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

## Response to Consultation

## South African Insurance Association

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 report should continue emphasising that AI governance is most effective when embedded within existing enterprise governance, risk management, compliance, operational resilience, information security, privacy, procurement, model risk, legal, internal audit and third-party risk management frameworks, rather than creating standalone AI governance structures. AI governance should be recognised as a board and executive accountability matter integrated within existing governance arrangements.

   While supporting the principles-based approach, the report should more clearly distinguish between baseline governance expectations applicable to all AI use cases and enhanced governance requirements applicable to material, customer-impacting, autonomous or higher-risk AI systems. Greater practical guidance on proportional implementation across different categories of AI would promote greater consistency across the industry.

   Greater emphasis should be placed on privacy, data protection and customer rights throughout the AI lifecycle. These considerations should extend beyond data quality and security to include lawful processing, transparency, data minimisation, retention, deletion, cross-border transfers, automated decision-making, explainability and protection of personal information across prompts, logs, retrieval-augmented generation (RAG) repositories, embeddings and other AI artefacts.

   The report would benefit from expanded guidance on governance of emerging AI capabilities, including GenAI, agentic AI, orchestration layers, AI gateways, tool-calling mechanisms and multi-model environments. These technologies introduce distinct governance, accountability and operational risks that extend beyond traditional model governance.

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

   The Executive Summary should reinforce that the Sound Practices are intended to support implementation through existing governance and risk management frameworks and encourage supervisory consistency while remaining proportionate to institutional risk

profiles. The wording should also reflect that baseline governance controls are expected for all AI use cases, with enhanced controls applied according to materiality and risk.

While the Sound Practices are not intended to establish binding international standards, they should be recognised as supporting supervisory convergence and consistent regulatory dialogue. The report may also clarify that implementation should occur within applicable domestic legal and regulatory frameworks.

Consider expanding the identified AI risk categories to include privacy and data protection, intellectual property, operational resilience, environmental considerations, data sovereignty, lawful cross-border transfers, vendor lock-in, substitutability, portability, geopolitical risks, concentration risk and systemic

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?**

While the Sound Practices are not intended to establish binding international standards, they should be recognised as supporting supervisory convergence and consistent regulatory dialogue. The report may also clarify that implementation should occur within applicable domestic legal and regulatory frameworks.

Consider expanding the identified AI risk categories to include privacy and data protection, intellectual property, operational resilience, environmental considerations, data sovereignty, lawful cross-border transfers, vendor lock-in, substitutability, portability, geopolitical risks, concentration risk and systemic dependencies associated with common AI service providers

Beyond disclosure obligations, the report should recognise explainability, meaningful human review, contestability and accessible redress mechanisms where AI materially influences customer outcomes or automated decision-making affecting customers' rights, eligibility, pricing, claims or access to financial services.

As agentic AI becomes more prevalent, additional guidance should address governance beyond traditional model controls. Institutions should establish controls for identity management, access rights, delegated authority, segregation of duties, behavioural monitoring, escalation, supervision, accountability and safe deactivation of AI agents. Governance should also recognise the need for effective human orchestration and supervision of agentic AI activities.

Practical examples of controls to manage shadow AI would assist implementation. These may include approved AI tool catalogues, awareness programmes, procurement controls, network restrictions, SaaS monitoring, data loss prevention measures, exception processes and provision of approved alternatives to discourage unauthorised AI use.

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

Consider providing guidance on the minimum management information (MI) that boards and executive committees should receive. This may include AI inventories, risk-tiering information, material AI use cases, incidents and near misses, third-party dependencies,

model performance, data quality issues, customer complaints, regulatory developments, control exceptions, remediation progress and benefits realisation. Such reporting would promote more consistent board oversight across institutions

Beyond disclosure obligations, the report should recognise explainability, meaningful human review, contestability and accessible redress mechanisms where AI materially influences customer outcomes or automated decision-making affecting customers' rights, eligibility, pricing, claims or access to financial services.

SAIA supports developing organisational capability for responsible AI adoption. AI training should be role-based, proportionate to risk exposure and periodically refreshed. Training should address permitted and prohibited AI use, data handling, confidentiality, privacy, prompt safety, bias, hallucinations, cyber risks, human oversight responsibilities and escalation obligations. Institutions should continue strengthening technical governance capabilities, including infrastructure monitoring, vulnerability management and secure AI engineering practices.

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.**

   Additional practical examples and sector-specific case studies, particularly for insurers and lower-materiality AI use cases, would assist institutions in applying the Sound Practices consistently across varying AI deployments.

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

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

   The glossary could be expanded to include additional terminology such as AI inventory, AI use case, high-risk AI, orchestration layer, tool-calling, guardrails, model cards, system cards, human oversight, contestability, AI incidents, prompt injection, embedded AI, frontier models and open-weight models. Common definitions would support more consistent implementation across the industry.

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

   Consider introducing a dedicated Sound Practice addressing governance of AI orchestration layers, gateways, multi-model routing and tool-calling capabilities. As institutions increasingly deploy multiple models and AI agents, this layer represents a distinct governance and control point requiring oversight, accountability and monitoring.

   Consider introducing guidance encouraging institutions to maintain the capability to independently suspend or disable production AI solutions where necessary to respond to incidents, unacceptable risk or operational failures.