FSB FINANCIAL STABILITY BOARD

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

Response to Consultation

Institute of International Finance (IIF)

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?

• Overall, we consider the report to have given a fair and balanced treatment of the benefits and risks of AI adoption by FIs. IIF members particularly welcome the recognition that not adopting AI may itself create risks, as this helps in providing a more balanced perspective on AI adoption; they would also call out that the final report could explicitly note that the use of AI for defensive purposes (e.g. cybersecurity, AML and fraud/scams prevention) will become increasingly relevant. Therefore, in addition to the risk of not adopting AI, rules that constrain or slow the use of AI can create an advantage gap for malicious actors who are not subject to such rules.

• As the FSB rightly identifies in the Consultation, AI is already delivering benefits across all finance industries . FIs are increasingly realizing benefits by using AI to automate repetitive and data-intensive processes, such as document analysis, transaction monitoring, compliance reviews, and customer service . These capabilities reduce operational costs, improve processing times, and enable employees to focus on higher-value activities that require judgment and expertise. Additionally, AI is increasingly being used to strengthen governance, risk and control functions, including automated control testing, policy compliance monitoring, model oversight, code review, and operational resilience monitoring. AI also strengthens firms' ability to triage and prioritize vulnerabilities, anomaly detection, fraud prevention, combat financial crime, and identify emerging risks by analyzing vast volumes of structured and unstructured data in real-time.

Beyond all of the above, AI is also helping FIs to (1) deliver better outcomes for customers and support innovation – enabling more personalized products and services, faster responses to customer inquiries, refined credit and risk assessments, improved financial inclusion, and more informed business decisions; and (2) enhance scalability, timeliness, and consistency by reducing false positives and false negatives in risk management processes, improving both customer experience and operational resilience beyond what is achievable through exclusively manual processes .

• The FSB recognizes that advances in AI models and coding agents are accelerating the discovery and exploitation of security vulnerabilities. As this capability continues to advance, the ability of firms and authorities to respond to the same underlying threat can be

constrained by the fact that vulnerability disclosure timelines, incident reporting requirements, and cross-border information sharing all depend on the legal and regulatory frameworks applicable in each jurisdiction, and cannot be assumed to be uniform or immediately available across borders. This is consistent with concerns the IIF has raised previously, which have highlighted that differing reporting requirements, definitions, and timeframes across jurisdictions can slow the ability of firms and authorities to respond to threats and have called for more consistent, bidirectional information sharing between authorities and financial institutions . Members would characterize the dynamic described here as a case for the same kind of coordination, rather than as grounds for a higher compliance burden on individual FIs.

Relatedly, the IIF would emphasize that the report's references to identifying and monitoring vulnerabilities, including in Sound Practice 11 (cyber and ICT risk management), are not to be read as a precursor to a future vulnerability reporting obligation specific to AI. Vulnerability reporting is conceptually distinct from, and generally more sensitive than, incident reporting as it concerns weaknesses that have not yet been exploited. A low reporting threshold could itself create risk, whether by aggregating information on unremediated vulnerabilities in one place and/or by generating a volume of disclosures that outpaces the ability of firms and authorities to act on them, particularly given the Consultation's own recognition that AI can accelerate vulnerability discovery.

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

• Overall, the IIF considers the sound practices to be sufficiently comprehensive and clear. FIs are already responsibly adopting AI, and AI is already strictly regulated by regimes including privacy, data governance, and operational risk regimes. As such, it is not a question of "enabling" responsible AI adoption as much as encouraging alignment between FI supervisors, and of fostering those sound practices over time.

• The IIF welcomes that in Sound Practice 8 (explainability and transparency) the FSB recognizes that there are different levels of explainability offered by different types of AI systems, and that approaches to explainability should be proportionate to the risks and materiality of the use cases and the type of AI being used. This is consistent with what the IIF has heard from its members , and the importance of reflecting proportionality mentioned earlier in this Consultation response.

The IIF would, however, suggest that Sound Practice 10 not be characterized as "human oversight" (emphasis added) to avoid confusion. Instead, it could be characterized as "human accountability and effective oversight". To explain, as the FSB rightfully observes in the description of this sound practice, oversight of AI can take different forms, some of which do not necessarily need to have direct oversight from a human, particularly in lower-risk use cases. Implying that all the forms of oversight need to be "human" could result in firms directing humans to perform low-value checking work in scenarios where it is not truly required nor value added, instead of mobilizing those resources for more critical tasks and riskier use cases . Indeed, it may be the case that over time, AI systems are found to be superior to humans in performing various oversight functions where long periods can elapse between meaningful events. Denoting that all oversight should recur to a human ignores that technology solutions – such as "validation agents", "LLM-as-a-Judge" and contrasting

outcomes against alternate models – could be sufficient, even more efficient, and/or even more consistent in their assessment of AI , particularly for lower-risk use cases. The sound practices could more clearly recognize that such oversight exists, while acknowledging that ultimately it is human beings – on behalf of the entities that employ them – that are responsible.

Considering that the Consultation is focused on the practices that the industry is using, the IIF would highlight that AI systems used to automate or support the governance of AI models is, in fact, a popular practice among FIs. Data from our upcoming annual survey on AI in financial services shows that, as of H1 2026, 81% of surveyed FIs leverage AI systems to automate or support the governance of AI models. This can be read as a direct response to a trend that the IIF is seeing: FIs are decisively moving to strategic execution and to scaling AI across the firm, which is resulting in an increase in inventories across all types of AI, pushing FIs to find efficient and secure ways to govern more AI applications than before , including by incorporating thoughtful performance metrics that, for example, reward the quality and assertiveness of oversight decisions made by AI systems.

• When assessing Sound Practice 2 (governance and accountability), the FSB states that "It is essential that audit, evaluation, risk management, and testing remain independent and functionally separate from development teams to maintain objectivity" (emphasis added). Though the IIF agrees that having lines of defense or other types of structures for defining clear roles and responsibilities for AI adoption is a sound practice for governing AI, it is worth clarifying that the separation of governance and development teams is not essential across all use cases.

Independence across different functions that participate in developing, testing, and launching an AI model can be useful. That said, helpfully, some industry participants are creating clusters of employees across business, technology and risk functions, with the purpose of streamlining the production of some value-added use cases. These examples suggest that it is best that risk professionals not become absorbed by developers of use cases or by business functions, but rather that the risk lens is present throughout the process of developing, testing and releasing some AI applications. An example of this approach, and in line with the proportionality discussed in the Consultation, is seen in some instances where firms embed the first line of defense across business and technology functions for some low-risk use cases, while keeping development and deployment assurances subject to independent validation, with internal audit providing independent third-line review. The FSB's final report could usefully acknowledge this operating-model flexibility without diluting the underlying independence, which remains central to established three-lines-of-defense expectations across major jurisdictions.

• Other governance mechanisms currently used by FIs include building upon existing model risk management concepts to expand post-deployment governance for frontier AI systems that are capable of adaptive behavior, tool use, or autonomous task execution. This type of governance includes ongoing monitoring, observability, auditability and runtime controls.

Some IIF members also noted that, for customer-facing use cases, AI governance increasingly incorporates conduct-risk considerations alongside other operational risks. This includes assessing whether AI-enabled interactions and decisions remain consistent with

applicable consumer protection and customer outcome expectations. Such mechanisms complement transparency and accountability measures while supporting customer confidence.

• When referring to Sound Practice 12 (third-party AI risk management), the FSB states that FIs could "continue providing a given AI-enabled service manually". However, this could prove unrealistic for high-volume automation and agentic AI deployed at scale, where the manual capacity to replace AI is understandably not held in reserve. This affirmation also sits in tension with the Consultation's own acknowledgment that rollbacks "can be challenging (or impossible) for some forms of AI models and systems". Instead, IIF members see that exit strategies for such use cases currently rest on substitutability (i.e., data portability, alternative providers, assured data access, etc.), which could be clearly recognized in the final report. This adjustment to the report merits attention by the FSB.

• Consistent with the overarching call for proportionality throughout the Consultation, it would be appropriate to understand Sound Practice 1 (strategic direction and oversight) in a proportionate manner that responds to the materiality, complexity and risk profile of the relevant AI use cases. This is determined by institutions within their broader governance and risk management frameworks – which include considerations around the involvement of governance and oversight bodies –, not by granular board-level monitoring or escalation of all AI applications. Such an approach could create operational inefficiencies and slow innovation without delivering commensurate risk management benefits, in addition to disproportionately affecting smaller institutions.

Relatedly, the Consultation seems to treat the terms "board" and "senior management" as synonyms or interchangeable concepts, though that is not reflective of current practices in the industry and could result in nudging board members to focus less on strategic business decisions and more on details about AI deployments that would be best addressed by others.

The same idea applies to AI literacy expectations. Appropriate awareness and understanding of AI-related opportunities, risks, and controls is important. However, these expectations do not necessarily need to be interpreted as requiring all board members to have professional AI technical expertise. Instead, FIs have flexibility to determine the appropriate allocation of expertise, support, and challenge functions within their governance structures. Relatedly, FIs are actively considering good practices for the maintenance of knowledge that could inadvertently get lost as staff uses AI; this pertains to practices to maintain knowledge about AI governance and validation, and also to knowledge about the business and operations that increasingly rely on AI.

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?**

• The IIF commends FSB efforts to study and recognize many of the nuances and difficulties that FIs face when engaging with AI developers and third-party providers. Overall, the IIF believes the sound practices to be appropriately balanced between all forms of AI and specific more complex forms. However, the Consultation also appears to assume in some sections that FIs have clear visibility into third-party data or models, or that they can

secure that visibility through contractual negotiations. Industry experience shows that this is not necessarily the case for the majority of firms.

According to results from the IIF's 2025 Annual Survey on AI Use in Financial Services, 58% of surveyed firms across the globe require the same level of validation for internally developed models as they do for externally developed ones. However, in practice, validation is often less thorough for third-party models due to limitations in accessing relevant information . In the as yet unpublished results of the 2026 edition of this survey, the IIF finds that more than half of surveyed FIs are still unable to perform equal validation for internally and externally developed models. Furthermore, results from this year's survey show that most surveyed FIs face information and validation constraints in four key areas when engaging with third parties: (1) limited access to model architecture or engineering details; (2) restricted insights into model updates or retraining processes; (3) limited access to training data; and (4) the inability to independently test model performance outcomes.

This challenge compounds further down the supply chain, as visibility into fourth, fifth, or nth-party vendors and open-source components is frequently outside any individual firm's reasonable ability to obtain. Further challenges stem from third-party providers increasingly deploying new AI capabilities, models updates, or autonomous functionalities into existing products and services (also known as "embedded AI") without prior notification to their clients. This further complicates established model risk management, creates challenges for explainability, transparency and auditability, and prevents FIs from accessing critical information in a timely manner.

Therefore, the IIF suggests that the sound practices explicitly recognize that firms' ability to implement certain governance practices may be constrained by the level of information shared by third-party providers, and that expectations should be interpreted in light of those practical limitations. Concurrently, further clarifying that responsible governance of third-party AI systems may need to rely on proportionate "compensating controls" and scalable assurance mechanisms would be welcomed. In this sense, a stronger recognition of the crucial role that AI model providers play in enabling effective risk management downstream – including through provider transparency and cooperation – could help establish a common industry expectation or baseline regarding the information that providers should make available to enable more effective third-party AI risk management.

In particular, the FSB could encourage tackling these challenges on two fronts related to third parties: data, and the models themselves. On the data front, these issues could be tackled by encouraging standardization or industry practices for third-party use of data in the context of AI. For example, development of "nutritional labels", standardized "data cards", or equivalent disclosures that provide FIs with sufficient information necessary to assess data quality and AI model performance would be beneficial.

In order to tackle some of the challenges observed in the models themselves, the IIF welcomes that the Consultation recognizes that whenever full model transparency from third-party providers is not achievable, institutions are relying on "compensating controls" – which are explicitly recognized in the Consultation and by the industry as well , and on standardized controls, such as model/system cards, certifications, or audit artifacts – which could be more broadly recognized in the Consultation. Other practices helping to govern third-party models without dramatically increasing the costs of governing the use of the

technology include defining Minimum Viable Controls (MVCs) for third-party models, including baseline disclosure expectations, and clear documentation standards . The IIF suggests recognizing these practices in the final version of the commented document.

IIF members welcome the report's recognition that risks posed by AI adoption are not all new; that said, further articulation of the AI-specific dimension of a risk may help clarify where existing risk management approaches may need to be adopted. This is again relevant in the context of third-party risk management. While the report appropriately recognizes that existing regulation and guidance on TPRM provide a suitable foundation, it could go further in explicitly recognize that existing TPRM frameworks and practices remain the appropriate foundation for managing third-party AI risks, with targeted adaptations where AI introduces genuinely distinct characteristics.

- IIF members are also broadly supportive of the AI Shared Responsibility model. This type of model would be useful in delineating accountability among AI vendors, their upstream suppliers, and the FIs that deploy the technology. This model also formally acknowledges that FIs cannot and should not be solely responsible for risks embedded deep within the AI supply chain.

- The report highlights the various AI-enabled types of attacks that can be facilitated by using AI techniques. However, it is worth noting that (1) these additional causes of cyber risk do not change the cyber security framework itself, but rather require additional skills and scrutiny on these types of risks; and (2) existing regulatory frameworks in various jurisdictions already allow to contain these risks.

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

- The sound practices are sufficiently flexible at this point. However, wording that may be interpreted as endorsing specific implementation approaches could be softened. For example, expressions such as "typically includes" could be replaced with more flexible language such as "could include actions such as".

- IIF members commend the FSB for highlighting the importance of public-private partnerships for better understanding the most advanced AI systems being used, for at least 3 reasons: (1) structured public-private dialogue helps ensure that regulatory objectives are achieved efficiently and without creating unnecessary barriers to innovation, competition, or financial inclusion; (2) effective policymaking benefits from continuous engagement between authorities and industry participants, allowing regulators to access technical expertise while developing proportionate, evidence-based frameworks; and (3) many financial services are inherently cross-border; the same happens with general-purpose technologies like AI. In these cases, cooperation between authorities and industry helps reduce fragmentation and promote greater consistency across jurisdictions, lowering costs while enhancing financial stability and market access.

- The IIF welcomes the FSB's interest in revisiting these practices from time to time to account for changes in technology and the accelerating pace of development of AI models.

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.**

No specific comments.

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

• Overall, the case studies provide useful illustrative examples of how sound practices can be implemented in practice. The IIF encourages the FSB to clearly present the case studies as illustrative examples rather than preferred implementation models, consistent with the FSB's stated intention to provide a non-binding menu of sound practices. This clarity will help preserve proportionality and reduce the risk of creating de facto supervisory expectations around specific tools, governance structures, or control arrangements if that is indeed not the intention of the FSB at this time.

Members welcome the direction of the report's recognizing that FIs incorporate AI cyber and ICT risk scenarios into their testing and exercises. At the same time, members have observed a broader pattern in supervisory conversations of scenario testing pushing toward increasingly extreme, low probability events, which in practice pulls investment and attention away from the patching, access management, and fundamental cyber hygiene measures that address the far more common, lower severity incidents firms are most likely to experience. IIF membership encourages the FSB to acknowledge this balance directly in its final report, and to note that such examples illustrate current practice without setting an expected model.

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

• As a global standard-setting body, the FSB is positioned to address regulatory fragmentation on key AI concepts. Some IIF members have suggested adjustments to some of the proposed definitions in the Consultation, such as:

o Materiality – the report uses this term extensively but does not define it in the Glossary (materiality and risk help drive the calibration of risk management requirements across all 12 sound practices). The IIF would appreciate if regulators recognize that IIF members may already consider, among other things, the following criteria to assess the materiality and risk dimensions in specific cases: adaptability, memory, orchestration ability, customer impact, financial and regulatory exposure, data sensitivity, scope of access to systems or data, reversibility, criticality to the institution/service continuity (e.g., use in critical or material functions/services), degree of substitutability, scale of usage, adversarial exposure, and degree of autonomy of action.

o Agentic AI – would be beneficial to include more specificity in the definition, including the following 3 components: (1) goal directed autonomy (including ability to plan and reason); (2) invokes other tools and systems to carry out a task; and (3) execute an action with real impact.

o Traditional AI – traditional AI does not include GenAI and beyond. It would also be beneficial to clarify that traditional statistical techniques, such as logistic regression, Generalized Linear Models (GLMs), standard scorecards or rules-based decision trees do not automatically fall within the AI perimeter when they are transparent, deterministic, or manually specified.

o Shadow AI – Consistent with the Consultation, this could be defined as "when employees adopt AI tools without authorization or for unapproved uses" (page 12).

o Concentration Risk – the report discusses this concept extensively in section 3.9 but does not provide a standalone definition. Clarifying whether this refers to provider concentration, model concentration, or both, would be useful, particularly since Sound Practice 12 appears to use the term to describe a financial institution's own exposure to a limited number of providers, while Section 3.9 separately touches on sector-wide or systemic concentration as a matter of financial stability risk. The IIF would suggest the glossary distinguish between these two uses so they are not read as interchangeable.

• A definition is a "statement of the meaning of a word, phrase or symbol" . However, some of the definitions proposed by the FSB go beyond that exercise, adding explanations about the reasons behind a certain concept, or adding some of the potential consequences of using specific AI models.

For example, while defining "foundation models", the FSB states that this is "An umbrella term referring to a diversity of models that are usually trained by applying deep learning to massive quantities of data, such as text and images. Because the expertise, time, and computing power involved in training foundation models from scratch are typically prohibitive for most non-specialist firms, these models are usually pre-trained and shared with end-users for further use and refinement" (emphasis added). In this case, the text in bold provides details that try to explain why foundation models are usually pretrained, but while doing so, it makes it seem as foundation models need to be, by definition, expensive in "terms of expertise, time and computing power", which is not necessarily the case for all foundation models at present or future times.

Also, the FSB defines "open-source models" as "AI models where the full training code, and in some cases the training data or its composition, is made publicly available for use and modification. These models offer customizability and reduce vendor lock-in but may introduce additional risks, such as security vulnerabilities, data privacy, or data quality concerns" (emphasis added). In this case, the bold text assumes that by using "open-source models", FIs would be able to reduce vendor lock-in, though this is not necessarily a consequence, nor a reason, for recurring to these models every time an FI does so.

• The FSB uses the concept of an "AI incident" in Sound Practices 4 and 11, in the context of post-incident reviews and information sharing. The IIF would ask the FSB to state plainly that references to AI incidents in the report are not intended to establish a new, freestanding reporting category, and that where an incident involving AI is, in substance, a cyber or ICT incident, it continues to be reported through a financial institution's existing cyber and ICT incident frameworks.

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

IIF's Cover Letter:

Dear Dr Schindler,

Sound Practices for Responsible Adoption of Artificial Intelligence – Response to the Financial Stability Board's Consultation Report

The Institute of International Finance (IIF) welcomes the opportunity to publicly respond to the Financial Stability Board's (FSB's) public consultation (Consultation) concerning the above topic, launched on June 10, 2026.

The IIF represents approximately 400 globally active financial institutions (FIs) from over 60 geographies, drawn from the banking, insurance, securities, asset management, payments and other sectors, including central and development banks. Many of our members are particularly interested in the cross-border impacts of Artificial Intelligence (AI), and related frameworks.

The IIF welcomes the FSB's objective of providing a menu of non-binding sound practices that are principles- and outcomes-based, proportionate, and technology neutral. The IIF also commends the FSB for explicitly stating in the Consultation that "the sound practices are not intended to establish an international standard, to impose a prescriptive approach for responsible AI adoption by financial institutions, nor to influence business decisions in adopting a certain AI technology". These parameters reflect that the FSB is focused on sound practices for the industry, rather than on more prescriptive alternatives such as principles or recommendations.

To guarantee that the Consultation continues to reflect its objectives, the FSB is encouraged to more clearly present the case studies throughout the Consultation as illustrative examples, which would be interpreted differently than preferred implementation models. This would be consistent with the FSB's stated intention to provide a non-binding menu of sound practices, rather than a prescriptive international standard. This added clarity would also play an important role in preserving proportionality and avoid creating de facto supervisory expectations around specific tools, governance models or control arrangements. For example, case studies describing dedicated AI inventories, specific documentation practices, governance structures, or cyber frameworks may be appropriate for some institutions but should not be interpreted as the FSB's preferred implementation model at this time. Similarly, references throughout the report to specific technical frameworks, methodologies or control practices could be clearly presented as illustrative examples of possible alternatives that market participants can use, rather than preferred or expected supervisory practices.

IIF members also agree that responsible adoption of AI can be achieved most effectively by integrating AI governance into existing firm-wide governance and risk management frameworks, rather than by creating parallel AI-specific regimes, rules, or duplicative obligations. Additionally, IIF members view responsible AI governance as not merely a tool used to limit risks, but also a tool for enabling safe scaling and experimentation. The financial industry believes FSB's report would benefit from making this enabling intent more structurally visible, for example by explicitly framing governance proportionality as a tool for accelerating responsible adoption, not only for limiting harm.

As AI advances, IIF encourages continued coordination among international standard-setting bodies and national authorities to promote consistency and interoperability across jurisdictions, which will better allow FIs to focus resources on managing material risks rather than on reconciling differing regulatory expectations.

Background

As in discussions with other authorities, the IIF would briefly ground any discussion of AI in the following:

• AI in finance is already regulated: AI applications in finance are already subject to regulation through sectoral and cross-sectoral regulations. Examples include rules around data protection, privacy, AML & KYC-specific measures, bank secrecy, risk assessment, consumer/investor protection, prudential requirements concerning data governance, cyber risk, third-party risk, and operational resilience. These regulations apply to the use of any general-purpose technology, like AI, in financial services, ensuring that the risks continue to be well managed and reducing the need for parallel AI-specific regimes.

• Existing risk frameworks and industry experience provide a strong foundation: FIs operate under national/regional prudential requirements, consistent with international standards, that require firms to develop their own sophisticated risk management governance frameworks, systems, and controls. These arrangements have served the industry well. FIs keep such arrangements under constant review and proactively make adjustments as needed. The industry therefore starts from a very strong risk management foundation – one that is commensurate with and contributes to the high trust placed in FIs by consumers and investors, and builds on the long-standing experience that the industry has demonstrated in developing, deploying, and governing AI over several years.

• A thorough comprehension of the AI value chain is crucial to appropriately govern AI: The share of FIs that use third parties for AI development has increased, from 53% in 2023 to 77% in 2025 , and 87% in 2026. As FIs increasingly rely on third-party solutions to continue having access to the most advanced iterations of AI models, these actors play a central role in the AI value chain and the governance of the technology's uses. This reality highlights that responsibilities cannot be assigned uniformly, as actors at different stages of the value chain – from cloud providers and model developers to deployers and end users – have varying levels of control, visibility, and ability to mitigate risks. This is further explained by the interconnected layers that exist throughout the AI lifecycle, which span through data sourcing, model training, fine-tuning, deployment and ongoing monitoring, to name a few, each of which can present a potential entry point for governance and oversight.

• Proportionality as a critical component for appropriate governance: Proportionality recognizes that the impact, risks, and benefits of AI, along with its complexities, can vary significantly across use cases, financial sub-sectors, the provenance of the AI systems, and the specific type of AI being referred to (e.g., Machine Learning, Traditional AI, Generative AI, Agentic AI, Open-Source Models, Open-Weights Models, etc.). Governance frameworks, therefore, do well by recognizing and calibrating factors such as the materiality of the function, the potential impact on customers and markets, the degree of autonomy of the technology used, and the nature of the underlying risks of each use case. In contrast, a one-size-fits-all approach would risk either imposing unnecessary burdens on low-risk

applications or failing to provide sufficient oversight for higher-risk deployments, in addition to stifling innovation.

Specific responses to some of the questions posed in the Consultation are discussed in the Annex of this document.

Further engagement

The IIF and its members stand ready to engage in additional discussions and consultations on these topics, or to clarify any aspect of our submission. For example, if desired, we could convene a meeting of interested IIF members, and/or our AI Experts Group comprised of senior practitioners across the global financial industry, with FSB officials to discuss industry practices and this submission further.

We would also be happy to provide an early briefing to FSB officials on the findings of the 2026 IIF-EY Annual Survey on AI Use in Financial Services, whose results are not yet publicly available but are scheduled for release in Q3 2026.

We thank you again for the opportunity to contribute to this important consultation.

IIF's specific response to the 8th question:

• IIF members understand that the Consultation is focused on providing a menu of non-binding sound practices for private sector financial institutions. That said, some IIF members would note that national and regional authorities seem to be pushing for sovereign AI initiatives without taking into consideration the tradeoffs that those frameworks entail for their own policy objectives, including impacts on competitiveness, innovation, security and operational resilience.

July 22, 2026

John Schindler
Secretary General
Financial Stability Board
Basel, Switzerland

Attention:

Submitted via email, annex also submitted via the Financial Stability Board's online form.

Dear Dr Schindler,

Sound Practices for Responsible Adoption of Artificial Intelligence – Response to the Financial Stability Board's Consultation Report

The Institute of International Finance (**IIF**) welcomes the opportunity to publicly respond to the Financial Stability Board's (**FSB's**) public consultation (**Consultation**) concerning the above topic, launched on June 10, 2026.

The IIF represents approximately 400 globally active financial institutions (**FIs**) from over 60 geographies, drawn from the banking, insurance, securities, asset management, payments and other sectors, including central and development banks. Many of our members are particularly interested in the cross-border impacts of Artificial Intelligence (**AI**), and related frameworks.

The IIF welcomes the FSB's objective of providing a menu of non-binding sound practices that are principles- and outcomes-based, proportionate, and technology neutral. The IIF also commends the FSB for explicitly stating in the Consultation that "the sound practices are not intended to establish an international standard, to impose a prescriptive approach for responsible AI adoption by financial institutions, nor to influence business decisions in adopting a certain AI technology". These parameters reflect that the FSB is focused on sound practices for the industry, rather than on more prescriptive alternatives such as principles or recommendations.

To guarantee that the Consultation continues to reflect its objectives, the FSB is encouraged to more clearly present the case studies throughout the Consultation as illustrative examples, which would be interpreted differently than preferred implementation models. This would be consistent with the FSB's stated intention to provide a non-binding menu of sound practices, rather than a prescriptive international standard. This added clarity would also play an important role in preserving proportionality and avoid creating *de facto* supervisory expectations around specific tools, governance models or control arrangements. For example, case studies describing dedicated AI inventories, specific documentation practices, governance structures, or cyber frameworks may be appropriate for some institutions but should not be interpreted as the FSB's preferred implementation model at this time. Similarly, references throughout the report to specific technical frameworks, methodologies or control practices could be clearly presented as illustrative examples of possible alternatives that market participants can use, rather than preferred or expected supervisory practices.

IIF members also agree that responsible adoption of AI can be achieved most effectively by integrating AI governance into existing firm-wide governance and risk management

frameworks, rather than by creating parallel AI-specific regimes, rules, or duplicative obligations. Additionally, IIF members view responsible AI governance as not merely a tool used to limit risks, but also a tool for enabling safe scaling and experimentation. The financial industry believes FSB's report would benefit from making this enabling intent more structurally visible, for example by explicitly framing governance proportionality as a tool for accelerating responsible adoption, not only for limiting harm.

As AI advances, IIF encourages continued coordination among international standard-setting bodies and national authorities to promote consistency and interoperability across jurisdictions, which will better allow FIs to focus resources on managing material risks rather than on reconciling differing regulatory expectations.

Background

As in discussions with other authorities, the IIF would briefly ground any discussion of AI in the following:

- **AI in finance is already regulated:** AI applications in finance are already subject to regulation through sectoral and cross-sectoral regulations. Examples include rules around data protection, privacy, AML & KYC-specific measures, bank secrecy, risk assessment, consumer/investor protection, prudential requirements concerning data governance, cyber risk, third-party risk, and operational resilience. These regulations apply to the use of any general-purpose technology, like AI, in financial services, ensuring that the risks continue to be well managed and reducing the need for parallel AI-specific regimes.
- **Existing risk frameworks and industry experience provide a strong foundation:** FIs operate under national/regional prudential requirements, consistent with international standards, that require firms to develop their own sophisticated risk management governance frameworks, systems, and controls. These arrangements have served the industry well. FIs keep such arrangements under constant review and proactively make adjustments as needed. The industry therefore starts from a very strong risk management foundation – one that is commensurate with and contributes to the high trust placed in FIs by consumers and investors, and builds on the long-standing experience that the industry has demonstrated in developing, deploying, and governing AI over several years.
- **A thorough comprehension of the AI value chain is crucial to appropriately govern AI:** The share of FIs that use third parties for AI development has increased, from 53% in 2023 to 77% in 2025[1], and 87% in 2026. As FIs increasingly rely on third-party solutions to continue having access to the most advanced iterations of AI models, these actors play a central role in the AI value chain[2] and the governance of the technology's uses. This reality highlights that responsibilities cannot be assigned uniformly, as actors at different stages of the value chain – from cloud providers and model developers to deployers and end users – have varying levels of control, visibility, and ability to mitigate risks. This is further explained by the interconnected layers that exist throughout the AI lifecycle, which span through data sourcing, model training, fine-tuning, deployment and ongoing monitoring, to name a few, each of which can present a potential entry point for governance and oversight.
- **Proportionality as a critical component for appropriate governance:** Proportionality recognizes that the impact, risks, and benefits of AI, along with its complexities, can vary significantly across use cases, financial sub-sectors, the provenance of

[1] IIF – EY 2025 Annual Survey on AI Use in Financial Services. October, 2025. Page 37. Available online, here.
[2] AI, Data, and the Cloud – Connectivity and Value Creation in Finance. IIF. May 2024. Available online, here.

the AI systems, and the specific type of AI being referred to (e.g., Machine Learning, Traditional AI, Generative AI, Agentic AI, Open-Source Models, Open-Weights Models, etc.). Governance frameworks, therefore, do well by recognizing and calibrating factors such as the materiality of the function, the potential impact on customers and markets, the degree of autonomy of the technology used, and the nature of the underlying risks of each use case. In contrast, a one-size-fits-all approach would risk either imposing unnecessary burdens on low-risk applications or failing to provide sufficient oversight for higher-risk deployments, in addition to stifling innovation.

Specific responses to some of the questions posed in the Consultation are discussed in the **Annex** of this document.

Further engagement

The IIF and its members stand ready to engage in additional discussions and consultations on these topics, or to clarify any aspect of our submission. For example, if desired, we could convene a meeting of interested IIF members, and/or our AI Experts Group comprised of senior practitioners across the global financial industry, with FSB officials to discuss industry practices and this submission further.

We would also be happy to provide an early briefing to FSB officials on the findings of the 2026 IIF-EY Annual Survey on AI Use in Financial Services, whose results are not yet publicly available but are scheduled for release in Q3 2026.

We thank you again for the opportunity to contribute to this important consultation.

Yours sincerely,

Jessica Renier
Managing Director, Digital Finance

Andrés Portilla
Managing Director, Regulatory Affairs

Annex
Answers to Consultation questions

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?

IIF submission text:

- Overall, we consider the report to have given a fair and balanced treatment of the benefits and risks of AI adoption by FIs. IIF members particularly welcome the recognition that not adopting AI may itself create risks, as this helps in providing a more balanced perspective on AI adoption; they would also call out that the final report could explicitly note that the use of AI for defensive purposes (e.g. cybersecurity, AML and fraud/scams prevention) will become increasingly relevant. Therefore, in addition to the risk of not adopting AI, rules that constrain or slow the use of AI can create an advantage gap for malicious actors who are not subject to such rules.

- As the FSB rightly identifies in the Consultation, AI is already delivering benefits across all finance industries[3]. FIs are increasingly realizing benefits by using AI to automate repetitive and data-intensive processes, such as document analysis, transaction monitoring, compliance reviews, and customer service[4]. These capabilities reduce operational costs, improve processing times, and enable employees to focus on higher-value activities that require judgment and expertise. Additionally, AI is increasingly being used to strengthen governance, risk and control functions, including automated control testing, policy compliance monitoring, model oversight, code review, and operational resilience monitoring. AI also strengthens firms' ability to triage and prioritize vulnerabilities, anomaly detection, fraud prevention, combat financial crime, and identify emerging risks by analyzing vast volumes of structured and unstructured data in real-time.

  Beyond all of the above, AI is also helping FIs to (1) deliver better outcomes for customers and support innovation – enabling more personalized products and services, faster responses to customer inquiries, refined credit and risk assessments, improved financial inclusion, and more informed business decisions; and (2) enhance scalability, timeliness, and consistency by reducing false positives and false negatives in risk management processes, improving both customer experience and operational resilience beyond what is achievable through exclusively manual processes[5].

- The FSB recognizes that advances in AI models and coding agents are accelerating the discovery and exploitation of security vulnerabilities. As this capability continues to advance, the ability of firms and authorities to respond to the same underlying threat can be constrained by the fact that vulnerability disclosure timelines, incident reporting requirements, and cross-border information sharing all depend on the legal and regulatory frameworks applicable in each jurisdiction, and cannot be assumed to be uniform or immediately available across borders. This is consistent with concerns the IIF

[3] IIF's AI/ML Experts Group – Briefing Note. May 2026. Available online, here.
[4] IIF – EY 2025 Annual Survey on AI Use in Financial Services. October, 2025. Page 37. Available online, here.
[5] IIF's AI/ML Experts Group – Briefing Note. May 2026. Available online, here.

has raised previously, which have highlighted that differing reporting requirements, definitions, and timeframes across jurisdictions can slow the ability of firms and authorities to respond to threats and have called for more consistent, bidirectional information sharing between authorities and financial institutions[6]. Members would characterize the dynamic described here as a case for the same kind of coordination, rather than as grounds for a higher compliance burden on individual FIs.

Relatedly, the IIF would emphasize that the report's references to identifying and monitoring vulnerabilities, including in Sound Practice 11 (cyber and ICT risk management), are not to be read as a precursor to a future vulnerability reporting obligation specific to AI. Vulnerability reporting is conceptually distinct from, and generally more sensitive than, incident reporting as it concerns weaknesses that have not yet been exploited. A low reporting threshold could itself create risk, whether by aggregating information on unremediated vulnerabilities in one place and/or by generating a volume of disclosures that outpaces the ability of firms and authorities to act on them, particularly given the Consultation's own recognition that AI can accelerate vulnerability discovery.

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

IIF submission text:

- Overall, the IIF considers the sound practices to be sufficiently comprehensive and clear. FIs are already responsibly adopting AI, and AI is already strictly regulated by regimes including privacy, data governance, and operational risk regimes. As such, it is not a question of "enabling" responsible AI adoption as much as encouraging alignment between FI supervisors, and of fostering those sound practices over time.

- The IIF welcomes that in Sound Practice 8 (explainability and transparency) the FSB recognizes that there are different levels of explainability offered by different types of AI systems, and that approaches to explainability should be proportionate to the risks and materiality of the use cases and the type of AI being used. This is consistent with what the IIF has heard from its members[7], and the importance of reflecting proportionality mentioned earlier in this Consultation response.

  The IIF would, however, suggest that Sound Practice 10 not be characterized as "**human** oversight" (emphasis added) to avoid confusion. Instead, it could be characterized as "human accountability and effective oversight". To explain, as the FSB rightfully observes in the description of this sound practice, oversight of AI can take different forms, some of which do not necessarily need to have direct oversight from a human, particularly in lower-risk use cases. Implying that all the forms of oversight need to be "human" could result in firms directing humans to perform low-value checking work in scenarios where it is not truly required nor value added, instead of mobilizing those resources for more critical tasks and riskier use cases[8]. Indeed, it may be the case that over time, AI systems are found to be superior to humans in performing various oversight functions where long

[6] IIF Letter to the FSB on the Format for Incident Reporting Exchange (FIRE) Consultation Report. December 19, 2024. Available online, here.
[7] IIF's AI/ML Experts Group – Briefing Note. May, 2025. Available online, here.
[8] IIF's AI/ML Experts Group – Briefing Note. February 2025. Available online, here.

periods can elapse between meaningful events. Denoting that all oversight should recur to a human ignores that technology solutions – such as "validation agents", "LLM-as-a-Judge" and contrasting outcomes against alternate models – could be sufficient, even more efficient, and/or even more consistent in their assessment of AI[9], particularly for lower-risk use cases. The sound practices could more clearly recognize that such oversight exists, while acknowledging that ultimately it is human beings – on behalf of the entities that employ them – that are responsible.

Considering that the Consultation is focused on the practices that the industry is using, the IIF would highlight that AI systems used to automate or support the governance of AI models is, in fact, a popular practice among FIs. Data from our upcoming annual survey on AI in financial services shows that, as of H1 2026, 81% of surveyed FIs leverage AI systems to automate or support the governance of AI models. This can be read as a direct response to a trend that the IIF is seeing: FIs are decisively moving to strategic execution and to scaling AI across the firm, which is resulting in an increase in inventories across all types of AI, pushing FIs to find efficient and secure ways to govern more AI applications than before[10] [11], including by incorporating thoughtful performance metrics that, for example, reward the quality and assertiveness of oversight decisions made by AI systems.

- When assessing Sound Practice 2 (governance and accountability), the FSB states that "It is **essential** that audit, evaluation, risk management, and testing remain independent and functionally separate from development teams to maintain objectivity" (emphasis added). Though the IIF agrees that having lines of defense or other types of structures for defining clear roles and responsibilities for AI adoption is a sound practice for governing AI, it is worth clarifying that the separation of governance and development teams is not essential across all use cases.

  Independence across different functions that participate in developing, testing, and launching an AI model can be useful. That said, helpfully, some industry participants are creating clusters of employees across business, technology and risk functions, with the purpose of streamlining the production of some value-added use cases. These examples suggest that it is best that risk professionals not become absorbed by developers of use cases or by business functions, but rather that the risk lens is present throughout the process of developing, testing and releasing some AI applications. An example of this approach, and in line with the proportionality discussed in the Consultation, is seen in some instances where firms embed the first line of defense across business and technology functions for some low-risk use cases, while keeping development and deployment assurances subject to independent validation, with internal audit providing independent third-line review. The FSB's final report could usefully acknowledge this operating-model flexibility without diluting the underlying independence, which remains central to established three-lines-of-defense expectations across major jurisdictions.

- Other governance mechanisms currently used by FIs include building upon existing model risk management concepts to expand post-deployment governance for frontier AI systems that are capable of adaptive behavior, tool use, or autonomous task execution.

---

[9] IIF's AI/ML Experts Group – Briefing Note. February 2025. Available online, here.
[10] IIF's AI/ML Experts Group – Briefing Note. February, 2026. Available online, here.
[11] IIF – EY 2025 Annual Survey on AI Use in Financial Services. October, 2025. Page 28. Available online, here.

This type of governance includes ongoing monitoring, observability, auditability and runtime controls.

Some IIF members also noted that, for customer-facing use cases, AI governance increasingly incorporates conduct-risk considerations alongside other operational risks. This includes assessing whether AI-enabled interactions and decisions remain consistent with applicable consumer protection and customer outcome expectations. Such mechanisms complement transparency and accountability measures while supporting customer confidence.

- When referring to Sound Practice 12 (third-party AI risk management), the FSB states that FIs could "continue providing a given AI-enabled service manually". However, this could prove unrealistic for high-volume automation and agentic AI deployed at scale, where the manual capacity to replace AI is understandably not held in reserve. This affirmation also sits in tension with the Consultation's own acknowledgment that rollbacks "can be challenging (or impossible) for some forms of AI models and systems". Instead, IIF members see that exit strategies for such use cases currently rest on substitutability (i.e., data portability, alternative providers, assured data access, etc.), which could be clearly recognized in the final report. This adjustment to the report merits attention by the FSB.

- Consistent with the overarching call for proportionality throughout the Consultation, it would be appropriate to understand Sound Practice 1 (strategic direction and oversight) in a proportionate manner that responds to the materiality, complexity and risk profile of the relevant AI use cases. This is determined by institutions within their broader governance and risk management frameworks – which include considerations around the involvement of governance and oversight bodies –, not by granular board-level monitoring or escalation of all AI applications. Such an approach could create operational inefficiencies and slow innovation without delivering commensurate risk management benefits, in addition to disproportionately affecting smaller institutions.

  Relatedly, the Consultation seems to treat the terms "board" and "senior management" as synonyms or interchangeable concepts, though that is not reflective of current practices in the industry and could result in nudging board members to focus less on strategic business decisions and more on details about AI deployments that would be best addressed by others.

  The same idea applies to AI literacy expectations. Appropriate awareness and understanding of AI-related opportunities, risks, and controls is important. However, these expectations do not necessarily need to be interpreted as requiring all board members to have professional AI technical expertise. Instead, FIs have flexibility to determine the appropriate allocation of expertise, support, and challenge functions within their governance structures. Relatedly, FIs are actively considering good practices for the maintenance of knowledge that could inadvertently get lost as staff uses AI; this pertains to practices to maintain knowledge about AI governance and validation, and also to knowledge about the business and operations that increasingly rely on AI.

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?

IIF submission text:

- The IIF commends FSB efforts to study and recognize many of the nuances and difficulties that FIs face when engaging with AI developers and third-party providers. Overall, the IIF believes the sound practices to be appropriately balanced between all forms of AI and specific more complex forms. However, the Consultation also appears to assume in some sections that FIs have clear visibility into third-party data or models, or that they can secure that visibility through contractual negotiations. Industry experience shows that this is not necessarily the case for the majority of firms.

  According to results from the IIF's 2025 Annual Survey on AI Use in Financial Services, 58% of surveyed firms across the globe require the same level of validation for internally developed models as they do for externally developed ones. However, in practice, validation is often less thorough for third-party models due to limitations in accessing relevant information[12]. In the as yet unpublished results of the 2026 edition of this survey, the IIF finds that more than half of surveyed FIs are still unable to perform equal validation for internally and externally developed models. Furthermore, results from this year's survey show that most surveyed FIs face information and validation constraints in four key areas when engaging with third parties: (1) limited access to model architecture or engineering details; (2) restricted insights into model updates or retraining processes; (3) limited access to training data; and (4) the inability to independently test model performance outcomes.

  This challenge compounds further down the supply chain, as visibility into fourth, fifth, or nth-party vendors and open-source components is frequently outside any individual firm's reasonable ability to obtain. Further challenges stem from third-party providers increasingly deploying new AI capabilities, models updates, or autonomous functionalities into existing products and services (also known as "embedded AI") without prior notification to their clients. This further complicates established model risk management, creates challenges for explainability, transparency and auditability, and prevents FIs from accessing critical information in a timely manner.

  Therefore, the IIF suggests that the sound practices explicitly recognize that firms' ability to implement certain governance practices may be constrained by the level of information shared by third-party providers, and that expectations should be interpreted in light of those practical limitations. Concurrently, further clarifying that responsible governance of third-party AI systems may need to rely on proportionate "compensating controls" and scalable assurance mechanisms would be welcomed. In this sense, a stronger recognition of the crucial role that AI model providers play in enabling effective risk management downstream – including through provider transparency and cooperation – could help establish a common industry expectation or baseline regarding the information that providers should make available to enable more effective third-party AI risk management.

[12] IIF – EY 2025 Annual Survey on AI Use in Financial Services. October, 2025. Available online, here.

In particular, the FSB could encourage tackling these challenges on two fronts related to third parties: data, and the models themselves. On the data front, these issues could be tackled by encouraging standardization or industry practices for third-party use of data in the context of AI. For example, development of "nutritional labels", standardized "data cards", or equivalent disclosures that provide FIs with sufficient information necessary to assess data quality and AI model performance would be beneficial.

In order to tackle some of the challenges observed in the models themselves, the IIF welcomes that the Consultation recognizes that whenever full model transparency from third-party providers is not achievable, institutions are relying on "compensating controls" – which are explicitly recognized in the Consultation and by the industry as well[13], and on standardized controls, such as model/system cards, certifications, or audit artifacts – which could be more broadly recognized in the Consultation. Other practices helping to govern third-party models without dramatically increasing the costs of governing the use of the technology include defining Minimum Viable Controls (MVCs) for third-party models, including baseline disclosure expectations, and clear documentation standards[14]. The IIF suggests recognizing these practices in the final version of the commented document.

IIF members welcome the report's recognition that risks posed by AI adoption are not all new; that said, further articulation of the AI-specific dimension of a risk may help clarify where existing risk management approaches may need to be adopted. This is again relevant in the context of third-party risk management. While the report appropriately recognizes that existing regulation and guidance on TPRM provide a suitable foundation, it could go further in explicitly recognize that existing TPRM frameworks and practices remain the appropriate foundation for managing third-party AI risks, with targeted adaptations where AI introduces genuinely distinct characteristics.

- IIF members are also broadly supportive of the AI Shared Responsibility model. This type of model would be useful in delineating accountability among AI vendors, their upstream suppliers, and the FIs that deploy the technology. This model also formally acknowledges that FIs cannot and should not be solely responsible for risks embedded deep within the AI supply chain.

- The report highlights the various AI-enabled types of attacks that can be facilitated by using AI techniques. However, it is worth noting that (1) these additional causes of cyber risk do not change the cyber security framework itself, but rather require additional skills and scrutiny on these types of risks; and (2) existing regulatory frameworks in various jurisdictions already allow to contain these risks.

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

IIF submission text:

- The sound practices are sufficiently flexible at this point. However, wording that may be interpreted as endorsing specific implementation approaches could be softened. For

[13] IIF's AI/ML Experts Group – Briefing Note. May, 2025. Available online, here.
[14] IIF's AI/ML Experts Group – Briefing Note. February, 2026. Available online, here.

example, expressions such as "typically includes" could be replaced with more flexible language such as "could include actions such as".

- IIF members commend the FSB for highlighting the importance of public-private partnerships for better understanding the most advanced AI systems being used, for at least 3 reasons: (1) structured public-private dialogue helps ensure that regulatory objectives are achieved efficiently and without creating unnecessary barriers to innovation, competition, or financial inclusion; (2) effective policymaking benefits from continuous engagement between authorities and industry participants, allowing regulators to access technical expertise while developing proportionate, evidence-based frameworks; and (3) many financial services are inherently cross-border; the same happens with general-purpose technologies like AI. In these cases, cooperation between authorities and industry helps reduce fragmentation and promote greater consistency across jurisdictions, lowering costs while enhancing financial stability and market access.

- The IIF welcomes the FSB's interest in revisiting these practices from time to time to account for changes in technology and the accelerating pace of development of AI models.

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.

IIF submission text: No specific comments.

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

IIF submission text:

- Overall, the case studies provide useful illustrative examples of how sound practices can be implemented in practice. The IIF encourages the FSB to clearly present the case studies as illustrative examples rather than preferred implementation models, consistent with the FSB's stated intention to provide a non-binding menu of sound practices. This clarity will help preserve proportionality and reduce the risk of creating *de facto* supervisory expectations around specific tools, governance structures, or control arrangements if that is indeed not the intention of the FSB at this time.

  Members welcome the direction of the report's recognizing that FIs incorporate AI cyber and ICT risk scenarios into their testing and exercises. At the same time, members have observed a broader pattern in supervisory conversations of scenario testing pushing toward increasingly extreme, low probability events, which in practice pulls investment and attention away from the patching, access management, and fundamental cyber hygiene measures that address the far more common, lower severity incidents firms are most likely to experience. IIF membership encourages the FSB to acknowledge this balance directly in its final report, and to note that such examples illustrate current practice without setting an expected model.

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

IIF submission text:

- As a global standard-setting body, the FSB is positioned to address regulatory fragmentation on key AI concepts. Some IIF members have suggested adjustments to some of the proposed definitions in the Consultation, such as:
    - Materiality – the report uses this term extensively but does not define it in the Glossary (materiality and risk help drive the calibration of risk management requirements across all 12 sound practices). The IIF would appreciate if regulators recognize that IIF members may already consider, among other things, the following criteria to assess the materiality and risk dimensions in specific cases: adaptability, memory, orchestration ability, customer impact, financial and regulatory exposure, data sensitivity, scope of access to systems or data, reversibility, criticality to the institution/service continuity (e.g., use in critical or material functions/services), degree of substitutability, scale of usage, adversarial exposure, and degree of autonomy of action.
    - Agentic AI – would be beneficial to include more specificity in the definition, including the following 3 components: (1) goal directed autonomy (including ability to plan and reason); (2) invokes other tools and systems to carry out a task; and (3) execute an action with real impact.
    - Traditional AI – traditional AI does not include GenAI and beyond. It would also be beneficial to clarify that traditional statistical techniques, such as logistic regression, Generalized Linear Models (GLMs), standard scorecards or rules-based decision trees do not automatically fall within the AI perimeter[15] when they are transparent, deterministic, or manually specified.
    - Shadow AI – Consistent with the Consultation, this could be defined as "when employees adopt AI tools without authorization or for unapproved uses" (page 12).
    - Concentration Risk – the report discusses this concept extensively in section 3.9 but does not provide a standalone definition. Clarifying whether this refers to provider concentration, model concentration, or both, would be useful, particularly since Sound Practice 12 appears to use the term to describe a financial institution's own exposure to a limited number of providers, while Section 3.9 separately touches on sector-wide or systemic concentration as a matter of financial stability risk. The IIF would suggest the glossary distinguish between these two uses so they are not read as interchangeable.

- A definition is a "statement of the meaning of a word, phrase or symbol"[16]. However, some of the definitions proposed by the FSB go beyond that exercise, adding explanations about the reasons behind a certain concept, or adding some of the potential consequences of using specific AI models.

  For example, while defining "foundation models", the FSB states that this is "An umbrella term referring to a diversity of models that are usually trained by applying deep learning to massive quantities of data, such as text and images**. Because the expertise, time, and computing power involved in training foundation models from scratch are typically prohibitive for most non-specialist firms,** these models are usually

[15] IIF's response to EC's consultation on the application of the definition of an AI system and the prohibited AI Practices Established in the AI Act. December, 2024. Available here.
[16] Merriam-Webster Dictionary. Definition of the word "definition". 2026. Available online, here.

pre-trained and shared with end-users for further use and refinement" (emphasis added). In this case, the text in bold provides details that try to explain why foundation models are usually pretrained, but while doing so, it makes it seem as foundation models need to be, by definition, expensive in "terms of expertise, time and computing power", which is not necessarily the case for all foundation models at present or future times.

Also, the FSB defines "open-source models" as "AI models where the full training code, and in some cases the training data or its composition, is made publicly available for use and modification. These models offer customizability **and reduce vendor lock-in but may introduce additional risks, such as security vulnerabilities, data privacy, or data quality concerns**" (emphasis added). In this case, the bold text assumes that by using "open-source models", FIs would be able to reduce vendor lock-in, though this is not necessarily a consequence, nor a reason, for recurring to these models every time an FI does so.

- The FSB uses the concept of an "AI incident" in Sound Practices 4 and 11, in the context of post-incident reviews and information sharing. The IIF would ask the FSB to state plainly that references to AI incidents in the report are not intended to establish a new, freestanding reporting category, and that where an incident involving AI is, in substance, a cyber or ICT incident, it continues to be reported through a financial institution's existing cyber and ICT incident frameworks.

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

- IIF members understand that the Consultation is focused on providing a menu of non-binding sound practices for private sector financial institutions. That said, some IIF members would note that national and regional authorities seem to be pushing for sovereign AI initiatives without taking into consideration the tradeoffs that those frameworks entail for their own policy objectives, including impacts on competitiveness, innovation, security and operational resilience.