American Bankers Association

Ryan T. Miller
Vice President & Senior Counsel
Innovation Policy
+1 (202) 663-7675
rmiller@aba.com

July 22, 2026

Secretariat to the Financial Stability Board
Bank for International Settlements
Centralbahnplatz 2
CH-4002 Basel
Switzerland

Via Online Submission

Re: ABA Comment Letter to the Financial Stability Board responding to its Sound Practices for Responsible Adoption of Artificial Intelligence Consultation Report

To Whom It May Concern:

The American Bankers Association[1] (ABA) appreciates the opportunity to respond to the Financial Stability Board's (FSB) Sound Practices for Responsible Adoption of Artificial Intelligence Consultation Report (Consultation Report).[2] Overall, our members found the report directionally sound and potentially of significant benefit to their institutions. However, we do have some targeted suggestions that we believe would make the deliverable even more effective, and we hope the FSB considers them in this spirit.

The following comments are based on discussions with ABA's AI Working Group, which is an interdisciplinary body of data scientists, technologists, compliance, legal, risk, security, audit, and human resources experts representing banks of all sizes. This group routinely comes together to formulate policy positions as well as to discuss operational best practices. Our comments are also informed by the work of the Financial Services Sector Coordinating Council's R&D Committee.[3]

Our members are particularly appreciative that the Consultation Report takes a proportionate, outcomes focused approach. ABA encourages the U.S. banking agencies to consider the FSB's sound practices as a helpful input into their ongoing AI-related guidance, supervisory, and policymaking efforts. The Consultation Report also aligns with outputs from standard-setting

[1] *The American Bankers Association is the voice of the nation's $26.1 trillion banking industry, which is composed of small, regional and large banks that together employ over 2 million people, safeguard $20.5 trillion in deposits and extend $13.7 trillion in loans.*
[2] https://www.fsb.org/uploads/P100626.pdf
[3] In 2026, FSSCC published a series of AI papers covering an AI risk management framework, identity & authentication controls, explainability, nutrition labels, and a lexicon; see https://fsscc.org/aieog-ai-deliverables/.

organizations and industry-led transparency and assurance mechanisms, including model cards, system cards, and similar documentation artifacts that can support banks' due diligence, oversight, and risk management activities.

To make the final report more beneficial to banks of all sizes, ABA recommends that the FSB make the changes addressed below.

Definition of "Machine Learning"

As the FSB is well aware, definitions of key terms are a foundational matter for the adoption of best practices. Among the definitions, "machine learning" is among the most integral. However, we believe the current version could be improved.

> **Current Definition of "Machine Learning"**
>
> A method of designing a sequence of actions, known as algorithms, to solve a problem, which optimise automatically through experience and with limited or no human intervention.[4]
>
> **Recommended Definition of "Machine Learning"**
>
> An AI learning method that enables computational systems to learn patterns, make predictions, and optimize decisions from large amounts of data without being explicitly programmed for each task. Machine learning encompasses supervised, unsupervised, and reinforcement learning paradigms, serving as the technical foundation for data-driven intelligence and automation.[5]

Indeed, the FSB might find it fruitful to align several of the definitions laid out in the AI lexicon linked in footnote 5. Lack of uniform definitions, both intra- and interorganizational, is an impediment to effective governance. The lexicon could serve the function of universalizing many of the common terms in this space. However, at minimum the FSB should revise the definition of "machine learning" as laid out above.

[4] Supra, note 2 at page 57 of the Consultation Report/page 61 of the PDF.
[5] Artificial Intelligence Executive Oversight Group AI Lexicon, page 6; available at https://fsscc.org/wp-content/uploads/2026/02/AIEOG-AI-Lexicon-February-2026.pdf.

Inclusion of additional Case Studies

ABA encourages the FSB to incorporate some or all of the below case studies offered by community banks (many of which are under $1 billion in assets) where appropriate in their final report. We would note that these use cases would be beneficial to banks of all sizes, with appropriate risk tailoring.

- Development of AI agents for various loan review functions, including identifying potential applicability for the Community Reinvestment Act, evaluating HMDA data, verifying income, confirming signatory requirements, determining whether the application is complete, etc. The agents' work would triage workflows for follow-up by human employees.

- Development of an AI agent that reviews commercial credit write-ups to identify conditions to close, approval requirements, and other material items, which would create a follow-up checklist for validation by lending employees. In addition, the AI agent would draft commercial loan covenants for clarity, tone, accuracy, and completeness, which would provide relationship managers with a starting point for review, modification, and formal approval.

- Creation of prompts to test for compliance with Regulation B and Regulation E. Validation focuses on whether the AI agent's actions are based on the appropriate factors and that they create a sufficient audit trail.

- Development of an AI reference assistant that organizes regulatory materials, supports employee training, and identifies potentially conflicting requirements, which would speed research while requiring compliance personnel to verify conclusions against authoritative sources. Similarly, AI can assist with the development/review of policies and procedures, as well as the action items for regulatory change management.

- Development of an AI agent that analyzes transaction activity against expected customer behavior to identify unusual patterns, which would prioritize potential fraud alerts for investigation and disposition by trained bank employees. A similar agent would review identity and access data in internal systems for unusual patterns and assists with user provisioning workflows, which would prioritize exceptions for investigation and approval by security personnel.

- Use of AI to prepare an initial draft of commercial trust resolutions from approved transaction details, which would reduce drafting time while requiring legal or lending personnel to verify authority, language, and completeness.

- Use of AI to prepare initial versions of marketing material, with a layer to check against compliance requirements. The product would still be routed to the appropriate second line review function; however, the AI would assist the first line business not only with the creative but also a preliminary assessment of passing regulatory muster.

Consolidation and Reorganization

ABA is aware the FSB intended to break the Sound Principles into thematic blocks covering 1) organization-wide AI governance, 2) managing specific AI use cases, and 3) the importance of managing cyber and third-party risks.[6] However, in practice many of the Sound Practices dovetail into one another and the inorganic division negatively impacts clarity.

As the FSB considers how to improve the deliverable to make it more accessible for banks of all sizes, it might consider merging some of the related concepts, resulting in fewer Sound Practices without sacrificing any substance. Because the document is not prescriptive and each bank will have to embody the principles under its own construct anyway, such a reconstitution will simplify the report's crystallization and application.

In addition, certain sections should be moved up in order of appearance to underscore their fundamental importance (e.g., Cyber and ICT Risk Management and Third-Party Risk Management).

After consolidation and reorganization, the list of Sound Practices would be as follows:

[New Name] (existing Sound Practice)

1. [Tailored Oversight] (comprised of Strategic Direction and Oversight/Materiality and Risk Assessment)
2. [Governance Program] (comprised of Governance and Accountability/Incorporation of AI Risks into Risk Management Framework/Organisational Adaptability)
3. Cyber and ICT Risk Management
4. [Due Diligence] (comprised of Selection/Third-Party Risk Management)
5. Data Governance

[6] Supra, note 2 at page 3 of the Consultation Report/page 7 of the PDF.

6. [Validation and Controls] (comprised of Performance Management/Human Oversight/Explainability and Transparency)

In addition, the final report should make it clear that it is an aspirational document that can help inform banks of all sizes what they *should* be doing versus what they *are* doing. Moreover, the text should emphasize that the programs implementing the principles should be commensurate with the size, complexity, and risk exposure of the bank. The report should also recognize that the bank's individual culture and organizational makeup are dispositive factors and AI governance must be built organically with these in mind, not in spite of them. There is no "one-size-fits-all" and the document should make this fact unambiguously clear.

In addition to the structural changes, ABA will address each of the sections below where targeted revisions can be made. If no substantive edits are indicated, the FSB should interpret that as the current text being satisfactory as is.

1. Tailored Oversight

Strategic Direction and Oversight (currently Sound Practice 1)

The FSB should consider softening this section to make it clear that the board and senior management involvement should follow a principle of proportionality and tailoring. When U.S. regulators consider adapting this practice for U.S. banks, ABA recommends that U.S. financial regulators recognize the distinction between the oversight role of a financial institution's board of directors and the operational and execution roles of senior management. Signoff should occur at the level appropriate given the overall risk exposure, the organization's appetite, and other relevant factors. For that reason, combining this section with "Materiality and Risk Assessment" (currently Sound Practice 5) would marry well.

In addition, tailoring is essential given the obligations imposed across various jurisdictions. Many banks operate at scale in interstate and international commerce, and the ability to adjust the compliance regime to meet differing requirements is a prerequisite for sound business practices.

Materiality and Risk Assessment (currently Sound Practice 5)

ABA supports the FSB's observation that materiality should be reassessed as appropriate after deployment.[7] As new AI tools are added, it results in more dependencies. Reassessment should

[7] Id. at page 27 of the Consultation Report/page 31 of the PDF.

be based on the use of the AI tool, or precipitating event (it need not be "regular"). For example, the phenomenon is exemplified with stress testing and BSA/AML monitoring.

Generally speaking, risks of AI are largely dependent upon the specific use case. Certain applications, such as bank employees using Retrieval-Augmented Generation (RAG) to interact with a repository of internal policies would be lower risk than an AI agent interacting directly with customers. As another example, a chatbot powered by static, keyword-based rules would present less inherent risk than one utilizing dynamic generative AI.

2. Governance Program

Governance and Accountability (currently Sound Practice 2)/Incorporation of AI Risks into Risk Management Framework (currently Sound Practice 3)

These sound practices are related and the FSB might consider combining them. Community and regional banks would benefit from additional insight into how to formalize existing cross-functional collaboration into more established governance practices. Meanwhile, working horizontally across business lines continues to be important for banks of all sizes and will take on additional importance as agentic AI scales.

Further, banks should have flexibility to leverage existing governance and risk management frameworks, and to proactively evolve the application of these frameworks in a targeted and focused way to address genuinely new or differentiated risks of AI systems.

Organisational Adaptability (currently Sound Practice 4)

ABA recommends that the FSB expand on the "maintain external awareness" portion under Organisational Adaptability (currently Sound Practice 4) treat "new supervisory expectations" as a separate component.[8] Further, it should be broadened to include new laws and regulations, not just supervisory expectations.

In particular, a best practice for banks is to anticipate policymaker developments. Technological systems and workflows should be flexible, not rigid. Modern bank operations, including but not limited to those involving AI, would ideally be nimble and resilient enough to adapt to evolving technologies, emerging threats, new consumer behaviors, and shifting legal requirements.

The FSB might also consider encouraging regulators to conduct symposia and/or have "office hours" to disseminate their views on how banks will be assessed for AI governance.

[8] Id. at page 24 of the Consultation Report/page 28 of the PDF.

3. Cyber and ICT Risk Management

Cyber and ICT Risk Management (currently Sound Practice 11)

ABA members note this is a critical area and only become more so after the advent of Anthropic's Mythos (and other powerful AI models). Accordingly, it should appear earlier on the list of Sound Practices. While being prescriptive is not desirable, the proliferation of a new caliber of threats may require a new assessment of patching cadences (industry standards are currently at 30-60 days).

Enhancing the speed of patching is a recognized priority for the sector given developments in frontier AI, but it remains necessary for firms to take a risk-based approach to vulnerability classification and patch deployment. Patching alone will not address the risks facing the sector. Nonetheless, some regulators have begun suggesting a "no-risk approach" to patching. This is not viable at scale and could have negative implications if firms are diverted from prioritizing those vulnerabilities and patches which entail the greatest potential harm.

In addition, ABA recognizes the importance of incorporating AI into tests and exercises, in particular, in determining what constitutes a severe but plausible scenario against which firms should be planning their resilience strategies. However, we caution that an overemphasis on extreme tail-risk scenarios can have negative implications in how firms prioritize and allocate resources. Most industry guidance and regulatory statements on frontier AI cyber risks have acknowledged the importance of fundamental cyber hygiene. This must be the priority for the industry. While it is worth further exploring solutions and strategies for recovering from an extreme cyber breach, this is a nascent area and many of the solutions which regulators have pointed to are not yet market ready or scalable for large banks. There is a real opportunity cost risk associated with prioritizing highly complex, untested solutions over the many tried-and-tested preventative controls upon which firms can still improve.

We welcome the report's inclusion of a "defense in depth" approach involving a set of layered controls that provide additional resilience to the firm's security posture. However, we note the reference to micro-segmentation.[9] Although segmentation is a valuable control, we encourage regulators to focus on the outcome they seek to achieve, namely preventing lateral movement of threat actors who have achieved an initial foothold in a victim's network. While segmentation can support this, in a large and complex network micro-segmentation can be highly complex to implement creating risks of both gaps in protection where legitimate traffic must continue to flow and operational incidents that can lead to service unavailability. Authorities should be

[9] Id. at page 45 of the Consultation Report/page 49 of the PDF.

aware of this and allow banks to evidence how they control the risk of lateral movement to within risk appetite.

Finally, we note that the Cyber and ICT Risk Management Sound Practice implies in multiple places the need for advanced AI models to undertake certain work, such as vulnerability identification and remediation. While using the most advanced models is advantageous, it should not be an excuse not to use existing models to do the same work. Many of the principles, architecture and security controls necessary for this work are applicable across all models and firms that delay in exploring and developing expertise in this area while they wait for access to the most advanced models are putting themselves at a significant disadvantage. We also note that the report currently covers only briefly some of the security considerations for using AI for this work. The importance of the appropriate security controls being in place before firms can make meaningful progress on vulnerability identification and testing should not be underestimated and may be worth more consideration in the final report.

For the FSB's awareness, ABA prepared an assessment guide to aid banks working with their service providers to detect and remediate cybersecurity vulnerabilities.[10]

4. Due Diligence

Selection (currently Sound Practice 6)/Third-Party Risk Management (currently Sound Practice 12)

In addition to moving up the Third-Party Risk Management entry higher on the list, the FSB should consider combining it with the related "Selection" (which covers due diligence) while expanding it beyond third parties to include additional actors. Service providers for third parties (and other forms of subvendors) have a significant impact on the overall risk exposure for a bank, and the FSB's framework should acknowledge this reality. (ABA recognizes that the "Selection" section is not purely geared towards third-party risk, although it does lean that way given the AI tech stack. For avoidance of doubt, we have no substantive commentary on "Selection" other than advising its consolidation with the Third-Party Risk Management section.)

ABA believes it would be beneficial to encourage the use of "notice and opportunity to optout" of third-party AI functionality to give the banks more control over their own systems. Banks also would seek additional leverage over transfer of data to jurisdictions outside the United States, which introduces significant compliance and information security concerns.

[10] https://www.aba.com/news-research/analysis-guides/ai-assisted-vulnerability-detection-and-remediation. ABA also maintains a landing page for an abundance of AI resources, found at www.aba.com/AI.

Moreover, the FSB should do more to distinguish between third-party provision of bank AI versus the third party's integration of AI for its own purposes. Both of these areas present important albeit distinct risk domains. The latter becomes particularly challenging with respect to ensuring adequate monitoring and audit rights, as many third parties push back against these demands on the grounds they are unduly burdensome. In any event, the method of verification is also questionable, as surveys are insufficient tools. Expectations, including on documentation and tracking, should therefore not be equally applied for this indirect AI exposure – we encourage the FSB to more directly note these stronger limitations and pragmatic expectations in this area.

We would welcome continued regulatory focus on improving transparency and information flows across the AI value chain, including practical expectations regarding the information deployers should be able to obtain from AI providers and other third-parties to support governance, assurance and regulatory compliance. Regulators should use their regulatory and supervisory authorities – including under the Bank Service Company Act in the US – to help promote additional controls and transparency for upstream providers.

Furthermore, the sound practices embodied in these sections are compatible with the concept of a standard-setting organization. ABA has consistently advocated for the development of voluntary standards and certification programs within the AI supply chain to allow for appropriate due diligence. Developers and deployers should not have to choose among explainability, risk management, and commercial best practices. Over time, these voluntary programs will mature and provide value to banks of all sizes (as well as the broader ecosystem).

Finally (and consistent with the above), the FSB should strongly encourage the development of standardized, high-quality model/system cards by AI developers to streamline the due diligence and monitoring frameworks employed by banks.

5. Data Governance

Data Governance (currently Sound Practice 7)

Data management and governance is a crucial factor in AI, since models are only as good as the training data. However, we encourage the FSB to avoid suggesting that there are "AI-specific considerations" in data governance, which could imply that AI requires a separate or bespoke data governance framework or that data in AI should be governed differently. In addition, we recommend that FSB acknowledge that data management and governance practices should be proportionate to materiality, risk, and intended use of the AI application.

Furthermore, some practices – such as determining data provenance -- are in nascent stages.

Thus, a component of an effective data governance program is to recognize when such controls are necessary, which means developing expertise for evaluating data quality, accuracy, integrity, and controls as practices mature.

In addition, we would recommend that the FSB clarify certain terms -- such as "unconventional data," and "authenticity" of data – that are not commonly used in the field, clarify the reference to metadata, and respectfully note that data recovery practices are not generally considered data security practices.

Moreover, many community banks are reliant on core service providers and are not able to access much of their data easily. This presents unique data governance challenges for them.

Finally, the FSB should recommend that banks embed a semantic layer within their data structure where feasible. This layer captures common business vocabulary, allowing the AI to interact directly with business concepts rather than raw database tables.

6. Validation and Controls

Performance Management (currently Sound Practice 9)

ABA members found it helpful that this section is consistent with the revised model risk management guidance (SR 26-2, et al).[11]

Human Oversight (currently Sound Practice 10)

The section on Human Oversight (currently Sound Practice 10)[12] is in general very well done, recognizing a wide-ranging spectrum of human controls. ABA would like to place some additional context on the "AI-in-the-loop" category.[13]

Firms should assign clear human accountability for agentic systems and workflows, with the accountability model calibrated to the risk assessment. Accountability refers to identifiable human ownership of the system, its approved use, its control environment, and material decisions or actions taken through the workflow.

Firms should also prioritize effective oversight of AI systems, calibrated to the risk assessment. Oversight is distinct from accountability: it concerns the mechanisms used to monitor, review, challenge, escalate, or intervene in AI-driven activity. Human oversight will be an important

[11]

[12] Id. at page 41 of the Consultation Report/page 45 of the PDF.

[13] Id. at page 42 of the Consultation Report/page 46 of the PDF.

component, particularly for higher-risk use cases, but effective oversight need not always require real-time human intervention whenever an AI system produces undesired content.

Further, when choosing among the human oversight controls, validators have to bear in mind what they are comparing to. For example, while the AI systems may not be perfect, they might perform more consistently than an aggregate of human employees.

The oversight discussion changes significantly when agentic AI is involved. Where appropriate for the risk profile, tailored governance of agents would include maintaining attribution and traceability of agent actions; constraining agent activity within approved environments; defining agent identity, delegated authority, and least-privilege access; applying zero-trust and runtime enforcement controls; establishing escalation boundaries for higher-risk actions; conducting evaluation, validation, and continuous assurance; implementing circuit breakers, resiliency controls, and human override mechanisms; and assigning clear human accountability.

The FSB should also consider the impact of how shutting down AI systems will impact business continuity. Regulatory and supervisory expectations should therefore recognize that, for some AI-enabled services, immediate vendor substitution or manual continuation may be impracticable. The FSB acknowledges this reality in its December 2023 TPRM toolkit, noting "the feasibility of an exit plan may depend on the circumstances of the financial institution, the relevant critical service and the service provider"[14] and that "some critical services are challenging to substitute and developing options to exit over a short-term may not always be feasible without undue costs or risk to the financial institutions."[15]

In such cases, resilience expectations should not overly-prioritize exit or replacement strategies; continuity planning, service degradation protocols, fallback controls, contractual protections, incident response, and other measures that enable financial institutions to maintain critical operations safely during periods of disruption should also be provided equal importance. The FSB also acknowledges this approach in its TPRM toolkit, noting "financial institutions may prioritise the continued performance of critical operations and identify all viable forms of exiting critical third-party services relationships..."[16]

Explainability and Transparency (currently Sound Practice 8)

ABA members believed this section was particularly well crafted and thorough. However, we would note that the contemplated use of banks offering customers an "optout" of AI-enabled

[14] https://www.fsb.org/2023/12/final-report-on-enhancing-third-party-risk-management-and-oversight-a-toolkit-for-financial-institutions-and-financial-authorities/ at page 28 of the report/page 32 of the PDF.
[15] Id.
[16] Id.

functions[17] would be very difficult to implement due to lack of control over underlying data and other dependencies (please refer to the limitations of using third-party AI tools addressed above in the "Due Diligence" section). Any such mechanisms should therefore be proportionate to the use case, the materiality of the AI to the customer outcome, and the potential impact on the customer.

In addition, the FSB should incorporate the following elements into this section to reflect principles that support transparency and monitoring:

- small language models to build field specific AI with transparent structure;
- retrieval augmented generation to index data transparently;
- built-in code monitoring underlying data structures/correlations and drift; and
- built-in code monitoring deviations/drift in results.

Conclusion

As can be observed from the above, ABA members received the Consultation Report in a very positive light. While there are several areas in which it can be improved, the Consultation Report as it stands is very much a solid foundation.

ABA and our members are grateful for the opportunity to provide input into Consultation Report and stand ready to work with policymakers throughout the publication process.

If you have any questions about this comment, please contact Ryan T. Miller via email at rmiller@aba.com or by phone at +1 (202) 663-7675.

Respectfully submitted,

Ryan T. Miller
Vice President & Senior Counsel, Innovation Policy

[17] Supra, note 2 at page 36 of the Consultation Report/page 40 of the PDF.