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

Response to Consultation

Deutsche Börse Group

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?

As the leading European market infrastructure provider, Deutsche Börse Group (DBG) broadly agrees with the benefits and risks identified in Sections 1.2 and 1.3 of the report. The mapping of AI use cases across the industry reflects the adoption patterns we observe across our own operations and within the broader market infrastructure ecosystem. We particularly welcome the report's recognition that FMIs occupy a distinct position in the financial system.

As an operator of financial market infrastructure across the entire value chain, in our experience, transitioning from machine learning-based document processing to LLM-based processing achieves significant improvements in time-to-insight compared to traditional approaches. We internally use AI and agent-based platforms to further optimize operations, analytical, and administration tasks. DBG operates an internal AI assistant built on a secure, governed platform with multiple, diversified existing and future model integrations (Google Gemini, Azure OpenAI, Anthropic), enabling employees across all business functions to leverage AI for drafting, research, information retrieval, and knowledge management.

Beyond the specific use cases noted in the report, we would like to draw attention to the following benefits that are particularly relevant to FMIs and that merit more explicit recognition:

The report's treatment of enterprise-wide AI productivity platforms could be expanded. The report describes "productivity-enhancing applications" primarily through examples of individual task automation. In practice, the category of enterprise-wide AI assistant platforms, governed by comprehensive technical and procedural controls and limited to supportive, non-critical processes, represents a qualitatively different and increasingly significant category of AI adoption. These platforms require institutional-scale governance that differs materially from the governance challenges of individual employee use of external AI tools. The FSB's description of productivity-enhancing applications would benefit from reflecting this distinction.

On the risk side, we broadly concur with the nine risk categories identified. However, we would like to highlight several areas where the treatment could be deepened:

First, on the topic of shadow AI, we share the FSB's concern regarding the use of unauthorized AI tools by employees. DBG has addressed this through a combination of approved platform provisioning, mandatory AI literacy training for all employees, and an AI Code of Conduct that explicitly limits AI use to authorized tools and prohibits personal use of AI systems for business purposes. Our framework provides a streamlined, risk-based governance pathway for agent development on approved platforms, enabling innovation while maintaining central oversight. We recommend that the report emphasizes more strongly that shadow AI risk is best mitigated through the provision of attractive, governed alternatives rather than through prohibition alone.

Second, regarding the risks specific to multi-cloud and multi-provider AI dependencies, the report addresses third-party dependencies and concentration but should recognize that financial institutions increasingly operate multi-provider AI stacks as a matter of operational reality. The report should avoid imposing governance requirements that assume a single-provider model and instead allow institutions the flexibility to manage cross-platform model governance, consistent security controls, and divergent update cycles across providers in a manner proportionate to their specific architecture.

Third, the report's treatment of agentic AI risks is welcome and timely. DBG is actively developing agentic AI capabilities across multiple platforms. From our practical experience, we would note that the report could more explicitly address the risk that arises when multiple agent frameworks coexist within a single institution, each with different security models, identity management approaches, and oversight mechanisms.

Fourth, while the report correctly identifies geopolitical risks, it does not explicitly address the risk of service disruption or restriction due to geopolitical tensions. For a European FMI operator, the dependence on predominantly US-based cloud and AI providers creates a category of strategic risk that goes beyond traditional third-party risk management. We would suggest that the FSB consider referencing this geopolitical dimension more explicitly, as it has direct implications for business continuity planning and exit strategy requirements.

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

Deutsche Börse Group considers the twelve Sound Practices to be broadly comprehensive and well-structured. They provide a useful high-level framework that is compatible with the EU AI Act requirements we are already implementing and with the governance structures we have established. The proportionality principle is particularly welcome, as it recognizes that a one-size-fits-all approach would be neither practical nor desirable given the diversity of financial institutions and AI use cases. That said, we would offer the following observations and recommendations for improving the clarity and practical applicability of some of the sound practices:

Sound Practice 1 (Strategic Direction and Oversight): We agree that board and senior management oversight is essential. DBG has embedded AI at the executive level, with joint Supervisory Board and Executive Board workshops on AI covering opportunities, implementation, and the EU AI Act. In our view, the FSB sound practices should avoid prescribing specific internal governance structures and instead allow institutions flexibility in

how they coordinate AI governance with existing IT and operational governance frameworks.

Sound Practice 2 (Governance and Accountability): We support the three lines of defense approach. The FSB sound practices should recognize that AI development is increasingly decentralized through citizen development and no-code platforms, and that institutions need the flexibility to apply risk-based differentiation in their approval processes accordingly.

Sound Practice 3 (AI Risk Management Framework and Documentation): We support the emphasis on a central AI Use Case Inventory. DBG maintains such an inventory since 2024, tracking use cases through a defined lifecycle with EU AI Act risk classification, assigned owners, and clearance. However, the FSB sound practices should allow for proportionate documentation requirements across the spectrum of use cases. For example, prompt versioning for rapidly iterating citizen-developed agents presents fundamentally different challenges than for centrally developed production systems. The documentation requirements for GenAI and agentic AI (prompt versioning, chain-of-thought logging) are directionally correct but should be formulated with greater awareness of practical constraints.

Sound Practice 4 (Organizational Adaptability): DBG has implemented mandatory AI training for all employees, an AI Ambassadors Programme as a group-wide network of trained multipliers, and joint board workshops conducted by internal and external experts. The FSB sound practices should recognize that effective approaches to building AI competence vary across institutions and avoid prescribing specific training formats.

Sound Practice 5 (Materiality and Risk Assessment): We agree with the distinction between materiality and risk. DBG applies the EU AI Act risk classification and has confirmed it does not operate prohibited or high-risk AI systems under this framework. The FSB sound practices should address how the FSB's materiality/risk framework relates to the EU AI Act's risk classification, to avoid parallel assessment obligations.

Sound Practice 6 (Selection): The FSB sound practices should recognize that institutions will operate a portfolio of AI tools and models, and that selection is an ongoing evaluation, not a one-time event. It should avoid framing model selection as a static decision, as institutions may need to replace foundation models due to cost, performance, or provider policy changes, and governance frameworks must retain the flexibility to accommodate this without triggering disproportionate regulatory obligations.

Sound Practice 10 (Human Oversight): DBG has developed Human in the Loop (HITL) guidance structured across three layers (policy, workflow, technology) with concrete implementations including automated performance alerts, mandatory human review for LLM outputs, and multiple-eye verification for agentic AI outputs. The FSB sound practices mention the risk of "rubber-stamping" but provides insufficient guidance on detection and prevention. In our view, effective countermeasures include analyzing oversight patterns for automated approval behavior, rotating reviewers, and implementing random challenge tests with known-incorrect outputs.

Sound Practice 11 (Cyber and ICT Risk Management): The FSB sound practices do not address a tension that is becoming acute as institutions scale agentic AI. Maintaining the

need-to-know principle while pooling data for AI consumption. Traditional access control architectures are designed around human users with static role assignments; agentic AI breaks this model, as an agent may query pooled data on behalf of users with different clearance levels, or chain data from multiple sources where the permissibility of the output depends on who is asking. The FSB should recognize this as a first-order security architecture challenge rather than assuming existing role-based access controls are sufficient for AI-mediated data access.

Sound Practice 12 (Third-Party AI Risk Management): The FSB sound practices should acknowledge that the relationship with foundation model providers is fundamentally different from traditional IT outsourcing, as model updates can silently alter the behaviors of deployed systems. Contractual provisions must specifically address advance notification of model changes, the right to pin model versions, and clear escalation paths.

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?**

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

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.**

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?**

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