July 22, 2026

Via Secure Online Form

Financial Stability Board
Basel, Switzerland
fsb@fsb.org

Re: Sound Practices for Responsible Adoption of Artificial Intelligence: Consultation Report

Ladies and Gentlemen:

The Bank Policy Institute[1] and the Institute of International Bankers[2] welcome the opportunity to provide feedback on the Financial Stability Board's consultation report *Sound Practices for Responsible Adoption of Artificial Intelligence.*[3] The Sound Practices will be instructive to banking regulators globally as the regulatory and supervisory frameworks for financial institutions' use of AI continues to evolve. The FSB has a unique opportunity to foster greater harmonization and interoperability in global regulatory frameworks by promoting consistent, coordinated implementation of the Sound Practices.

Given the rapid evolution of AI technologies, it is critical that regulators promote competition and economic growth by providing banks with the ability to remain agile and move quickly as technology changes. Regulators should avoid imposing restrictive, bank-specific

[1] BPI is a nonpartisan public policy, research and advocacy group that represents universal banks, regional banks, and major foreign banks doing business in the United States. BPI produces academic research and analysis on regulatory and monetary policy topics, analyzes and comments on proposed regulations, and represents the financial services industry with respect to cybersecurity, fraud, and other information security issues.

[2] The Institute of International Bankers represents the U.S. operations of internationally headquartered financial institutions from more than 35 countries around the world. The membership consists principally of international banks that operate branches, agencies, bank subsidiaries, and broker-dealer subsidiaries in the United States. The IIB works to ensure a level playing field for these institutions, which are an important source of credit for U.S. borrowers. These institutions also enhance the depth and liquidity of U.S. financial markets and contribute significantly to the U.S. economy through direct employment of U.S. citizens, as well as through other operating and capital expenditures.

[3] FINANCIAL STABILITY BOARD, SOUND PRACTICES FOR RESPONSIBLE ADOPTION OF ARTIFICIAL INTELLIGENCE (June 10, 2026), https://www.fsb.org/uploads/P100626.pdf (hereinafter the "Report").

requirements that place financial institutions at a competitive disadvantage relative to nonbank firms. Subjecting banks to heightened and prescriptive supervisory expectations while nonbank competitors face no equivalent framework risks discouraging banks from responsibly adopting AI and shifting activity to less regulated parts of the market. Moreover, any systemic risks posed by AI adoption are just as likely to arise in non-banking areas of the economy; if regulatory intervention is believed to be an effective way to manage this risk, it should be extended to critical infrastructure generally rather than confined to banks. A level playing field requires regulators to apply comparable standards to comparable risks across both bank and nonbank providers.

Banks are committed to using AI responsibly in their businesses and operations and have long used traditional AI and machine learning tools to support critical risk management functions. The emergence of and rapid advancements in generative AI, agentic AI and frontier models have created new opportunities for financial institutions to improve AML, fraud prevention, cybersecurity, customer service and internal productivity, among other capabilities. As Federal Reserve Board Vice Chair for Supervision Michelle W. Bowman recently noted, "AI has become an integrated part of [the FRB's] daily experience. Financial institutions are developing their own applications and implementing vendor-assisted tools. Banks of all sizes benefit from its greater efficiency, speed, and content generation. Whether used in targeted modeling or enterprise-wide tools, AI will become a force multiplier for the financial system, and in the broader U.S. economy."[4]

The Sound Practices address issues of critical importance to financial institutions, their customers, financial regulators and the broader financial system. The Report appropriately recognizes that financial institutions are rapidly adopting AI to reap efficiency gains, improve customer service and enhance risk management capabilities, including cyber defense.[5] We view the Report as a positive step toward adopting a strong, fit-for-purpose and durable regulatory approach to AI use and responsible innovation in the financial sector.

In particular, we support the Report's emphasis on proportionality[6] and "materiality and risk"-based safeguards as guiding pillars for safe and sound AI adoption. Including these principles in any regulatory framework governing financial institutions' AI use is critical to ensuring that firms retain the flexibility to leverage AI-enabled tools, systems and capabilities to meet evolving customer demands and combat rapidly evolving threats. For example, recent advancements in frontier models, like Anthropic's Mythos, accelerate the pace of cyber

[4] Vice Chair Michelle W. Bowman, Remarks at the Federal Reserve Board, FEDERAL RESERVE (May 1, 2026), https://www.federalreserve.gov/newsevents/speech/bowman20260501a.htm ("Bowman Remarks").

[5] Report, *supra* note 2 at 1.

[6] Proportionality, as the term is used in the Report and this letter, refers to consideration of an institution's "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" in the calibration of regulatory and supervisory expectations relating to a financial institution's use of AI. Report, supra note 2 at 4.

operations and require increased focus on fundamental cyber hygiene, including using AI tools for effective cyber defense.[7] These developments underscore the need for regulators to adopt an approach that does not undermine financial institutions' ability to adapt to and combat risks with agility and scale as AI continues to evolve at a rapid pace. In this regard, we agree with the FSB that "[f]inancial institutions may also need to consider the implications of *not* adopting AI."[8]

At the same time, how the Sound Practices are *implemented* in bank regulatory regimes could determine whether a particular regulatory framework stifles responsible innovation or subjects institutions to overlapping or inconsistent requirements across jurisdictions. For example, operationalizing the Sound Practices as inflexible, AI-specific checklists and documentation requirements would slow adoption and divert risk management resources away from critical areas. When regulators and supervisors fail to appreciate the benefits and risk profiles of AI use cases or interpret rules or guidance in an overly-prescriptive, "one-size-fits-all" manner – banks can be discouraged from using AI to innovate and improve business processes and service delivery.[9] A regulatory framework that instead implements the Sound Practices by emphasizing proportional and risk-tailored guardrails, prioritizing public-private collaboration and information sharing, and avoiding duplicative or unnecessary compliance requirements would encourage safe and sound AI adoption and innovation by financial institutions.

The United States faces a critical inflection point in AI development and its role in the financial regulatory landscape. Project Glasswing[10] and Anthropic's Mythos model have changed the tone of AI policy discussions in the U.S.[11] In response to these and other developments, U.S. policymakers increasingly rely on public-private coordination to shape their approach to AI in financial services, particularly critical cybersecurity and infrastructure

---

[7] For example, FRB Governor Lisa Cook recently stated that "[v]ery powerful AI tools, such as Anthropic's Mythos Preview model, have demonstrated the ability to detect previously undetectable vulnerabilities in software applications that support important and widely used computer systems." Governor Lisa D. Cook, *AI, the Economy, and the Financial System*, Remarks at the Federal Reserve Board, FEDERAL RESERVE (May 27, 2026), https://www.federalreserve.gov/newsevents/speech/cook20260527a.htm.

[8] Report, *supra* note 2 at 3.

[9] *Response to Request for Information; Regulatory Reform on Artificial Intelligence*, BANK POLICY INSTITUTE (Oct. 27, 2025), https://bpi.com/wp-content/uploads/2025/10/BPI-OSTP-AI-RFI-Response-10.27.25.pdf ("Specifically, some BPI members have experienced difficulty engaging with supervisors and examiners who may not appreciate the benefits and risk profiles of AI use cases. This can lead to hesitation, if not risk-aversion, across organizations when it comes to AI adoption.")

[10] Project Glasswing is Anthropic's limited release, controlled defensive initiative around Mythos. *See* https://www.anthropic.com/project/glasswing.

[11] *See*, *supra* note 3 Bowman Remarks ("Today, we are facing the rapid evolution of AI tools much earlier than many expected, and the risks and benefits are now more tangible and clear. Anthropic's Mythos—an AI model that identifies cyber vulnerabilities—highlights the dynamic nature of this technology and the rapid pace that its capability can evolve.")

functions.[12] U.S. federal banking regulators also are reconsidering their approach to model risk management,[13] third-party risk management[14] and other AI-related risk vectors. For example, earlier this year, the FRB, the Federal Deposit Insurance Corporation, and the Office of the Comptroller of the Currency issued revised MRM guidance, superseding 2011 guidance predating the advent of GenAI and agentic AI models.[15] The revised MRM guidance, like the 2011 guidance, sets forth supervisory expectations regarding banks' development, use and testing of "models." However, unlike the 2011 guidance, the revised MRM guidance expressly excludes GenAI and agentic AI models from its scope, recognizing that such models are "novel and rapidly evolving."[16] The OCC has indicated that the U.S. federal banking "agencies plan to issue in the near future a request for information that addresses model risk management generally and considers, in particular, banks' use of AI, including generative AI and agentic AI and AI-based models."[17]

---

[12] For example, a June 2, 2026 Executive Order titled "Promoting Advanced Artificial Intelligence Innovation and Security" directed the U.S. Secretary of the Treasury, in consultation with the National Cyber Director, the Director of the National Security Agency, and the Director of the Cybersecurity and Infrastructure Security Agency, to, within 30 days, form an "AI cybersecurity clearinghouse"— a voluntary collaboration with the AI industry and operators of critical infrastructure to coordinate and deconflict software vulnerability scanning, validate such vulnerabilities, and prioritize remediation and distribution of vulnerability patches. *See* Executive Order No. 14,409, Promoting Advanced Artificial Intelligence Innovation and Security, 91 Fed. Reg. 34,565 (June 5, 2026).

[13] *See Supervisory Guidance on Model Risk Management*, SR 26-2, FEDERAL RESERVE (April 17, 2026), https://www.federalreserve.gov/supervisionreg/srletters/SR2602a1.pdf; *Model Risk Management: Revised Guidance*, OCC 2026-13, OFFICE OF THE COMPTROLLER OF THE CURRENCY (April 17, 2026), https://www.occ.gov/news-issuances/bulletins/2026/bulletin-2026-13.html; *Agencies Revise the Interagency Model Risk Management Guidance*, FIL-15-2026, FEDERAL DEPOSIT INSURANCE CORPORATION (April 17, 2026), https://www.fdic.gov/news/financial-institution-letters/2026/agencies-revise-interagency-model-risk-management-guidance (hereinafter, "Revised MRM Guidance").

[14] Bowman Remarks, *supra* note 3 ("We are also working to update and simplify our third-party risk-management guidance to reflect actual and future risk.")

[15] *See Guidance on Model Risk Management, SR 11-7,* FEDERAL RESERVE (Apr. 4, 2011) https://www.federalreserve.gov/boarddocs/srletters/2011/sr1107.pdf; *Supervisory Guidance on Model Risk Management*, OCC 2011-12, OFFICE OF THE COMPTROLLER OF THE CURRENCY (Apr. 4, 2011); *Adoption of Supervisory Guidance on Model Risk Management*, FDIC FIL-22-2017, FEDERAL DEPOSIT INSURANCE CORPORATION (June 7, 2017), https://www.fdic.gov/news/financial-institution-letters/2017/fil17022.html*; see also* Response to Request for Information; Regulatory Reform on Artificial Intelligence, BANK POLICY INSTITUTE (Oct. 27, 2025), https://bpi.com/wp-content/uploads/2025/10/BPI-OSTP-AI-RFI-Response-10.27.25.pdf (expressing concerns that the 2011 MRM guidance was "not the appropriate framework under which to evaluate non-deterministic AI models" because it was "based on traditional validation techniques … [that] are ill-suited to GenAI or more advanced AI models, like agentic AI models.")

[16] Revised MRM Guidance, *supra* note 11 at note 3.

[17] *See* OCC, Press Release, *Model Risk Management: Revised Guidance* (April 17, 2026), https://www.occ.gov/news-issuances/bulletins/2026/bulletin-2026-13.html.

Against this backdrop, we recommend that U.S. financial regulators work closely with the banking industry, AI developers and global standard setters, including the FSB, to understand the technology as it continues to develop, and to put in place a pragmatic, flexible and streamlined approach to regulating financial institutions' use of AI. Accordingly, this submission advocates three key messages that we hope prove helpful to the FSB, global financial regulators and U.S. policymakers when considering how to incorporate the final Sound Practices into future regulatory frameworks:

(i) AI regulation should acknowledge both the risks and the benefits of AI, reflect a "materiality and risk"-based approach that does not treat AI as a risk in and of itself and recognize the risks of *not* adopting AI.

(ii) Regulators should apply proportionality when determining regulatory and supervisory expectations for financial institutions' use of AI in a manner that ensures banks have the flexibility to leverage existing governance and risk management frameworks and proactively evolve the application of these frameworks to address new or differentiated risks of AI systems.

(iii) Supervisory expectations established in the regulatory framework should be calibrated to recognize the limitations on financial institutions' ability to control, diligence, and monitor third-party AI vendors given provider concentration and increasing implementation and embedding of AI across the supply chain (*i.e.,* to 4th and nth party vendors).

1) AI regulation should acknowledge both the risks and the benefits of AI, reflect a "materiality and risk"-based approach that does not treat AI as a risk in and of itself and recognize the risks of *not* adopting AI.

We support the flexible approach to materiality and risk assessment in Sound Practice 5, which provides that "[f]inancial institutions can develop a specific framework for assessing the materiality and risk of their AI use cases, or adapt and enhance relevant frameworks to take account of AI-specific considerations and risks."[18] As AI technologies, including AI-enabled threats, continue to evolve, the risk to financial institutions from delayed adoption of advanced AI tools in areas ranging from cyber defense, fraud detection, anti-money laundering effectiveness and operational resilience capabilities will increase and will be more severe. A balanced regulatory approach therefore requires weighing the risks of financial institutions' inaction on AI alongside the risks of adoption.

In particular, we appreciate the FSB's nuanced approach to explainability and transparency (Sound Practice 8), which aligns with our view that there is no one-size-fits-all solution and acknowledges that trade-offs exist between model performance and

[18] Report, *supra* note 2 at 28.

explainability.[19] We view Sound Practice 8 as a strong example of how materiality and risk-based calibration can be applied to a specific risk vector.

We agree that "the importance of explainability may vary depending on the level of autonomy, materiality, and risk of different AI models and systems" and that "[t]he optimal level of transparency may vary depending on the intended recipients of the information."[20] Appropriate transparency should be risk-based and context-specific, and the level and type of disclosure may vary depending on the recipient, nature of the product or service, extent to which AI materially influences the relevant decision or outcome and/or applicable legal or regulatory requirements. For certain AI systems, particularly foundation models (*e.g.,* large language models), effective oversight may rely more heavily on outcome-based review, continuous monitoring and use-case level controls, rather than full transparency into the inner workings of a model.

Banks currently calibrate explainability controls to risk, with higher expectations for credit-related, customer facing or regulatory compliance-related functions and lighter expectations for low-risk productivity tools. And, to promote transparency across stakeholders, banks maintain role-appropriate materials such as technical documentation for model owners and validators, business-level summaries for governance forums, and customer-facing disclosures and notices consistent with regulatory requirements. Consistent with Sound Practice 8, when setting expectations regarding explainability and transparency based on materiality and risk, financial regulators should take into account the particular circumstances to which they are being applied.[21] We caution against imposing blanket explainability or transparency obligations on financial institutions triggered solely by using AI (or a particular

[19] Report, *supra* note 2 at 33.

[20] Report, *supra* note 2 at 33, 36.

[21] Relatedly, regulators should consider distinguishing between expectations and obligations applicable to developers and deployers of AI tools, given the distinct spheres of visibility, control and risk profiles inherent in these roles with respect to a given AI use case. The Colorado Automated Decision-Making Technology Act ("CADMT Act") reflects such a distinction. *See* Colorado Automated Decision-Making Technology Act, C.R.S. § 6-1-1701 et seq. The CADMT Act, set to take effect on January 1, 2027, distinguishes between obligations applicable to developers of automated decision-making technologies used to materially influence a consequential decision ("covered ADMT") and those applicable to deployers of such technologies based on their respective roles and risk-management capabilities. Specifically, under the CADMT Act, developers of covered ADMT must provide certain technical documentation regarding the covered ADMT to deployers of the technology, including descriptions of categories of training data, known limitations and instructions for appropriate use and human review. CADMT Act, C.R.S. § 6-1-1702. Developers are also required to notify deployers of material updates or modifications to the covered ADMT. *Id.* Deployers, on the other hand, are subject to use- and consumer-facing obligations, including notice requirements, post-adverse-outcome disclosures, recordkeeping obligations and consumer rights of data correction and to request meaningful human review following a consequential decision made by the covered ADMT. CADMT Act, C.R.S. §§ 6-1-1703 – 1705. This delineation reflects the distinct roles and risk management capabilities of AI developers, who control model design and have greater visibility into training data, design choices and limitations, and deployers, who control context of use and have greater access to end-users.

form of AI). For example, the Report's suggestion that banks offer customers an "opt-out" for AI-enabled functions may not be appropriate or practicable in all circumstances.[22]

Further, we appreciate and reiterate the Report's acknowledgment that AI-powered tools themselves can be used by financial institutions as part of their cyber and information and communication technology risk-management (Sound Practice 11) and human oversight (Sound Practice 10) programs. Banks have seen tangible benefits from using AI tools to manage risk in these and other areas. For example:

- *Human Oversight (Sound Practice 10*): Financial institutions incorporate AI, including agentic processes, into their AI oversight programs and practices by integrating materiality and risk-based assessments, differentiating between oversight expectations for bounded agents performing constrained processes and for more autonomous agents. For bounded agents, financial institutions often rely on more standard testing, observability and outcome analysis. For agents with additional autonomy, oversight processes define more detailed decision lineage and autonomy policies specifying where agents may act without approval, where human sign-off is mandatory and what "kill switch" mechanisms must exist.

  Across the materiality and risk spectrum, firms have found "AI-in-the-loop" oversight (*e.g.,* the use of AI agents to oversee other agents) – with appropriate and established triggers and human monitoring at defined stages – to be even more accurate than human oversight in certain contexts. When leveraging AI agents for oversight, financial institutions determine guardrails and metrics that trigger timely human intervention and support "kill switch" mechanisms. These include action-level log monitoring, decision traceability, outcome metrics and execution velocity or resource consumption thresholds. Firms continuously adapt and evolve these guardrails and metrics as AI tools evolve, consistent with Sound Practice 4 (Organisational adaptability).

  Where appropriate for the materiality and risk profile of a given AI use case, governance of agents could 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.

- *Cyber and ICT Risk Management (Sound Practice 11):* Financial institutions adapt their cyber and ICT risk management practices on an ongoing basis to address both the risks and benefits of novel AI tools and systems, including frontier models, to cyber operations. For example, to support identity and access management, certain institutions assign AI agents distinct identities and manage them like human system users, with role-based access, permissions, data-use controls and performance monitoring aligned to the applicable use

[22] Report, *supra* note 2 at 36.

case. However, we caution regulators and standard-setters against over-emphasizing extreme tail-risk scenarios and note the significant opportunity costs associated with inefficient prioritization and resource allocation in this area. For example, although we welcome the Report's inclusion of a "defense in depth" approach involving multiple layers of cyber defense, we note that implementing controls like micro-segmentation in large, complex networks is resource-intensive and can sometimes lead to unintentional protection gaps or operational incidents.[23] While segmentation is a valuable control, we instead encourage the regulators to focus on the security outcome they seek to achieve, namely preventing threat actor lateral movement once compromising a victim's network.

- *Performance Management (Sound Practice 9)*: Financial institutions are adapting their performance management practices and frameworks in light of the probabilistic nature of GenAI and agentic AI, which do not lend themselves to validation and other traditional performance monitoring techniques. Rather than independently reviewing model inputs, processing and outcomes through traditional validation processes, which are ineffective for probabilistic AI models that can yield different outputs based on the same input, financial institutions define metrics and tracking mechanisms to measure AI performance against use-case objectives, risk thresholds and business outcomes. The Report correctly recognizes that "GenAI models or systems produce a range of potentially acceptable outputs but not necessarily a single, correct output."[24] However, human involvement and oversight to monitor a probabilistic AI model's performance is not always the most effective approach. Rather than sample-based review, financial institutions often use AI tools to monitor AI agent performance to enable near real-time oversight, tracking and analysis of 100% of activity. As a result, we caution against the imposition of prescriptive human oversight obligations as a compensating control. Instead, regulatory frameworks should permit flexibility to determine the best control strategy.

2) **Regulators should apply proportionality when determining regulatory and supervisory expectations for financial institutions' use of AI in a manner that ensures banks have the flexibility to leverage existing governance and risk management frameworks and proactively evolve the application of these frameworks to address new or differentiated risks of AI systems.**

We strongly support the Report's emphasis on proportionality. Proportionality should be a guiding pillar for any future regulatory framework. Given the diversity of AI use cases and technologies, from traditional machine learning to GenAI and highly autonomous agentic systems, and the range of institutions that use AI, from small, regional community banks to GSIBs, a one-size-fits-all approach to any of the Sound Practices would stifle innovation, undermine banks' ability to protect against and mitigate risks, and misallocate compliance resources. By promoting consistent cross-jurisdictional implementation of the Sound Practices

[23] Report, *supra* note 2 at 45.

[24] Report, *supra* note 2 at 37.

and ensuring that institutions have the flexibility to leverage existing risk management frameworks where appropriate, the FSB and global regulators will reduce fragmentation, improve implementation effectiveness and allow institutions to focus resources on managing material risks rather than reconciling differing regulatory expectations.

To avoid overly prescriptive implementation of the Sound Practices, we believe the following should be further clarified to emphasize the necessity of proportionality in practice:

- *Sound Practice 1 (Role of the Board and Senior Management).* We support the FSB's focus on "organization-wide AI governance." We understand that, given the range of governance systems to which the Report is relevant,[25] it does not delineate the responsibilities of a financial institution's board of directors and its senior management for AI governance. Instead, the Report may be interpreted to set an expectation that both the board and senior management participate in the day-to-day management of an institution's AI adoption.[26] In the United States, "[a] central tenet of good corporate governance is the distinction between the board's responsibility for *oversight* of the business and affairs of [a financial institution] and the board's delegation to management of the responsibility for the day-to-day *operations* of the" financial institution.[27] We recognize that there are different approaches to organizational governance in the United States compared to other jurisdictions. Nevertheless, we recommend that U.S. financial regulators distinguish between the oversight role of a financial institution's board and senior management's operational and execution roles governing an institution's AI adoption, use and strategy. For example, a financial institution's board should not be expected to "establish clear boundaries for AI adoption," as this function is appropriately performed by senior management.[28]

- *Sound Practice 3 (Risk Management and Documentation).* Documentation and tracking obligations for AI use, particularly through third-party vendors, should be proportional and principles-based. As AI is embedded across the systems and processes used by financial institutions, tracking and documenting every instance of AI use is not a realistic expectation,[29] particularly as vendors embed AI in their products and services, including in

[25] *See* Report, *supra* note 2 at note 33 (acknowledging the "significant differences in legislative and regulatory frameworks across jurisdictions regarding the functions of the board of directors and senior management.")

[26] *See, e.g.,* Report, *supra* note 2 at 18-19 ("The board and senior management are engaged end-to-end, from strategy through implementation…").

[27] Exposure Draft, *Guiding Principles for Enhancing U.S. Banking Organization Corporate Governance*, BANK POLICY INSTITUTE (Jan. 12, 2021).

[28] Report, *supra* note 2 at 19.

[29] Experience in the E.U. shows that over-prescriptive reporting and documentation rules can quickly become counterproductive. The European Commission's own "fitness check" found that supervisory reporting requirements were not fully aligned, occasionally overlapping or inconsistent with one another, while recent E.U. work in this area has moved toward simplification because the burden was not always delivering better

ancillary tools and systems they may not view as sufficiently "material" to disclose. As a result, it could become excessively difficult or impossible for financial institutions to obtain meaningful information about the ways in which 3rd, 4th and nth parties implement AI in these products and services. Effective AI risk controls should prioritize output management, including management of intermediate outputs in agentic workflows, rather than exhaustive inventories and static documentation rules, and focus on actual outcomes rather than procedural requirements.

- *Sound Practice 4 (Organisational Adaptability).* We fully agree that it is necessary for financial institutions to "learn, adapt and adjust their oversight, governance, risk management practices, and capabilities as AI evolves."[30] This requires a regulatory framework that is adequately flexible and principles-based so institutions can leverage and evolve their existing oversight, governance and risk management practices to keep up with rapid advances in AI technology. Rules or guidance with granular and specific requirements will quickly become outdated, with little benefit to risk management, as technology and practices evolve. The bureaucratic processes created by rigid, sector-wide documentation standards lacking proportionality serve as one example of requirements that rapidly become obsolete, delay AI adoption and hinder healthy innovation.

  Further, to facilitate organizational adaptability, we strongly support public-private collaboration to enable shared learning and capability-building between industry and the public sector, helping both sectors "learn, adapt, and adjust" as AI continues to evolve.[31]

- *Sound Practice 7 (Data Governance and Data Management).* We agree that effective data management and data governance are essential to AI risk management. However, we encourage regulators and the FSB to avoid suggesting that all AI requires a separate or bespoke data governance framework or that data in AI necessarily must be governed differently. In addition, we recommend that regulators and the FSB acknowledge that data management and governance practices should be proportionate to materiality, risk, and intended use of the AI application, and that some practices – such as determining data provenance – are in nascent stages.

---

supervisory insight. *See* Commission Staff Working Document, Fitness Check of EU Supervisory Reporting Requirements (June 11, 2019).

[30] Report, *supra* note 2 at 23.

[31] The Financial Stability Oversight Council's recent Artificial Intelligence Innovation Series are important case studies in this regard. *See Artificial Intelligence Innovation Series*, U.S. DEPARTMENT OF THE TREASURY, *available at* https://home.treasury.gov/policy-issues/financial-markets-financial-institutions-and-fiscal-service/financial-stability-oversight-council/council-work/artificial-intelligence-innovation-series. We view the Innovation Series as a sound example of public-private engagement to support shared learning and the collective understanding of the risks, benefits and challenges presented by new technologies to inform the development of new rules and guidance.

1) **Supervisory expectations established in the regulatory framework should be calibrated to recognize the limitations on financial institutions' ability to control, diligence, and monitor third-party AI vendors given provider concentration and increasing implementation and embedding of AI across the supply chain (i.e., to 4th and nth party vendors).**

The Report's focus on financial institutions' leveraging existing technology-neutral third-party risk management guidance is a sound starting point.[32] However, the Report does not fully reflect the practical limitations that financial institutions face across the AI lifecycle when managing a third-party. Expanded AI adoption by vendors that financial institutions rely on means risks can originate not only from third parties selected and engaged by a financial institution, but from nth parties far beyond a financial institution's reasonable visibility or control. Regulators and policymakers should recognize that it may not be feasible for a financial institution to assess risk across every link in complex, evolving vendor supply chains, particularly when there is limited visibility into a third party's proprietary model design and training process. We therefore encourage the FSB to calibrate supervisory expectations to reflect those practical constraints and recognize alternative risk management approaches commensurate with the level of visibility and control available to the institution. In addition, we have previously recommended that U.S. regulatory guidance on financial institutions' third-party obligations focus on assessing the third-party's own vendor risk management program.[33] At the same time, regulators should use their supervisory authorities, including under the Bank Service Company Act in the United States, to help promote additional controls and transparency for upstream providers.

We would welcome continued regulatory focus on improving transparency and information flows across the AI lifecycle, including practical expectations on the information deployers need to obtain from AI developers and other third-parties to support governance, assurance and regulatory compliance. While financial institutions remain accountable for the responsible AI deployment and use, accountability should not require full visibility or control over every component of the AI lifecycle. Effective risk management will require shared responsibility across AI developers, infrastructure (*e.g.,* cloud) providers and deployers, with expectations calibrated to each participant's ability to obtain information, exercise oversight and influence outcomes. Specifically, we support the development of the "AI Shared Responsibility Model" referenced in the Report to help clarify roles across financial institutions, AI developers and infrastructure partners.[34] Such a framework can enhance transparency and risk management, particularly for third-party and foundation models. To this end, it is essential

[32] Report, *supra* note 2 at 48.

[33] Letter from BPI to the FRB, FDIC and OCC, re: Proposed Interagency Guidance on Third-Party Relationships: Risk Management 40–42 (Oct. 18, 2021), https://bpi.com/wp-content/uploads/2021/10/BPI-Issues-Comment-Letter-in-Response-to-Proposed-Interagency-Guidance-on-Third-Party-Relationships.pdf.

[34] Report, *supra* note 2 at 49.

that responsibilities are assigned to the parties best positioned to manage the underlying risks, recognizing that financial institutions may have limited visibility into proprietary model design and training processes. Any such approach should also be supported by appropriate transparency and standardized assurance mechanisms to ensure consistency and practical implementation.

Beyond limited visibility into upstream dependencies, financial institutions face similar challenges with validating and substituting third-party AI systems. The Report's recommendation in Sound Practice 12 (Third party AI risk-management) that "financial institutions develop robust exit strategies which may include alternative third-party AI providers or the ability to continue providing a given AI-enabled service manually in the event of disruption or termination" may not be practical.[35] Switching to an alternative provider potentially introduces different performance characteristics, biases and operational behaviors (and risks) into a financial institution's systems. Substitution is further complicated by concentrated dependence on a limited number of technology vendors including large language model developers and cloud service providers. Regulatory and supervisory expectations should reflect an understanding that, in certain cases, substitution of AI vendors may not be a viable option for a financial institution without significantly compromising their AI capabilities and essential business functions (including customer service and ICT risk mitigation).

Accordingly, the Report should clarify that vendor substitutability may not always be practical, and that other aspects of an institution's third-party risk management program – such as continuity planning, service degradation protocols, fallback controls, contractual protections and incident response – can be adequate where vendor substitution is not feasible. The FSB acknowledged 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" 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."[36]

Relatedly, the onus should not only be on financial institutions to manage concentration risks from AI vendors because this is a sector-wide issue of systemic concern that banks have limited ability to manage or mitigate. Given that many financial institutions (as in other industries) are increasingly reliant on a small group of large, powerful AI vendors, a single financial institution's ability to effectively negotiate robust safeguards, transparency measures and control rights over critical AI vendors is inherently limited.

---

[35] Report, *supra* note 2 at 50.

[36] See FINANCIAL STABILITY BOARD, ENHANCING THIRD-PARTY RISK MANAGEMENT AND OVERSIGHT: A TOOLKIT FOR FINANCIAL INSTITUTIONS AND FINANCIAL AUTHORITIES at 29 (Dec. 4, 2023), https://www.fsb.org/2023/12/final-report-on-enhancing-third-party-risk-management-and-oversight-a-toolkit-for-financial-institutions-and-financial-authorities/.

These limitations may also impact a financial institution's ability to implement certain governance practices, including those highlighted in the Sound Practices, due to constraints on the information provided by third-party AI vendors. Sound Practice 7, for example, appropriately highlights the challenges that financial institutions face when relying on third-party AI models where training data is proprietary and inaccessible.[37] Regulators and standard setters, such as the FSB, should therefore support financial sector resilience against AI supply chain concentration risks by standardizing data transparency practices for AI vendors. Developing "nutritional labels," standardized "data cards," "AI Bill of Materials" or equivalent disclosures that provide financial institutions with sufficient information to assess third-party data quality in a standardized format would be beneficial. Regulators also should consider how to encourage harmonization across AI vendor security guardrails and control interfaces, as financial institutions face significant operational friction due to fragmented security controls across major AI providers.

Finally, we encourage financial regulators and the FSB to provide practical, realistic steps financial institutions can take to manage third-party risks with vendors over which they have limited visibility and contractual bargaining power. To do so, regulators might leverage lessons learned from developing supervisory approaches to cloud service and financial market infrastructure providers, both of which share analogous concentration and systemic risk considerations.

* * * * *

The Bank Policy Institute and Institute of International Bankers appreciate the opportunity to comment. If you have any questions, please contact Heather Hogsett at Heather.Hogsett@bpi.com or Michelle Meertens at mmeertens@iib.org.

Respectfully submitted,

/s/ Heather Hogsett

Heather Hogsett
Executive Vice President & Head of BITS
Bank Policy Institute

/s/ Michelle Meertens

Michelle Meertens
Deputy General Counsel
Institute of International Bankers

[37] *See* Report, *supra* note 2 at 15.