July 22, 2026

Via online submission portal

Secretariat of the Financial Stability Board
c/o Bank for International Settlements
CH-4002, Basel
Switzerland

Re: MFA Comments on FSB Consultation Report on Sound Practices for Responsible Adoption of Artificial Intelligence (June 10, 2026)

Dear Sir or Madam:

MFA[1] appreciates the opportunity to submit comments to the Financial Stability Board ("**FSB**") on its consultation report on Sound Practices for Responsible Adoption of Artificial Intelligence, published June 10, 2026 (the "**Consultation Report**" or the "**Proposed Sound Practices"**). We support the FSB's objective of fostering responsible artificial intelligence ("**AI**") adoption across the financial sector and welcome the Consultation Report's recognition that AI presents both significant opportunities and risks that warrant thoughtful governance. We offer the following comments with the goal of helping ensure that the final report is as effective and durable as possible and provides guidance that is calibrated to the actual diversity of institutions and use cases in the financial system.

Central to that goal is preserving the flexible, principles-based character of the Proposed Sound Practices. The FSB has been clear that these practices are not intended to establish an international standard or impose a prescriptive approach. We believe that framing is essential. AI technology is evolving rapidly, the range of institutions adopting it is vast, and the use cases

[1] Managed Funds Association (MFA), based in Washington, DC, New York, Brussels, and London, represents the global alternative asset management industry. MFA's mission is to advance the ability of alternative asset managers to raise capital, invest, and generate returns for their beneficiaries. MFA advocates on behalf of its membership and convenes stakeholders to address global regulatory, operational, and business issues. MFA has more than 180-member fund managers, including traditional hedge funds, credit funds, and crossover funds, that employ a diverse set of investment strategies. Member firms help pension plans, university endowments, charitable foundations, and other institutional investors to diversify their investments, manage risk, and generate attractive returns over time.

vary enormously. We support the FSB's recognition of proportionality given the variety of and sizes of financial institutions. We have concerns, however, that the specificity of the Proposed Sound Practices and case studies is in tension with the report's flexible framing. The Proposed Sound Practices skew heavily toward governance models characteristic of large institutional banks that are neither reflective of current market practice across the sector nor practicable for smaller, less complex, and more nimble organizations. AI is a transformative technology with the potential to democratize sophisticated capabilities across the financial sector. Alternative asset managers already have compelling use cases that demonstrate this potential. However, a one-size-fits-all approach modeled on the infrastructure and organizational capacity of large banking institutions risks slowing this progress, effectively limiting meaningful AI adoption to the largest, most well-resourced institutions. Such an outcome runs counter to the FSB's stated goal of proportionate, responsible AI adoption across a diverse financial system.

Further, in certain areas, the Proposed Sound Practices do not fully reflect the current realities of how AI technology is developed, offered, and deployed in the market today.

Executive Summary

Our comments address five principal concerns:

1. Reinforce the flexible, non-prescriptive intent of the Proposed Sound Practices throughout the body of the final report by revising provisions that imply a single or preferred implementation path and confirming that multiple approaches may satisfy each underlying principle;

2. Revise language throughout the Proposed Sound Practices that implies a governance model or organizational structure that is characteristic of large institutional banks and not reflective of the broader financial sector or proportional for many segments of the industry;

3. Clarify that the case studies included in the final report are illustrative only and do not represent the expected or preferred approach to implementing a given sound practice;

4. Modify the Proposed Sound Practices to reflect current AI market realities, in particular with respect to third-party dependency and concentration risk as well as explainability; and

5. Direct that the Proposed Sound Practices build upon, rather than duplicate, financial institutions' existing model risk management, data governance, and cybersecurity frameworks, and align AI-specific terminology and definitions with existing industry

usage, including through cross-border harmonization, to avoid creating competing standards.

I. The Proposed Sound Practices Should Reinforce Their Flexible, Non-Prescriptive Intent

MFA strongly supports the FSB's stated intent that the Proposed Sound Practices should serve as a flexible, high-level framework that complements rather than supersedes existing regulatory regimes. We urge the FSB to carry that intent consistently into the body of the final report, ensuring that the text as a whole, not just the framing language in the introduction, reflects the flexibility to which the FSB has committed. As currently written, the Proposed Sound Practices risk effectively imposing mandatory requirements through detailed implementation expectations, granular documentation standards, and language that leaves little room for proportionate application. The cumulative effect is a framework that, while nominally flexible, threatens to carry the weight of binding regulation.

Sound Practice 3 (Incorporation of AI risks into risk management framework) illustrates the tension between the introduction's stated flexibility and the prescriptive nature of the Proposed Sound Practices themselves. Its premise, that financial institutions should incorporate AI risk considerations into existing risk management frameworks, is sound and proportionate on its face. In application, however, the practice calls for financial institutions to "track and document their AI use throughout the lifecycle," with an accompanying list of documentation elements spanning ten separate categories even before noting "enhanced documentation and logging for GenAI and agentic AI use cases." It further directs institutions to establish bespoke "contractual arrangements for vendors to provide timely information and compensating controls" where logging or rollback capabilities are vendor managed. Large financial institutions may have the internal resources to build and maintain a centralized AI use-case inventory (as the report's own case study describes) and the negotiating leverage to secure customized contractual terms from AI vendors. Many alternative asset managers, however, cannot readily replicate these expectations, nor would it be necessary or proportionate for use cases that are relatively low risk.

Sound Practice 9 (Performance management) illustrates the same problem from a different angle. Although the practice appropriately calls for performance evaluation "proportionately" to a use case's "materiality and risk," it subsequently directs financial institutions to "specify thresholds across all relevant dimensions of performance" and enumerates five factors for performance assessment. This framework assumes an institution is building or materially customizing its own models and has the testing infrastructure and technical staff needed to implement multi-dimensional performance testing. Many smaller alternative asset managers instead deploy off-the-shelf or lightly configured third-party AI tools, where there is no realistic ability or need to run this type of testing.

Taken together, these examples show that the Proposed Sound Practices are not, in practice, the flexible principles institutions can adapt to their own circumstances that the introduction describes. Rather, in these instances they function as prescriptive operational and structural requirements that leave little practical room for institutions to tailor their approach to their businesses.

Notably, asset managers are not operating in a regulatory vacuum with respect to their adoption of AI. As fiduciaries, asset managers are already subject to overarching duties of care and loyalty that require them to act in their clients' and investors' best interests, including in the selection, deployment, and oversight of the technology they use. These duties are principles-based rather than prescriptive, and that flexibility allows asset managers to calibrate their governance and oversight practices to the specific AI tools they use, the specific risks those tools present, and the specific context in which they are deployed. This is particularly well-suited to the current stage of AI adoption given how rapidly it is evolving. A rigid, checklist-style framework imposed at this moment risks becoming obsolete or counterproductive as the technology matures, whereas the existing fiduciary framework is inherently adaptive.

We recommend that the FSB review the body of the Proposed Sound Practices for language that implies a single or preferred implementation path and revise that language to make clearer that multiple approaches may satisfy the underlying principle. The Proposed Sound Practices should consistently reinforce that the relevant risk or governance objective should be achieved by means appropriate to the institution and use case in question.

II. The Proposed Sound Practices Should Reflect the Diversity of the Financial Sector

Our concerns about the lack of flexibility are compounded by the fact that the Proposed Sound Practices, in their substantive requirements, baseline assumptions, and the case studies selected to illustrate them, are oriented heavily toward large, systemically important banking institutions. Despite calling for proportionality, the requirements are often disproportionate to what could or should be expected of the broader range of financial institutions the FSB intends to address. Alternative asset managers and other non-bank financial intermediaries operate under fundamentally different business models, risk profiles, and governance structures than global systemically important banks. References to “size” and “complexity” are insufficient without concrete guidance on how smaller or more specialized firms should adjust their posture.

The FSB correctly acknowledges the importance of proportionality in framing the Proposed Sound Practices:

> Proportionality should be considered when applying the sound practices in this report. Each financial institution should consider the relevance and applicability

of these sound practices based on its business model, complexity, role in the financial system, structure, size, and the extent to which AI is used or deployed in their critical (or material) functions. Financial institutions may also refer to or adapt certain sound practices, depending on the materiality and risk of a specific AI use case, impact and outcomes of the use case, and the type of AI being adopted.

The Proposed Sound Practices, however, place the onus on financial institutions to revise and adjust prescriptive requirements that appear crafted with only one type of financial institution in mind.

The FSB's treatment of board and senior management oversight in Sound Practice 1 (Strategic direction and oversight) illustrates this concern. Many financial institutions outside the large bank sector do not have a board of directors in the traditional sense, and even where senior leadership bodies exist, it is neither practical nor effective for day-to-day AI oversight to rest with the most senior people or governance bodies within an organization. Instead, we typically observe that senior leadership appropriately sets AI strategy, articulates risk appetite, and weighs in on key decisions, but ongoing oversight responsibility is delegated to individuals, working groups, or committees, as makes sense for the size and organizational structure of the asset manager, with clear accountability and appropriate expertise.

Sound Practice 2 (Governance and accountability) presents a related concern. The practice begins uncontroversially with the notion that financial institutions should "define clear roles and responsibilities and maintain an appropriate governance framework." However, it then directs institutions toward the three lines of defense model as the reference structure, going so far as to state that it is "essential" that audit, evaluation, risk management, and testing functions remain independent and functionally separate from development teams. This may be a reasonable articulation for large financial institutions with the staff and organizational depth to support fully separated first-, second-, and third-line functions. For many alternative asset managers, however, this level of functional separation is neither practical nor achievable, and mandating it as "essential" deprives institutions of the optionality to devise organizational structures and reporting lines proportionate to their size and complexity.

This orientation carries through to the case studies, which are drawn almost exclusively from large, internationally active banks, G-SIBs, and similarly large insurers. None features an alternative asset manager or comparably-scaled nonbank financial institution. Because the Proposed Sound Practices are meant to be illustrated "in practice" through these examples, the effect is to implicitly frame large-bank infrastructure as the reference model for responsible AI adoption, rather than as one proportionate approach among several.

The lack of proportionality in the Proposed Sound Practices is further exacerbated by the fact that they seem not to meaningfully differentiate among levels of risk, instead crafted for very high risk use cases. By contrast, other regulatory frameworks, such as the EU AI Act, distinguish among unacceptable, high, limited, and minimal risk uses of AI and calibrate obligations accordingly. The Consultation Report states that "authorities . . . stress the importance of responsible AI adoption via strong governance, clear accountability, and sound risk management practices especially in high materiality and high risk use cases", but the Consultation Report never defines or otherwise elaborates on what constitutes a "high risk" use case for these purposes. Nor do the Proposed Sound Practices themselves give concrete guidance or illustrations of adjustments for lower risk use cases. Expressly limiting the application of certain Proposed Sound Practices to higher-risk use cases would better align the framework with its own stated proportionality principle.

Sound Practice 10 (Human oversight) presents a related concern. Its core requirement, that institutions implement human oversight calibrated to the materiality, risk, autonomy, complexity, and explainability of each AI use case, is, on its face, appropriately proportionate. However, the Consultation Report's own discussion of agentic AI acknowledges that the impracticality of real-time human monitoring grows as agent use scales and suggests that effective monitoring and detection may "require augmentation with another AI agent or other forms of AI." For a large institution, standing up a second AI system to oversee a first may be a viable option. For most alternative asset managers, however, building or procuring a supplementary layer of AI purely to monitor another AI system is neither proportionate nor, in many cases, feasible, particularly for use cases that are neither highly material nor high risk.

In light of the foregoing, we recommend that the FSB revise the Proposed Sound Practices to more consistently reflect the diversity of institutions operating in the financial sector. First, the FSB should revise language throughout the Consultation Report that assumes a particular organizational structure, such as a traditional board of directors or a fully segregated three lines of defense model, and replace it with objective-based language that identifies the governance outcome sought while leaving institutions latitude to determine the structure, personnel, and reporting lines appropriate to their size and complexity. Second, the FSB should apply the proportionality principle within the body of each sound practice itself, rather than relying on the introduction to do that work; a brief proportionality caveat accompanying each practice would go a long way toward ensuring the framework functions as guidance rather than as a checklist. Third, the FSB should clarify prominently, both in the introduction and immediately preceding each case study, that the case studies are illustrative examples of one approach among several and not a reference implementation or benchmark, and should soften language within the case studies themselves that may otherwise be read as suggesting a single correct way to implement the sound practices. These revisions would allow the final report to be more applicable to the full range of institutions operating in today's financial sector.

III. The Proposed Sound Practices Should Be Revised to Reflect Current AI Market Realities

Several of the Proposed Sound Practices reflect assumptions about how AI is built, delivered, and governed that do not accurately describe the AI market as it exists today. Financial institutions of all types and sizes are increasingly reliant on a small number of large technology companies for foundation models, cloud infrastructure, and application-layer AI products. The commercial, technical, and contractual realities of that market bear little resemblance to the framework the Proposed Sound Practices appear to contemplate in certain areas, in particular with respect to third-party dependency and concentration risk as well as explainability.

A related, fundamental gap is the Proposed Sound Practices' lack of differentiation between "developers" and "deployers" of AI. Other regulatory frameworks distinguish between entities that build or substantially customize an AI model or system and those that procure and use a model or system largely as offered by a third party, and calibrate obligations accordingly. The Proposed Sound Practices do not draw this distinction in any comparable way. This omission is particularly consequential because most financial institutions, including the great majority of alternative asset managers, are deployers rather than developers of the AI models and systems they use at this stage. Proposed Sound Practices premised on the assumption that an institution controls or can meaningfully influence a model's underlying architecture, training data, or ongoing development, are difficult or impossible for a deployer to satisfy in the same way. Sound Practice 8's expectation that institutions "build explainability into model development" and Sound Practice 12's expectations regarding negotiated contractual protections and substitutability are both illustrative of requirements that are premised on a level of access, control, or negotiating leverage that most deploying institutions simply do not have.

Sound Practice 12 (Third-party AI risk management) expects financial institutions to manage AI supply chain risk, negotiate contractual protections, and maintain substitutability arrangements with key AI providers. In practice, the foundation model and cloud infrastructure market is highly concentrated, terms of service are often standardized and non-negotiable for all but the very largest institutions, and genuine substitutability (i.e., the ability to switch from one leading AI provider to another) is not commercially feasible for many institutions. Concentration in this market is a structural feature of the current technology landscape, not a consequence of inadequate procurement practices by individual financial institutions.

Sound Practice 8 (Explainability and transparency) states that where compensating controls cannot adequately mitigate explainability risk, institutions must consider "deploying a more explainable AI or non-AI model or system, even if performance is inferior.” This provision effectively prescribes reverting to a less capable system rather than leaving institutions to determine the appropriate risk management response within their own governance frameworks. It creates potential limitations on AI adoption by requiring institutions to forgo best-available tools where full explainability cannot be achieved.

We urge the FSB to review the Consultation Report carefully and align the assumptions and recommendations more closely with the realities of the current AI landscape.

IV. The Final Report Should Build on Existing Governance Frameworks Rather Than Create Competing Definitions and Standards

The Proposed Sound Practices largely proceed as though AI governance is a novel discipline requiring an entirely new compliance architecture, rather than an extension of governance frameworks financial institutions have refined over decades. As such, they risk creating competing and duplicative standards. The same dynamic operates across borders. Many alternative asset managers and other financial institutions operate across multiple jurisdictions and are therefore subject to a range of overlapping, and at times inconsistent, AI-related regulatory standards. Global firms must comply with various standards across borders, and the risk of competing definitions and frameworks described throughout this section is magnified as individual jurisdictions adopt their own AI-specific requirements.

To promote consistent implementation and reduce regulatory fragmentation across jurisdictions, the FSB could further clarify key terminology and definitions related to AI in the Proposed Sound Practices. For example, terms such as "AI," "traditional AI,” “machine learning," "generative AI," "agentic AI," "foundation models," and "AI models" could be defined more precisely and distinguished from one another, as vague or overlapping terminology may lead to inconsistent application both within and across jurisdictions.

This is apparent, in particular, in areas that closely track long-standing model risk management practices, including those governing quantitative and algorithmic trading models, which already provide mature approaches to model validation, testing, benchmarking, and ongoing performance monitoring. Notably, the FSB itself recognizes elsewhere in the Consultation Report that AI risk assessment can build on "financial authorities' respective regulations or guidance on AI risk management, model risk management, operational resilience, and third-party risk management." Left as is, the current approach risks producing parallel, and at times competing, definitions and frameworks that address substantially the same underlying risks as existing regimes, without reference to or reconciliation with them.

This risk is evident in Sound Practices 8 and 9. Financial institutions, and quantitative investment managers in particular, already validate and monitor proprietary models using wellestablished model risk management concepts such as conceptual soundness, outcomes analysis, benchmarking, challenger models, and sensitivity and stress testing. Sound Practice 9's dimensions of performance (accuracy, range of output, robustness, stability, and sensitivity testing) and Sound Practice 8's explainability methods (including "post-hoc explainability," "triangulation," and the distinction between "global" and "local" explainability) describe substantially the same discipline using new, AI-specific terminology. Where these concepts diverge from or only partially overlap with the model risk management vocabulary institutions already use, financial institutions are left maintaining two parallel taxonomies, and in some instances two parallel validation processes, for what is fundamentally a single underlying activity.

A related concern arises under Sound Practice 7 (Data governance). While the Proposed Sound Practices acknowledge that data governance frameworks "may be pre-existing," the AI-specific elements that follow, including new data classification and labelling schemes, a formal "data lineage" concept, and "policy aware" metadata instructions, are introduced without reference to the data governance, data mapping, and recordkeeping frameworks financial institutions already maintain for other regulatory purposes. As a result, Sound Practice 7 risks functioning as a new, parallel data governance taxonomy layered on top of institutions' existing frameworks, rather than as guidance on how to extend those existing frameworks to cover AI-specific data uses.

The same concern applies to Sound Practice 11 (Cyber and ICT risk management). MFA and its members recognize that AI introduces genuine, AI-specific cybersecurity risks, and we welcome the FSB's attention to them. However, cybersecurity and operational resilience are already heavily regulated, with mature frameworks addressing technology risk, third-party and supply chain risk, incident detection and reporting, and business continuity. Without careful coordination, Sound Practice 11 risks becoming a competing standard that runs alongside, rather than within, those existing frameworks, creating duplicative compliance burdens and, potentially, conflicting definitions of common cyber and ICT risk concepts.

We urge the FSB to revise the final report to make clearer that the Proposed Sound Practices, in many respects, can be extensions of existing model risk management, data governance, and cybersecurity and operational resilience frameworks, rather than as freestanding, AI-specific structures. Where existing frameworks already address a risk present in an AI context, compliance with those existing requirements should satisfy the relevant Sound Practice without the need for duplicative AI-specific measures or divergent AI-specific definitions.

V. Conclusion

MFA appreciates the FSB's work on this important initiative and the opportunity to provide these comments. We strongly support the goal of responsible AI adoption across the financial sector and believe that well-calibrated, genuinely flexible guidance will do more to advance that goal than prescriptive requirements that do not reflect the diversity of the institutions, markets, and technologies they are intended to address. We urge the FSB to consider the recommendations set forth above as it finalizes the Proposed Sound Practices, and we welcome the opportunity to discuss any of these points further. If the FSB has questions or comments, please do not hesitate to contact Jill R. Whitelaw at JWhitelaw@mfaalts.org or the undersigned at JFlores@mfaalts.org with any questions regarding this letter.

Respectfully submitted,

/s/ Jillien Flores

Jillien Flores

Chief Advocacy Officer