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

Response to Consultation

Association for Savings and Investment in South Africa (ASISA)

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?**

   ASISA members broadly agree with the benefits and risks of AI adoption identified in the report. The report provides a comprehensive and balanced assessment of AI's potential to improve operational efficiency, customer experience, fraud detection, financial crime monitoring, risk management, regulatory compliance, decision-making and operational resilience, while appropriately recognizing key risks relating to governance, data quality, explainability, human oversight, cyber and ICT risk, third-party dependency, concentration risk, consumer harm, and the emerging challenges associated with GenAI and agentic AI.

   The following observations are worth noting:

   • Workforce readiness and human risk: recognition of the risk of workforce de-skilling and capability atrophy resulting from widespread AI adoption and the responsibility that financial institutions have to retain sufficient expertise to oversee AI systems, challenge outputs and meet governance and fiduciary obligations. Skills development, continuous learning and change management should be recognised as critical enablers of responsible AI adoption.

   • Data governance: in the adoption of AI, proper attention should be given to data protection and privacy laws that apply in different jurisdiction in which a financial institution operate – considering international best practices. Convergence with the implementation of data privacy laws should be considered with the adoption of AI.

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

   ASISA members agree sound practices are comprehensive, well-structured and sufficiently clear to support responsible AI adoption by financial institutions. The distinction between organization-wide governance and AI lifecycle controls is particularly helpful, and the principles-based, proportional approach appropriately accommodates different institution sizes, business models and levels of AI maturity. The framework addresses the key elements required for responsible AI adoption, including governance, accountability, data management, human oversight, explainability, performance monitoring, cybersecurity and third-party risk management.

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?**

   ASISA members agree the sound practices strike an appropriate balance between managing risks associated with AI generally and addressing the additional risks presented by emerging technologies such as GenAI and agentic AI.

   The principles-based and technology-neutral approach is a particular strength, as it ensures that the framework remains relevant as AI technologies evolve. The report appropriately focuses on risk characteristics such as materiality, autonomy, complexity, explainability, data dependency and customer impact, rather than creating separate governance frameworks for different AI technologies. This supports proportional and risk-based implementation across a wide range of use cases.

   At the same time, the report appropriately recognizes that GenAI and agentic AI can create heightened risks relating to hallucinations, explainability, human oversight, cybersecurity, data governance, confidentiality, third-party dependency and autonomous decision-making. The guidance on governance, monitoring, guardrails, logging, restrictions on system access and kill-switch mechanisms provides a sound foundation for managing these risks.

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

   ASISA members agree the Sound Practices appear sufficiently flexible to accommodate newer forms of AI and support responsible adoption over time. Their principles-based, technology-neutral, risk-based and lifecycle-oriented approach is a significant strength, as it focuses on enduring governance outcomes rather than prescribing controls for specific technologies that may quickly become outdated.

   The emphasis on governance, accountability, proportionality, materiality, human oversight, monitoring, data governance and risk management provides a framework that can be applied across traditional AI, machine learning, GenAI, agentic AI and future AI technologies that may emerge.

   The inclusion of organisational adaptability and the recognition that AI technologies and risks will continue to evolve further support the framework's long-term relevance. The report appropriately avoids a one-size-fits-all approach and allows institutions to tailor controls according to the nature, scale, complexity and risk profile of individual AI use cases.

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 are useful, practical and effective in demonstrating how financial institutions can realize benefits from responsible AI adoption while implementing appropriate governance, risk management and oversight controls. The examples illustrate a range of AI

applications, including credit risk management, fraud detection, customer engagement, operational efficiency and AI governance.

An ASISA member submitted the below practical example to illustrate how the principle of proportionality, which underpins the sound practices, may be applied to a lower-risk, decision-support AI use case within a smaller financial institution. It is believed that this could further enhance the practical application of proportionality within the paper.

Case study: AI-assisted investment research

A South African asset manager has deployed an AI-enabled assistant to support its investment team with research and information retrieval. The investment team previously spent significant time reviewing large volumes of information from a range of internal and external sources as part of the investment process.

The solution assists in identifying, retrieving and analyzing relevant information, reducing the manual effort involved. The AI solution is intended solely as a decision-support tool. AI-generated responses are used to support research and do not constitute investment recommendations.

The institution has implemented governance controls proportionate to the nature and risk of the use case. The AI solution has no authority to execute trades, publish research, communicate externally or amend investment records. Access is restricted to authorized users, AI-generated outputs are verified before being relied upon, and all investment decisions remain subject to the independent judgement and approval of the portfolio manager.

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

Yes, the case studies provide useful and generally actionable insights for financial institutions seeking to adopt AI responsibly. They help move the discussion beyond high-level principles by demonstrating how governance, testing, human oversight, explainability, monitoring, phased deployment, data management and third-party controls can be applied in practice.

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

The glossary is generally clear, comprehensive and aligned with current industry terminology and emerging AI developments. It appropriately covers many of the key concepts relevant to responsible AI adoption, including generative AI, agentic AI, foundation models, model drift, data drift, explainability and data/model poisoning. The glossary provides a useful foundation for establishing a common understanding of AI-related concepts across jurisdictions, sectors and stakeholders.

ASIA members propose the following for consideration:

- The definition of 'Agentic AI' would benefit from greater precision -

The current definition does not distinguish between single-agent and multi-agent systems, nor does it capture the concept of tool use and environmental interaction that is central to the report's own analysis of agentic AI risks in Section 1.3.

Suggested addition: 'Agentic AI systems may operate as single agents or as networks of multiple interacting agents (multi-agent systems). They may access and use tools, APIs, databases, and other external systems to achieve their objectives, and may dynamically adapt their approach based on environmental feedback.'

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

1. The numbering of the paragraphs will assist usage with the ability to reference to specific aspects/concepts in the document.

2. As the use of AI is developing and the technology is evolving, it might be required to review these principles over time to ensure everything is still relevant, or if certain aspects/concepts need to be enhanced or if any new aspects/concepts need to be included.