FSB FINANCIAL STABILITY BOARD

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

Response to Consultation

Nakhile Consulting

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?**

   Broadly yes. One risk warrants explicit treatment: the absence of a defined governance object at the level of the individual consequential decision. The report identifies inadequate governance, blurred accountability in automated interactions, and reduced accountability in AI-driven processes, but treats these as organisational or use-case properties. In our advisory practice with financial institutions, the risk that materialises in deployment is more specific: a single AI use case (for example claims handling or credit management) contains decision classes of very different consequence to the customer and the institution, and where governance attaches to the use case rather than the decision, the highest-consequence decisions within a use case routinely carry no named human owner. International precedent, notably the Dutch childcare benefits case, illustrates how individually ungoverned automated decisions convert at scale into institutional and systemic harm without recourse. We suggest the final report name this explicitly as a risk: consequential decisions without identified ownership.

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

   The practices are comprehensive at the two levels they address: organisation-wide governance (SP1 to 4) and use-case lifecycle management (SP5 to 10). We respectfully suggest they are incomplete at a third level where, in our applied experience, risk actually enters: the individual consequential decision. Three refinements would close the gap. First, SP5's materiality assessment should be conducted at the level of decision classes within a use case, not only the use case in aggregate; a claims use case may combine trivial routing decisions with repudiation decisions of the highest customer consequence, and a single use-case rating obscures this. Second, SP2's accountability framework should require a named human owner for each high-materiality decision class, accountable for its outcomes and its recourse pathway; the three lines of defence allocate functions, but functions are not owners of decisions. Third, SP10's human oversight should be calibrated per decision class rather than per use case, which also resolves the report's own observation that real-time oversight of agentic systems is impractical: oversight scales when it attaches to decision classes with defined escalation thresholds rather than to systems as a whole. These refinements are operationalisable: our published Decision Materiality Index (DOI 10.5281/zenodo.21134247) classifies automated decisions by consequence and assigns named ownership, and has been applied in live engagements at a short-term insurer and an unsecured lender in South Africa, both operating high-materiality decision systems. We would be glad to share methodological detail as supplementary material.

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?**

   No comments beyond those provided under Questions 1 and 2.

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

   No comments beyond those provided under Questions 1 and 2.

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 skew toward large, internationally active institutions in advanced jurisdictions. We suggest the final report include at least one case study from an emerging market and one from a smaller nonbank, for two reasons. First, the FSB's proportionality principle needs demonstration, not only assertion: smaller institutions reading only G-SIB case studies conclude the practices are not for them. Second, emerging-market institutions often deploy consequential AI (credit scoring on alternative data, claims automation) with thinner governance resources, making them the environments where sound practices are most needed and least illustrated. From our practice we can offer anonymised applied material from South African financial services, including a short-term insurer and an unsecured lender implementing decision-level governance, and would be glad to provide this via fsb@fsb.org as supplementary material if useful.

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

   No comments

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

   No comments

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

   We welcome the report and support its finalisation. Our comments reflect a single perspective, offered from applied practice: the sound practices govern well at the level of the organisation (SP1 to 4) and the AI use case (SP5 to 10), but the unit at which AI risk actually crystallises for customers, institutions and ultimately financial stability is the individual consequential decision: the credit declined, the claim repudiated, the account suspended. Where materiality assessment, accountability and human oversight attach to decision classes rather than to use cases in aggregate, the practices become materially easier for institutions of all sizes to operationalise, and the report's own proportionality principle gains a concrete mechanism: proportionality by consequence of the decision, rather than by size of the institution alone. This perspective draws on live AI governance engagements in South African financial services, an emerging market context underrepresented in the report's case studies, and on a published framework for classifying automated decisions by consequence (Decision Materiality Index, DOI 10.5281/zenodo.21134247). We thank the FSB for the opportunity to comment and would welcome any follow-up via the contact details provided.