

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

Response to Consultation

World Savings and Retail Banking Institute - European Savings and Retail Banking 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?

WSBI-ESBG broadly agrees with the benefits and risks of AI adoption by financial institutions as described in the FSB consultation report. The report rightly highlights AI's potential to improve operational efficiency, strengthen fraud detection, enhance AML/CFT and KYC processes, support credit risk assessment... These benefits are indeed already visible across areas such as transaction monitoring, customer support, regulatory compliance, portfolio management and operational optimisation.

At the same time, WSBI-ESBG welcomes the report's recognition that AI adoption is not only a technological development, but also a broader governance, risk management, operational resilience and consumer protection issue. AI systems may influence many key banking activities and is therefore important that their adoption remains embedded within institutions' existing frameworks, and that implementation remains proportionate to each financial institution. Indeed, AI adoption across jurisdictions worldwide remains highly uneven, as some larger entities are deploying advanced AI systems, while many banks in emerging markets are still at an exploratory stage.

Benefits:

WSBI-ESBG believes the report could, however, place greater emphasis on several additional benefits, notably:

1. AI's potential regarding advancing financial inclusion. AI could provide alternative credit assessment methods, personalised financial planning and education, simplified onboarding, and improved access to financial services for underserved populations.
2. AI's role as a multilingual tool, where automated translation, speech recognition and AI assisted response across multiple language can increase inclusion in linguistically diverse markets
3. AI's contribution to regulatory compliance, as AI can support regulatory reporting, sanctions screening, internal audit planning, policy compliance checks, model monitoring

and supervisory analytics. In addition, AI can support ESG and climate-related reporting by extracting and classifying relevant information from internal systems, documents, public sources and borrower disclosures

Risks:

WSBI-ESBG also agrees with the key risks identified in the report, including governance challenges, model risk, data quality issues, explainability limitations, inadequate human oversight, cyber risks, third-party dependencies, and consumer protection concerns. Nonetheless, several risks merit further attention.

1. Concentration risk arising from growing reliance on a limited number of cloud providers, foundation model developers and AI vendors may create systemic dependencies, vendor lock-in, data sovereignty concerns and operational resilience challenges.
2. Capability concentration, as large institutions have more resources for to develop proprietary AI capabilities, attract specialised talent and implement sophisticated governance frameworks. If proportionality is not operationalised effectively, responsible AI requirements could unintentionally widen the gap between larger institutions and smaller or mid-sized institutions
3. AI readiness, following with capability concentration, uneven AI adoption across jurisdiction could widen the competitiveness gaps between banking sectors, particularly where access to digital infrastructure, skilled professionals or supervisory capacity differ.
4. Risks associated with shadow AI, data leakage and the unauthorised use of external AI tools by staff.
5. Bias and financial exclusion risks, particularly where multiple institutions rely on similar datasets or models, potentially leading to the exclusion of vulnerable customer groups.
6. Advances in generative AI, which are increasing the sophistication of fraud schemes, including deepfakes, voice cloning, synthetic identities, phishing and social engineering attacks.

WSBI-ESBG also suggests placing more emphasis on reputational risk, customer trust and contestability. Retail banking relationships, especially with underserved or vulnerable customers, depend on perceptions of fairness, transparency and consistency. If AI-supported decisions are perceived as opaque, arbitrary or difficult to challenge, customer confidence may deteriorate even where models perform well technically. AI-related failures at one institution could also undermine confidence in AI-enabled financial services more broadly. Clear communication, human escalation channels, complaint-handling mechanisms and recourse arrangements are therefore essential. Responsible adoption should therefore include staff training, AI literacy and customer support as well as safeguards to ensure that AI does not reduce accessibility for customers who rely on physical, assisted or non-digital channels.

Overall, while AI shows significant benefits, challenges and risks still need to be overcome through responsible adoption frameworks, which take into account the varying pace of adoption by smaller and medium size institutions.

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

WSBI-ESBG considers the proposed sound practices to be broadly comprehensive, clear and well-structured. The framework appropriately addresses the key elements of responsible AI adoption, including governance and accountability, risk management, data governance, explainability, performance monitoring, human oversight, cyber and ICT risks, and third-party risk management. We particularly support the FSB's principles-based, risk-based and technology-neutral approach. This is preferable to a highly prescriptive international standard, given the rapid evolution of AI technologies and the diversity of financial institutions' business models, sizes, risk profiles and levels of technological maturity. The principles-based approach allows institutions to implement the sound practices in a manner proportionate to their size, complexity, resources and the materiality of their AI use cases. This is especially important for retail, savings, inclusion-focused and mid-sized institutions, which may adopt AI in lower-risk customer service, operational efficiency or financial inclusion use cases and may not have the same internal resources as large internationally active institutions.

In addition, the distinction between enterprise-wide governance and AI lifecycle management is also helpful, and the emphasis on aligning AI adoption with business strategy, risk appetite and existing governance structures is appropriate. The framework also rightly recognises the importance of proportionality, as not all AI use cases carry the same level of risk or require the same level of oversight.

Nonetheless, while the sound practices are clear at a principles level, additional practical implementation guidance would still enhance their effectiveness.

Firstly, financial institutions would benefit from additional guidance illustrating how implementation expectations may differ according to institutional size, technological maturity, materiality of AI use cases, degree of autonomy, customer impact and systemic relevance. Responsible AI adoption should be viewed as a progressive journey rather than a binary state of compliance. Banks should be encouraged to strengthen governance frameworks incrementally as AI adoption expands, allowing governance capabilities to evolve alongside the increasing materiality and complexity of AI use cases. Such guidance could take the form of an indicative implementation matrix, a light-touch maturity model, or practical examples showing how low, medium and high-risk AI use cases may be subject to different levels of governance, documentation, testing and oversight. This would support consistent implementation while avoiding unnecessary compliance burdens, particularly for smaller and mid-sized institutions.

Secondly, WSBI-ESBG suggests that the report could better connect the sound practices with existing regulatory, operational risk, model risk management, outsourcing and corporate governance frameworks wherever possible. Financial institutions already operate under mature requirements related to operational resilience, model risk management, outsourcing, cybersecurity, data protection, consumer protection and conduct risk. An informative mapping between the FSB sound practices and widely used frameworks, , would help institutions integrate AI governance into existing structures rather than building duplicative parallel processes.

Thirdly, the report should explicitly acknowledge other implementation models. For many banks, particularly smaller and medium-sized ones, AI adoption will primarily occur through third-party solutions rather than internally developed models. The sound practices should explicitly acknowledge this implementation model and clarify how governance, due diligence and oversight expectations should be applied when institutions rely predominantly on third parties i.e. external technology software and/or infrastructure providers

Finally, WSBI-ESBG also believes that the report could also provide further guidance on operationalising proportionality, clarifying different forms of human oversight, and encouraging information-sharing and collaboration mechanisms where appropriate. Future iterations could also more explicitly address data sovereignty, customer recourse mechanisms, workforce reskilling and the environmental impact of large-scale AI adoption.

Overall, the sound practices provide a robust foundation for responsible AI adoption. Their flexibility and technology-neutral design are key strengths, but the framework could be further strengthened through practical implementation examples, proportionality guidance and additional tools to support institutions at different stages of AI maturity.

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?

WSBI-ESBG generally agrees that the sound practices strike an appropriate balance between managing risks associated with 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. The technology-neutral and risk-based approach is a key strength of the framework, as it provides a common set of governance principles that can be applied across different AI use cases while remaining flexible enough to accommodate future technological developments. The report appropriately recognises the importance of proportionality and avoids imposing overly prescriptive requirements that could quickly become outdated.

We support the report's recognition that not all AI applications carry the same level of risk and that governance expectations should reflect factors such as materiality, complexity, autonomy, explainability, customer impact and operational significance. A proportionate approach is particularly important to ensure that lower-risk internal applications are not subject to the same level of controls as AI systems used in customer-facing activities, financial decision-making or other critical business processes.

The report also appropriately acknowledges risks that are more specific to GenAI and agentic AI, including hallucinations, prompt injection, data leakage, model drift, autonomous actions and challenges related to human oversight and accountability. However, as AI systems become more autonomous, some risks may require further consideration. Agentic AI systems can coordinate and execute multi-step actions across systems and databases, creating new risks related to workflow complexity, compounding errors, unintended outcomes, authority limits and operational resilience. For this reason, the final report could provide clearer worked examples or a companion table showing how implementation expectations differ between traditional machine learning, GenAI and agentic AI. Such

guidance would be particularly useful in relation to human oversight, explainability, performance monitoring, documentation, audit trails, authority limits and containment mechanisms.

Furthermore, human oversight remains a critical safeguard, but the report could recognise the practical limitations of applying traditional oversight mechanisms to increasingly autonomous systems. For highly automated or agentic applications, continuous real-time human monitoring may not always be feasible, and institutions may need alternative controls such as approval thresholds, escalation mechanisms, defined authority limits, interruption capabilities and robust monitoring frameworks. Oversight expectations should evolve in line with the degree of autonomy of the system, while continuing to prioritise foundational risk management concerns such as data governance, cybersecurity and third-party risk management, which remain relevant across all forms of AI and represent the most immediate implementation priorities for financial institutions.

Nonetheless, governance expectations related to emerging forms of AI should remain proportionate to actual deployment. Institutions that are not developing or deploying highly autonomous AI systems should not be expected to implement governance arrangements designed primarily for frontier AI technologies. The level of governance and control should reflect not only the materiality of the use case, but also the institution's technological readiness and the maturity of the AI capabilities being deployed. This would support a gradual evolution of governance as institutions progress from simpler analytical models to more autonomous AI applications. That is why, banks should be encouraged to evaluate emerging AI capabilities through controlled pilots - such as supervised sandboxes, testing environments and phased deployment before integrating them into material business processes. Such an approach would allow institutions to build experience while maintaining appropriate financial stability safeguards

WSBI-ESBG also believes additional guidance would be helpful on topics such as

1. Explainability for autonomous decision-making (including prompts and output governance);
2. Authority boundaries for AI agents;
3. Testing of agentic systems in controlled environments, as the final report should encourage institutions to establish suspension, rollback or shutdown mechanisms to stop, contain or reverse an AI system where it behaves unexpectedly, produces unreliable outputs, creates customer harm or generates operational risk;
4. Management of third-party foundation models, institutions relying on vendor-provided foundation models, commercial AI platforms or cloud-based AI services, should have contractual and operational arrangements to manage model changes, version updates, performance degradation, changes in data usage practices, service discontinuity and concentration risk
5. Appropriate disclosure when customers interact directly with AI systems, such as chatbots, virtual assistants or AI-supported advisory tools, including easy escalation to human staff, and clear complaint-handling and recourse arrangements. Customers should

not be disadvantaged because an AI system was involved in the interaction or decision-making process

Overall, WSBI-ESBG considers that the report achieves an appropriate balance and provides a sound foundation for the responsible adoption of both traditional and emerging forms of AI. At the same time, further practical guidance on the governance and oversight of highly autonomous systems would help institutions implement the framework consistently as AI technologies continue to evolve.

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

WSBI-ESBG agrees that the sound practices are sufficiently flexible to accommodate newer types of AI and support responsible adoption over time. Their principles-based, technology-neutral and risk-based design provides a durable framework that can adapt to evolving technologies without requiring constant revision. The emphasis on proportionality is particularly important, as it allows institutions to tailor governance arrangements according to their size, complexity, business model, risk profile and level of AI maturity.

We also welcome the fact that the framework does not prescribe a single governance model and can be applied through different organisational structures. This flexibility is especially valuable for smaller, specialised and inclusion-focused institutions, which may rely on existing governance, risk and compliance functions rather than dedicated AI governance teams. The framework therefore supports responsible innovation while avoiding unnecessary compliance burdens. Nonetheless, flexibility should be accompanied by clear minimum governance expectations. Institutions should be encouraged to maintain AI inventories, review governance frameworks periodically, monitor emerging risks and ensure appropriate oversight of material AI use cases. Regular engagement between regulators, financial institutions and technology providers will also be important to ensure that governance practices evolve alongside AI capabilities.

WSBI-ESBG also believes that the framework could be further strengthened by explicitly encouraging horizon scanning and ongoing monitoring of technological developments, emerging risks and evolving regulatory expectations. Responsible AI adoption should not be treated as a static control exercise, but as an ongoing process of reassessing risks, governance structures, organisational capabilities, control effectiveness and technological developments. Organisational adaptability however, should not be interpreted solely as building internal capabilities. Many institutions, particularly smaller and mid-sized banks, will achieve effective governance through collaboration with industry associations, shared service providers, technology partners, academic institutions... Recognising these collaborative approaches would support proportionality and encourage broader responsible AI adoption across the financial sector

Finally, the report could more explicitly reflect broader consumer and societal consideration for AI adoption, including staff training, costumer awareness of AI, environmental consideration and data sovereignty issues. Incorporating these dimensions would help ensure that responsible AI adoption remains aligned with broader public policy objectives and the public trust, key for savings and retail banks.

Overall, we consider the sound practices to be sufficiently flexible and future oriented. Their principles-based approach provides a strong foundation for responsible AI adoption, while ongoing review, horizon scanning and practical guidance will help ensure continued relevance as AI technologies and use cases evolve.

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 provide useful and practical illustrations of how responsible AI adoption can deliver measurable benefits in areas such as fraud detection, risk management, customer service, operational efficiency, compliance and cybersecurity. They demonstrate the importance of strong governance, human oversight and risk management, while showing that AI can generate tangible business and customer outcomes when implemented responsibly.

However, the current selection of case studies is weighted towards larger and more technologically advanced institutions. While these examples are informative, they do not fully reflect the diversity of institutions represented across the financial sector, including smaller and mid-sized banks, savings and retail banks, postal banks, cooperative banks, microfinance institutions, payment institutions... Additional examples reflecting a broader range of business models, institutional sizes and levels of AI maturity would improve the report's practical relevance, as these institutions often have different operating models, customer bases, resource constraints and technology strategies compared with large international banks, and may rely more heavily on third-party solutions or collaborative governance arrangements.

In particular, the report would benefit from standardised formats for the case studies and more examples involving retail, savings and inclusion-focused institutions. AI can play an important role in expanding access to financial services, improving digital financial literacy and enabling more personalised financial guidance. Case studies showing how AI can support financial inclusion would help demonstrate that responsible AI adoption can contribute not only to efficiency and risk management, but also to broader public policy objectives, including access, affordability and customer protection.

Additional case studies could also include:

- Examples of collaborative governance models, covering shared model validation capabilities, pooled third-party assessments, joint AI governance frameworks, industry-wide monitoring utilities, fintech partnerships, supervisory sandboxes or sector-wide approaches to AI testing.
- Practical banking case studies, including the use of AI by savings or retail banks for deposit behaviour and liquidity forecasting; early warning indicators in retail loan portfolios; customer complaint classification and conduct risk monitoring; regulatory reporting quality checks; AML/CFT alert optimisation; internal audit planning; and branch or channel productivity analytics.

- Examples where governance prevented or mitigated adverse outcomes, anonymous “near miss” cases involving biased models, third-party failures, unexpected model behaviour, cybersecurity incidents, weak audit trails or customer complaints, which could provide highly transferable lessons
- Additional case studies related to emerging markets, which would enrich the report by illustrating how banks have adopted AI under different operational, regulatory and resource constraints. Such examples should demonstrate that responsible AI adoption is achievable across a wide variety of institutional and market contexts
- Examples of supervisory initiatives, such as innovation hubs, supervisory sandboxes, capacity-building programs or public-private collaboration initiatives, which would provide useful insights into how supervisory frameworks can evolve alongside AI adoption

Overall, the existing case studies are useful and relevant, but a broader range of examples covering smaller institutions, non-bank financial institutions, financial inclusion use cases and practical implementation experiences would make the report more representative of the diversity of the global financial sector and enhance its value as a practical implementation resource.

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

WSBI-ESBG considers that the case studies provide useful and actionable insights for financial institutions seeking to adopt AI responsibly. They demonstrate how AI can be applied in practical business and control environments, highlight the importance of governance and human oversight, and show that responsible AI adoption can generate tangible benefits in areas such as fraud detection, risk management, compliance, customer service and operational efficiency.

We welcome the fact that the case studies go beyond theoretical principles and illustrate how governance frameworks can support real-world implementation. They help demonstrate that AI can be deployed not only in customer-facing applications, but also in internal control functions, documentation processes, compliance activities, cybersecurity, operational resilience, knowledge management and management decision-making. This makes the report relevant for institutions at different stages of AI maturity and with different strategic priorities

However, the practical value of the case studies could be further enhanced through additional implementation detail. Institutions would benefit from greater insight into how AI use cases were governed, approved, tested, monitored and integrated into existing risk management frameworks. More information on governance arrangements, materiality assessments, validation approaches, performance metrics and human oversight mechanisms would help institutions translate the examples into their own implementation plans. Indeed, the challenge is not to replicate sophisticated AI implementation experiences, but to understand how to start this journey in a gradual, realistic way. Therefore, the case studies would provide greater practical value for banks in our region if they described not

only the final AI implementation, but also the institution's adoption journey. Illustrating how governance, organizational capabilities and internal controls evolved over time would help banks at earlier stages of AI adoption understand how responsible implementation can be achieved progressively. Each case study should also clearly identify the sound practices it illustrates. A simple mapping between the case study and the relevant governance principles would help readers understand how the twelve sound practices translate into operational implementation.

In addition, not all banks will start their journey with GenAI. The report could therefore differentiate case studies according to the maturity of the institution and the complexity of the AI use case, which would enable interested banks to identify examples that are more relevant to their own stage of AI adoption and facilitate a more proportionate implementation approach

Furthermore, WSBI-ESBG believes there needs to be greater transparency regarding governance and accountability arrangements. Additional information on the roles of the Board, senior management, risk functions, compliance teams, internal audit and technology teams would help institutions understand how responsibilities were allocated in practice. Particular interest was expressed in how institutions organised cross-functional coordination, documentation standards, reporting arrangements, oversight mechanisms and the application of the three lines of defence

WSBI-ESBG also believes there would be value in including implementation challenges and lessons learned. Practical examples of how institutions addressed issues such as data quality limitations, legacy system integration, skills shortages, third-party dependencies and organisational change would provide particularly useful guidance for institutions at earlier stages of AI adoption. Examples where governance frameworks helped prevent or mitigate adverse outcomes could be as valuable as examples of successful implementation. This could be coupled with practical examples of dialogue between financial institutions and supervisory authorities during implementation phases, to provide insights into how institutions and regulators can develop shared understandings and best practices.

Finally, distinguishing more clearly between low, medium and high-risk AI use cases, and including practical checklists or key takeaways for each case study, would further support proportional implementation across institutions with different levels of AI maturity. Continued collection and sharing of case studies over time would also help institutions learn from evolving industry practices.

Overall, we consider the case studies to be relevant, practical and informative. Their value could be further strengthened through more detailed implementation information, standardised formats, clearer links to the sound practices, lessons learned, and a broader focus on governance and operational challenges encountered during deployment.

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

WSBI-ESBG believes that the glossary provides a useful and broadly aligned foundation for a common understanding of AI governance and risk management. The definitions are generally clear, reflect current industry terminology, and are consistent with concepts widely

used by financial institutions, regulators and international standard-setting bodies. The glossary helps establish a common language across business, risk, compliance, technology and audit functions, supporting more consistent implementation of the sound practices.

However, given the rapid evolution of AI technologies, several definitions could be expanded or clarified to improve their practical applicability. In particular, there needs greater precision around emerging concepts such as agentic AI, foundation models, model drift, AI assurance, explainability, AI supply chains and human oversight mechanisms. Periodic updates to the glossary will be important to ensure continued alignment with technological developments and evolving industry practices.

Agentic AI:

Indeed, the glossary would benefit from a more detailed treatment of agentic AI. Rather than defining agentic AI solely as systems capable of autonomous action, the glossary could distinguish between different levels of autonomy, ranging from assistive and semi-autonomous systems to fully autonomous agents capable of planning, coordinating and executing multi-step tasks. Clearer distinctions between single-agent and multi-agent systems would also support more consistent governance and oversight approaches.

Generative AI:

This could also be expanded to GenAI and associated risks. Relevant concepts include hallucinations, prompt injection, retrieval-augmented generation (RAG), prompt engineering, synthetic data and foundation models. Clarifying these terms would help institutions better understand the risks and controls associated with modern AI systems, particularly in customer-facing and decision-support applications.

Human oversight:

Further clarification of human oversight concepts would also be beneficial. In practice, institutions may employ different approaches, such as human-in-the-loop, where human approval is required before a critical decision or action is executed; human-on-the-loop, where an AI system operates under monitoring with the possibility of human intervention; and human-in-command, where humans retain strategic responsibility and accountability even where operational processes are automated.

AI Governance:

The glossary could include or further develop definitions for concepts such as shadow AI, AI incidents, contestability, model cards, system cards, bias, data drift and material AI use cases. These concepts are increasingly relevant to AI governance in financial institutions and would help clarify expectations relating to accountability, customer protection, operational resilience and risk management.

Material AI use cases:

WSBI-ESBG believes that additional clarification around the concept of material AI use cases could improve consistency of implementation. Illustrative criteria such as customer

impact, operational criticality, financial significance or decision autonomy could assist institutions in applying governance expectations proportionately

Overall, the glossary should remain technology-neutral while also facilitating interoperability with existing international and national frameworks. Where appropriate, definitions could be aligned with terminology used by other international organisations and supervisory authorities to reduce ambiguity and support consistent implementation across jurisdictions. ESBI-ESBG also considers that the glossary should be understandable not only by AI specialists, but also for board, senior management, risk functions and supervisory authorities. Clear and operational definitions will support more consistent implementation and supervisory dialogue across jurisdictions with different levels of technical expertise.

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

WSBI-ESBG warmly welcomes the FSB's initiative to develop internationally consistent sound practices for the responsible adoption of AI in the financial sector. The report provides a balanced, principles-based and risk-oriented framework that appropriately seeks to support innovation while reinforcing governance, operational resilience, consumer protection and financial stability. The sound practices are comprehensive, adaptable and relevant to a broad range of financial institutions and jurisdictions.

We strongly support the report's technology-neutral and proportionate approach. Financial institutions differ significantly in terms of size, complexity, business model, AI maturity and available resources. Maintaining flexibility will therefore be essential to ensure that responsible AI adoption remains achievable for a wide range of institutions, including smaller, specialised and inclusion-focused institutions, while avoiding unnecessary compliance burdens or the concentration of AI capabilities among a limited number of large market participants.

As the report moves towards finalisation, additional practical implementation support could further enhance its usefulness. Examples of AI inventories, governance structures, risk classification methodologies, reporting templates and implementation guidance would help institutions translate high-level principles into operational practices. Such tools would be particularly valuable for institutions that are still building governance capabilities or adopting AI gradually.

We also encourage greater recognition of financial inclusion as a strategic opportunity arising from responsible AI adoption. AI has the potential to improve accessibility, affordability and the quality of financial services, expand access to underserved populations and strengthen financial literacy. Responsible AI governance should therefore not only focus on mitigating risks but also support socially beneficial innovation and broader public policy objectives.

Finally, we encourage continued coordination with existing international frameworks and supervisory initiatives to promote consistency and reduce fragmentation. Responsible AI adoption is an evolving journey, and regular dialogue between regulators, financial institutions and technology providers will be essential to ensure that governance frameworks remain effective, adaptable and supportive of innovation over time.