

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

Response to Consultation

Italian Banking Association (ABI)

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?

ABI broadly agrees with the benefits and risks described in the report. The analysis reflects many of the areas where banks are already using or exploring AI, including fraud detection, AML/CFT, cybersecurity, customer assistance, credit processes, internal controls, operational efficiency, software development and knowledge management. AI can also act as an important enabler for control functions, risk management and compliance, as well as a tool to scale organisational knowledge and decision-support capabilities across large and complex institutions.

At the same time, ABI considers that some potential benefits of AI remain more prospective and should be assessed with appropriate caution. This is particularly the case for business model transformation and for the use of more autonomous AI systems in critical operations, core decision-making processes or customer-impacting use cases, where deployment is expected to remain gradual and subject to robust governance, testing, human oversight and risk management.

ABI also agrees that AI may introduce or amplify risks relating to data quality, explainability, model performance, human oversight, cyber and ICT security, third-party dependencies and concentration. These risks are well understood by banks and are already being addressed through existing governance, risk management and control frameworks.

In this context, particular attention should be paid to dependencies on third-party AI providers and on the broader AI supply chain. These may involve limited transparency over underlying models, vendor-driven changes to models or systems, uncertainty around data handling and reuse, and concentration across key layers such as cloud infrastructure, hardware, foundation models, data providers and deployment platforms. Such dependencies may create operational and resilience risks, but also challenges for effective oversight, accountability and business continuity, especially where AI is used in critical services or material functions.

At the same time, the risk of not adopting AI should also be recognised. In areas such as fraud prevention, cyber defence, anomaly detection and operational resilience, internal controls and compliance monitoring. AI is becoming increasingly relevant to maintain an

effective level of protection. Responsible AI adoption should therefore be seen not only as a source of new risks, but also as a necessary component of sound risk management.

The report could also acknowledge certain longer-term challenges associated with the development and deployment of advanced AI models, including increasing computational and energy demands and their potential broader economic implications.

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

ABI considers the sound practices broadly comprehensive and useful as a high-level reference framework. Their structure, distinguishing between organisation-wide governance and AI lifecycle management, is consistent with the direction already taken by many banks. At conceptual level, the practices are generally clear and provide a valuable common reference for responsible AI adoption. However, some further operational clarification could support more consistent implementation in practice across jurisdictions and institutions, provided that such clarification remains non-prescriptive and principle-based.

It is essential that the sound practices remain principle-based, proportionate and sufficient flexible to accommodate different institutions, business model and use cases. In this respect, the proportionality principle is key. The level of controls, documentation and oversight should depend on the materiality, risk, autonomy, customer impact and criticality of the specific AI use case. Further clarification would be particularly useful on how proportionality should be applied in practice, including how institutions may calibrate governance, documentation, testing and approval requirements for low-risk, medium-risk and high-risk AI use cases. Low-risk or low-materiality applications should not be subject to the same level of requirements as high-impact or highly autonomous AI systems used in critical functions.

Similar clarification would be beneficial with respect to the calibration of human oversight. The appropriate form and intensity of human oversight should depend on the specific use case, including the level of autonomy of the AI system, the degree of reliance on its outputs, the potential impact on customers or markets, and the feasibility of effective human intervention. This is particularly relevant for more autonomous forms of AI, including agentic AI, where governance arrangements may need to address escalation mechanisms, limits on permitted actions, monitoring of intermediate steps, and clear accountability for outcomes.

Banks already operate under extensive requirements on governance, internal controls, ICT risk, cyber risk, outsourcing, operational resilience, data protection, model risk and prudential supervision. The FSB practices should therefore avoid generating duplicative obligations and should be interpreted as a tool to leverage and enhance existing frameworks rather than create parallel governance structures.

Further clarification should support consistent implementation in practice, while preserving the principle-based, proportionate and flexible nature of the sound practices.

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?

ABI broadly agrees with the proposed balance. It is appropriate to identify practices that are relevant to AI adoption in general, while recognising that emerging forms of AI, including GenAI and agentic AI, may require additional attention due to their evolving capabilities, increasing autonomy, interaction with external tools and potential cyber implications.

At the same time, ABI underlines that the level of governance and risk management should not be determined solely by the technology label. Assessment should remain focused on the specific use case, including its purpose, level of autonomy, materiality, risk profile, data sensitivity, reliance on third parties, impact on customers and degree of human oversight.

The report could also acknowledge that AI development and deployment may become increasingly decentralised within financial institutions, including through mature low-code and no-code tools that allow business users to configure or adapt AI-enabled solutions more autonomously. While this can support innovation and faster experimentation, it may also increase governance, oversight and documentation challenges, including the need to maintain adequate visibility over AI use cases and to prevent or detect unauthorised uses, such as shadow AI.

Traditional AI, statistical models, decision-support tools, internal productivity solutions and low-impact applications should not be treated in the same way as highly autonomous systems operating in critical or customer-impacting processes. A proportionate and context-based approach remains essential to avoid unnecessary complexity and to preserve the ability of banks to innovate responsibly.

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

ABI considers the practices sufficiently flexible, provided that they remain non-prescriptive, technology-neutral and proportionate over time. AI technologies and adoption models continue to evolve rapidly. Banks are already adapting their organisational models, strengthening AI governance, developing AI factories or centres of expertise, investing in skills and integrating AI lifecycle controls into existing risk management and control frameworks. This evolution requires flexibility, not rigid requirements.

Future developments of the sound practices should therefore preserve their current nature as a high-level reference framework. They should not become detailed operational rules, nor should they create retrospective expectations for systems and frameworks already implemented under existing regulatory and supervisory requirements.

A flexible approach is particularly important to support innovation at scale. The main challenge for banks is increasingly to move from experimentation to controlled industrialisation, integrating AI into processes while maintaining strong governance, risk management and accountability.

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.

ABI considers the case studies useful as illustrative examples of responsible AI adoption. They show that AI can generate value in several areas, including risk management, operational efficiency, customer service and financial crime prevention, while highlighting the importance of governance, accountability and oversight arrangements.

At the same time, the case studies should be regarded as illustrative and non-exhaustive. While the examples are helpful and particularly relatable from a banking perspective, the report could further acknowledge that coverage of insurers, asset managers, payment firms and other non-bank financial institutions is more limited. AI adoption is evolving rapidly and use cases continue to expand across financial institutions. It would therefore be neither realistic nor desirable to aim for a complete or static mapping of possible applications.

For this reason we do not consider it necessary to add further case studies. The current examples already provide a useful overview of how responsible AI adoption can support a range of business and risk-management objectives.

In addition, the case studies should not be used as benchmarks, supervisory expectations or preferred organizational models. Financial institutions differ significantly in terms of business models, operating environment, technological maturity and AI strategies. The primary value of the case studies lies in supporting knowledge-sharing and illustrating possible approaches to responsible AI adoption rather than establishing standardised implementation models.

The most important message is that responsible AI adoption is already taking place within banks through structured governance, prioritisation mechanisms, lifecycle controls, risk assessments, human oversight and performance monitoring. Case studies can support knowledge-sharing, but they should not become a basis for additional regulatory expectations.

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

ABI considers that the case studies provide useful insights, especially insofar as they show that responsible AI adoption requires governance, accountability, controls, documentation, testing, monitoring, human oversight and integration with business processes. They also provide practical illustrations of how institutions can translate responsible AI principles into operational governance arrangements throughout the AI lifecycle.

However, their value should remain practical and illustrative. Financial institutions differ in size, business model, organisational structure, technological maturity and AI strategy. Therefore, the case studies should not be read as standardised implementation models.

In ABI's view, the key actionable insight is that AI governance should enable scale, not constrain it. Banks are increasingly focusing on selecting AI initiatives ex ante according to strategic coherence, expected benefits, feasibility, scalability and risk profile, and on measuring results ex post through indicators of effectiveness, efficiency, performance, quality, human feedback and risk control.

This confirms that the banking sector is already moving towards mature AI adoption models, where governance, measurement and control are embedded into the AI lifecycle. The FSB sound practices should support this direction without adding unnecessary procedural burdens.

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

ABI considers the glossary generally useful. Clear terminology is essential to support common understanding and supervisory convergence.

At the same time, ABI underlines the importance of avoiding overly broad definitions that could create uncertainty or include traditional statistical, deterministic or rule-based systems already widely used in banking and already subject to sectoral regulation. Definitions should also be sufficiently future-proof and technology-neutral, so that terminology remains applicable as AI technologies evolve and does not anchor supervisory expectations to specific technologies or implementation models. In addition, technical terms should be clearly framed within the AI context to avoid ambiguity with similar terminology used in other domains.

Definitions should be aligned, as far as possible, with existing regulatory frameworks and international standards, including the EU AI Act. In this regard, further alignment would be useful on specific terms such as agentic AI and AI agents, different forms of model openness, including open-source and open-weight models, and the relationship between foundation models and the regulatory concept of general-purpose AI models. Greater consistency across jurisdictions and standard-setting bodies would help reduce legal uncertainty, minimise compliance complexity and support a more coherent implementation of AI governance frameworks, particularly for internationally active financial institutions.

In particular, clarity is important on the distinction between AI systems and traditional models, between decision-support systems and automated decision-making systems, and between different levels of autonomy and materiality.

A clear and proportionate terminology is essential to avoid legal uncertainty, regulatory overlap and inconsistent implementation. This is particularly relevant for banks, where many analytical models are already governed by robust prudential, risk management, validation and supervisory frameworks.

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