



FINANCIAL
STABILITY
BOARD

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

Response to Consultation

Jiajun M (independent response)

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 with two additions.

1. Risk is stage-dependent. The report assesses risk use case by use case, never at the level of the institution's overall adoption. A hallucinating pilot chatbot used by five analysts and a hallucinating system embedded in portfolio decisions carry the same label but are different risks. Probability and blast radius both grow as an institution moves from exploring to embedded, and fastest where AI feeds decisions rather than drafts. Naming this would sharpen the materiality logic of SP5.

2. The report misses the adoption gap between fund managers and the asset owners who fund them. Its own census puts sophisticated trading firms and asset managers among the most advanced adopters; asset owners do not appear in the report at all. Owners move slower for structural reasons, mainly fiduciary duty and heavier governance overhead per decision. The consequence matters for stability: asset owners are the monitoring layer for the capital they allocate. An owner that cannot evaluate its managers' AI-driven sourcing, valuation, or risk processes cannot supply the discipline that private markets depend on, and capital keeps flowing to strategies whose model behaviour, including correlated positioning on common models, the funding side can no longer assess. Without a staged measure of adoption on both sides, this gap stays invisible.

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

Clear, yes. Complete, no. The practices govern conduct but do not measure where an institution stands.

Three practices assume that knowledge. SP1 aligns AI adoption with strategy; alignment needs a starting point. SP4 adapts as AI evolves; adaptation needs a baseline. SP5 assesses use-case materiality at inception and thereafter; materiality shifts as adoption deepens, and the report gives no scale for that. "Maturity" and "adoption stage" appear nowhere in the report. The Section 3.1 lifecycle is a different axis: it tracks a use case from inception to retirement, not the institution. The report concedes the problem itself: the

Section 2.1 case study notes institutions setting adoption targets found "finding suitable metrics remains challenging."

The FSB has named this gap at system level: the October 2025 monitoring report identified data gaps in tracking adoption. Firm-level staged self-classification is the micro complement; institutions on a common scale let supervisors aggregate what they cannot currently observe. Treasury's AI Innovation Series (concluded 24 June 2026) framed the sector's challenge as operationalising AI at scale, a stage-transition problem.

I recommend the addition of one sentence in the final report recognising staged adoption classification as an implementation aid for SP1, SP4 and SP5.

Moreover, SP12 scopes third-party risk to performance, transparency, data quality, supply chain and concentration, and continuity. Liability allocation is missing. Standard vendor terms limit provider liability, so responsibility stays with the institution whatever the contract says. Acknowledging this strengthens the case for compensating controls and human oversight (SP8 to SP10).

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?

Adequate. Institutions move through the same four stages whether the underlying systems are predictive, generative or agentic. Agentic AI changes the cost of not knowing your stage: autonomy compounds errors faster than assistance. That strengthens the case for measurement without altering any of the twelve practices.

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

Yes, and a stage-anchored report ages better than a technology-anchored one. Tools named in 2026 look dated by 2028; institutions will still sit between exploring and embedded whatever runs underneath. Stage-anchoring is the cheapest flexibility available, and fits the direction of IOSCO's May 2026 supervisory toolkit: durable, comparable oversight across jurisdictions.

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.

No. The case studies skew to banks and insurers; nonbank asset owners (public pension funds, endowments, sovereign wealth funds) are absent. They face the same adoption decisions under different constraints: fiduciary duty to beneficiaries, thinner technology staffing, public accountability most banks do not carry. Governance-first by necessity, they are a distinct case-study category, not a smaller bank. Their absence also matters beyond coverage: asset owners oversee the fund managers they finance, so owner-side capability is a financial-stability question. A menu only helps readers who can find themselves in it.

Such case studies can be built from public disclosures alone, where scaling institutions publish metrics, and piloting institutions publish intentions. My classification (<https://jima1835.github.io/ai-adoption-in-finance>) documents adoption arcs at sovereign funds and public pension managers in the anonymised format the report already uses; the FSB is welcome to draw on it.

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

Mostly. I recommend tagging each case study with adoption stage and implementation weight (team size, build time, infrastructure). Readers look for institutions one step ahead of them. Without a stage tag, a piloting-stage reader cannot tell which examples apply now; without a weight tag, a lightly staffed institution cannot tell a G-SIB-grade build from a lightweight path. Some case studies already mention this in passing (one agentic system: in-house, three months). Making it systematic costs little and operationalises the report's proportionality principle.

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

Yes with two suggestions. First, add "adoption stage" as the glossary defines forms of AI but no degrees of adoption.

Second, the only agentic entry defines agentic AI by limited human oversight, and there is no term for what institutions deploy first. I recommend a definition on "human-supervised agentic AI" as a subset of agentic AI with agents restricted to approved use cases behind human-in-the-loop approval and a responsibility gate. The Section 2.3 case study describes exactly this (agent certification within defined boundaries) without naming it. Autonomy is the report's main risk driver; the glossary should name its gradations. "Human-supervised" avoids clashing with the existing supervised-learning entry.

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

I submit this response in my personal capacity. The views are my own, not my employer's or any institution's, and are drawn entirely from public sources and the consultation report itself. I do not request confidential treatment and consent to publication.

In summary, the twelve practices govern conduct but do not measure position. Nothing in the report lets a board or supervisor answer the prior question of where an institution actually stands in AI adoption, yet three practices presuppose that information. Strategic direction (SP1) needs a starting point, adaptability (SP4) needs a baseline, and materiality assessment (SP5) shifts as adoption deepens. Staged adoption classification closes the gap, works today from public data, and connects the sound practices to the data gaps the FSB's own October 2025 monitoring report acknowledged.

Two recommendations. One sentence in the final report recognising staged adoption classification as an implementation aid for SP1, SP4 and SP5; it names a tool and costs nothing in flexibility. And case-study coverage extended to nonbank asset owners, whose

absence is a stability blind spot as much as a coverage gap; material for this category can be built from public disclosures alone.

Jiajun Ma, ISO/IEC 42001 Lead Implementer

References:

1. FSB, Sound Practices for Responsible Adoption of AI: Consultation report, 10 June 2026 (<https://www.fsb.org/uploads/P100626.pdf>).
2. FSB, Monitoring Adoption of AI and Related Vulnerabilities in the Financial Sector, October 2025 (<https://www.fsb.org/2025/10/monitoring-adoption-of-artificial-intelligence-and-related-vulnerabilities-in-the-financial-sector/>).
3. US Department of the Treasury, AI Innovation Series conclusion, 24 June 2026 (<https://home.treasury.gov/news/press-releases/sb0540>).
4. IOSCO, Supervisory Toolkit, May 2026 (<https://www.iosco.org/library/pubdocs/pdf/IOSCOPD823.pdf>).
5. Ma, J., AI Adoption in Finance, live dashboard (<https://jima1835.github.io/ai-adoption-in-finance>).