traditional and GenAI use cases in mind. For agentic AI systems, the relevant record looks materially different: it must capture not only that a system was deployed, but that its action scope was authorized, by whom, under what authority, and with what controls active at the time of authorization — not reconstructed after the fact.

The practices would be strengthened by including guidance on how institutions should adapt their governance documentation requirements as AI capabilities evolve — specifically, what elements of the execution-time record should be revisited when a new AI form is adopted.

One thing that does not change as AI evolves: the core accountability question. Who authorized this AI to act, under what authority, and what record exists to demonstrate it? That question is durable across every AI type the FSB is likely to encounter.

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 effectively illustrate the breadth of AI adoption across institution types and the diversity of governance approaches being developed. We particularly note Case Study 8, which describes the spectrum from institutions maintaining centralised AI inventories with full lifecycle documentation to institutions relying on spreadsheets — and the emergence, in mature organisations, of "agent certification," where an agent is reviewed and approved for use within defined boundaries with its certified status and permitted scope recorded as inventory attributes.

We note one gap across the case studies: while they describe what institutions are doing with AI, they do not illustrate what the governance record looks like. The case studies describe governance processes and institutional choices. They would be more actionable if at least one showed an institution demonstrating, rather than describing, its governance record — specifically, what an execution-time governance artifact contains and what it enables an institution to show.

We offer the following practitioner case study for the FSB's consideration:

Case Study: Execution-Time AI Governance Documentation — Yeh and Johnson Consulting LLC

Context: Yeh and Johnson Consulting LLC developed EDMUND™ — an AI decision governance methodology — in response to the governance gap that the FSB's case studies describe: organisations deploying AI in consequential decision contexts that have governance policies but cannot produce execution-time evidence of governance.

The gap addressed: When an organisation deploys an AI system in a consequential decision context — underwriting, claims handling, fraud detection, credit decisioning — the governance question is not only "do we have a policy" but "what is the execution-time record of who authorized this AI to make this class of decision, under what authority, with what controls active, at what date and time?" This record is what Sound Practice 3 calls for, what Sound Practice 10's responsibility matrices point toward, and what Sound Practice 8 identifies as the precondition for accountability. Most organisations cannot produce it.

What the governance record contains: A decision receipt produced by EDMUND™ captures: a decision statement; the decision type and risk tier (AUTOMATE / AUGMENT / HUMAN-ONLY); the key risks identified; the required actions taken; an authority reality

check documenting who holds formal authority, informal authority, and veto power over the decision; a re-litigation prevention clause recording what evidence exists to prevent the decision from being relitigated without new information; monitoring triggers; and the originator designation (HUMAN / AGENT-AUGMENT / AGENT-AUTOMATE). Each receipt is dated, versioned, and logged to a governed ledger.

Mapping to the FSB's proposed practices:

• Sound Practice 3 (documentation): The receipt is the documentation artifact — capturing use case description, accountability allocation, lifecycle status, and validation findings in a single governed record.

• Sound Practice 10 (human oversight): The receipt captures which oversight form was designated, who designated it, and what authority that designation rested on — the record that distinguishes meaningful oversight from tick-the-box compliance.

• Sound Practice 8 (explainability enabling accountability): The receipt is not an explainability tool; it is the accountability record that explainability is intended to enable. Where explainability is limited, the receipt documents the compensating controls applied and by whom.

Practitioner observation: The FSB's Case Study 8 describes the emergence of "agent certification" in mature organisations — reviewing and approving agents for use within defined boundaries, with certified status and permitted scope recorded as inventory attributes. The EDMUND™ architecture produces this class of artifact: an execution-time receipt documenting authorization, scope, risk tier, and accountability at the moment of governance, not retrospectively. This case study offers evidence that the gap the FSB has identified is real, that practitioners are actively building toward it, and that the proposed practices would benefit from more explicit guidance on what the execution-time record should contain.

We also note that the current case studies are weighted toward large, internationally active institutions. A practitioner case study from a mid-sized institution — for example, an insurance carrier, regional bank, or specialty insurer — navigating AI governance with proportionate resources would make the practices significantly more accessible to the institutions that are currently underserved by governance frameworks calibrated to G-SIBs.

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

The case studies provide valuable context and illustrate the diversity of approaches institutions are taking. We offer two observations to increase their practical utility.

First, the current case studies describe governance processes. They would be more actionable if they illustrated what the governance artifacts those processes produce look like — what an institution can show to a regulator, an auditor, or a court to demonstrate that governance was exercised. The FSB's own language in Sound Practice 3 states that proactive record-keeping from inception promotes effective AI risk management. The case

studies would reinforce this principle more concretely if at least one showed the record, not only the process that was supposed to create it.

Second, several case studies describe institutions that have deployed AI in consequential contexts — credit underwriting, claims handling, fraud detection — with documented governance processes and internal review steps. These cases are instructive. They would be more actionable if they included a description of what an institution in that position would produce if asked by an examiner: "Show us the record of who authorized this AI to make this class of decision, when, and under what controls." That question is already being asked in safety-and-soundness examinations in several jurisdictions. The case studies do not yet prepare institutions to answer it.

We recommend the FSB consider adding to at least one case study a description of the governance artifact — the execution-time record — that the governance process is intended to produce, and what that record should be capable of demonstrating to external reviewers.

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

The glossary definitions for AI types — machine learning, generative AI, large language models, agentic AI, and related terms — are clear and aligned with current industry usage. The distinction between AI types is appropriately drawn and consistent with how these terms appear in the practices themselves.

We note two definitional gaps on the governance side.

First, the glossary does not define what a governance record is. Sound Practice 3 calls for documentation; the practices throughout reference accountability, audit trails, and record-keeping. But the report does not define the artifact those requirements are intended to produce. "Documentation" and "record-keeping" appear as requirements without a shared definition of what a compliant record contains or what it should be capable of demonstrating. A glossary definition — even a brief one — would reduce the implementation variance that undefined terms create across institution types and jurisdictions.

Second, the glossary would benefit from a clearer distinction between human oversight as a process and the record of oversight. Sound Practice 10 addresses meaningful versus nominal oversight thoughtfully, but the glossary currently defines oversight in functional terms only. Distinguishing between the oversight mechanism (what humans do) and the oversight record (what an institution can produce to demonstrate that oversight was exercised) would make the practices significantly more actionable for compliance and audit purposes.

Both gaps point to the same underlying issue: the glossary is well-calibrated to AI technology terms and underdeveloped on the governance artifact side. Strengthening the latter would materially improve how institutions and supervisors interpret the practices consistently.

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

We appreciate the opportunity to contribute to this consultation as a practitioner voice. Our responses across Questions 1 through 7 reflect a single consistent observation: the FSB's

proposed sound practices are well-designed at the policy level and would be strengthened by greater specificity at the record level — what institutions should be able to produce, to whom, and under what circumstances, to demonstrate that governance was actually exercised.

The distinction between a governance framework and a governance record is not a technical detail. It is the difference between an institution that has a policy and an institution that can prove it followed one. That distinction will become the central question in regulatory examinations, litigation, and supervisory review as AI adoption scales.

We believe this consultation represents a meaningful step toward international alignment on responsible AI governance. We welcome the FSB's continued engagement with practitioner perspectives as the final report is developed, and we are available to provide additional input if useful.

Respectfully,

Sue Yeh Johnson

Founder, Yeh and Johnson Consulting LLC