

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

Response to Consultation

Naresh R (independent response)

1. Do you agree with the benefits and risks of AI adoption by financial institutions described in this report? Are there any substantive benefits or risks not covered?

Overall, I agree that the report provides a balanced and comprehensive discussion of both the benefits and risks of AI adoption by financial institutions.

Strengths:

1. Balanced coverage of AI benefits and risks across the financial sector. Sections 1.2 and 1.3 appropriately recognize the potential for AI to improve operational efficiency, customer service, fraud detection, cybersecurity, and risk management, while also addressing risks related to data, explainability, governance, third-party dependencies, and emerging agentic AI capabilities.
2. Recognition that AI risks extend beyond traditional model-performance concerns. The report appropriately reflects that AI-enabled systems may create operational, conduct, compliance, cyber, and customer risks and contribute to financial-stability risks, in addition to risks arising from model accuracy or performance.
3. Forward-looking treatment of agentic AI and interconnected systems. The discussion of autonomy, unauthorized actions, interactions among multiple AI agents, and rapidly evolving third-party dependencies provides a strong foundation for considering how AI risk may change as systems gain greater capabilities and operational influence.

Areas for Clarification:

1. Further distinguish model-output risk from AI execution risk. The report already discusses unauthorized actions by AI agents. It could further illustrate the distinction between risks arising from inaccurate or biased outputs and risks arising when an AI-enabled process executes unauthorized, incorrectly sequenced, excessive, or difficult-to-reverse actions. This distinction may help institutions identify the different controls needed for model performance, permissions, transaction limits, approvals, intervention, and reversibility.
2. Further emphasize the complete AI-enabled business process as the object of governance. The report recognizes AI systems, agents, tools, and interactions among

multiple AI agents. It could additionally note that institutions may benefit from assessing the full AI-enabled process, including models, agents, tools, data, external services, business rules, permissions, and human decision points, rather than reviewing components only in isolation. This may help identify interaction and orchestration risks that separate component-level reviews may not reveal.

3. Provide additional illustrations across the spectrum of third-party AI arrangements. The report already provides thoughtful coverage of third-party AI risk and differing levels of transparency and control. Additional examples could show how governance considerations may vary across institution-hosted open-source models, customized vendor solutions, external model APIs, foundation-model services, and AI embedded within broader vendor products. This may help institutions tailor due diligence, testing, monitoring, contractual protections, and contingency planning proportionately.

4. Further distinguish institution-level controls from shared ecosystem risks. The report appropriately discusses concentration risk and third-party dependencies. It could further illustrate the distinction between risks that institutions can reasonably manage through their own governance frameworks and shared dependencies, such as common foundation models, cloud providers, or subproviders, that may also benefit from supervisory dialogue, industry collaboration, or information sharing. This distinction may support both effective institutional risk management and broader financial-system resilience.

Overall, the report provides a strong foundation for understanding the benefits and risks of AI adoption. The clarifications suggested above are intended to provide additional implementation context that may support consistent application while preserving the report's flexible, proportionate, and technology-neutral character.

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

Overall, I believe the 12 Sound Practices provide a comprehensive, well-structured, and principles-based framework for the responsible adoption of AI in financial services. They appropriately balance flexibility with sound governance expectations and should remain relevant as AI technologies and use cases continue to evolve.

Strengths:

1. Comprehensive coverage of organization-wide governance and the AI lifecycle. The Sound Practices collectively address strategy, governance, accountability, materiality, system selection, data, testing, monitoring, explainability, human oversight, cyber and ICT risk, and third-party AI risk in a coherent and logical manner.

2. Appropriate emphasis on proportionality and a technology-neutral approach. The framework appropriately avoids prescriptive technical requirements and instead encourages institutions to tailor governance according to the materiality, complexity, and risk of specific AI use cases.

3. Strong recognition of enterprise-wide governance and accountability. Practices 1–3 appropriately position AI governance as an organization-wide responsibility, with roles for

the board, senior management, business owners, technology teams, and independent control functions, rather than solely as a technology or model-development activity.

Areas for Clarification:

1. Provide additional implementation considerations for materiality assessments. Practice 5 already provides strong principles for assessing AI materiality and risk. The report could further illustrate a non-exhaustive set of considerations, such as customer impact, business criticality, action authority, autonomy, reversibility, data sensitivity, and third-party dependence. This may promote greater consistency while preserving institution-specific judgment.

2. Further emphasize governance of the complete AI-enabled business process. Practices 3, 5, 6, and 9 collectively address AI identification, materiality and risk assessment, system selection, testing, and ongoing monitoring. The report could further illustrate that institutions may benefit from assessing the complete AI-enabled business process, including interacting models, agents, tools, data, external services, business rules, permissions, and human decision points, rather than reviewing individual AI components in isolation. This may help identify interaction and orchestration risks that component-level reviews may not reveal.

3. Provide additional illustrations of coordinated assurance across existing governance functions. Practice 2 already provides strong guidance on governance, accountability, the three lines of defense, and cross-functional coordination. Additional implementation examples showing how business owners, model risk management, operational risk, compliance, legal, cyber, data, privacy, third-party risk, and internal audit may coordinate their respective responsibilities could help institutions strengthen enterprise-wide governance while maintaining clear accountability and avoiding unnecessary duplication.

4. Further illustrate governance across the spectrum of third-party AI arrangements. Practice 12 already recognizes differing levels of institutional control and visibility. Additional examples could demonstrate how governance approaches may vary across institution-hosted open-source models, customized vendor solutions, external model APIs, foundation-model services, and AI embedded within broader vendor products. This may help institutions tailor due diligence, testing, monitoring, contractual protections, and contingency planning according to the level of transparency and control available.

5. Provide additional examples of reporting that supports effective board oversight and senior-management decision-making. Practices 1 and 2 provide strong governance expectations for the board and senior management. Additional illustrations of information that may be useful to them, such as aggregate AI exposure, material AI use cases, significant incidents, third-party concentrations, or material capability changes, may help institutions translate these principles into proportionate governance practices while remaining adaptable to different jurisdictional governance structures.

Overall, the Sound Practices provide a strong foundation for responsible AI governance across the financial sector. The implementation considerations suggested above are intended to support more consistent application while preserving the framework's principles-based, proportionate, and technology-neutral character.

3. Do the sound practices strike an appropriate balance between managing risks relating to all forms of AI, and addressing some of the risks relating to emerging and new complex forms of AI, such as GenAI and agentic AI?

Overall, I believe the report strikes an appropriate balance between technology-neutral guidance and discussion of emerging AI technologies. The principles-based approach should enable the Sound Practices to remain relevant as AI capabilities, architectures, and deployment models continue to evolve.

Strengths:

1. Appropriate balance between enduring principles and emerging technologies. The report successfully applies consistent governance principles across different forms of AI while providing targeted discussion of GenAI and agentic AI where their capabilities introduce additional governance considerations.
2. Strong recognition of emerging capabilities associated with agentic AI. The discussion of autonomy, planning and execution, interactions among AI agents, ongoing monitoring, and human oversight appropriately recognizes that governance challenges increasingly arise from how AI systems operate in practice rather than from the underlying models alone.
3. Technology-neutral guidance supports long-term applicability. By focusing on governance outcomes rather than specific technologies, model architectures, or implementation techniques, the framework is well positioned to remain useful as AI continues to evolve.

Areas for Clarification:

1. Provide additional implementation illustrations based on AI capabilities rather than technology labels. As AI terminology continues to evolve, institutions may benefit from examples organized around capabilities such as planning, memory, tool use, delegation, external connectivity, and action authority rather than particular model types. This could help preserve the report's technology-neutral approach while supporting more consistent implementation.
2. Further illustrate that autonomy exists along a continuum. The report appropriately discusses agentic AI and human oversight. Additional examples showing how governance considerations may evolve as systems progress from assistance to recommendation, constrained execution, and broader operational authority could help institutions apply proportionate controls without creating rigid classifications.
3. Reinforce the distinction between model explainability and operational traceability. The report already recognizes the importance of transparency and traceability of AI-enabled activities. It could further illustrate that, particularly for complex agentic or third-party AI, operational traceability, including logs, tool usage, execution pathways, and decision records, may complement model explainability as an important governance mechanism.
4. Illustrate how combinations of capabilities may influence governance. Rather than considering autonomy alone, institutions may also evaluate combinations of capabilities, such as access to sensitive data, the ability to invoke external tools, interaction with

customers, or authority to initiate actions. Additional examples could help institutions apply proportionate governance based on the overall risk presented by an AI-enabled process.

Overall, the report successfully balances technology-neutral guidance with discussion of emerging AI capabilities. The implementation considerations suggested above are intended to help institutions apply the Sound Practices consistently as AI technologies continue to evolve, while preserving the framework's flexible, principles-based, and future-oriented approach.

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

Overall, I believe the Sound Practices are sufficiently flexible to accommodate future developments in AI. Their principles-based, technology-neutral, and proportionate approach should enable the framework to remain relevant as AI capabilities, deployment models, and operating environments continue to evolve.

Strengths:

1. Principles-based guidance provides a durable foundation. By focusing on governance outcomes rather than specific technologies or implementation methods, the framework is well positioned to accommodate future advances in AI without requiring frequent revisions.
2. Recognition that AI governance must evolve over time. Practice 4 appropriately encourages institutions to continuously strengthen their AI governance capabilities and adapt their frameworks as technologies, business use cases, and risk profiles evolve.
3. Appropriate emphasis on ongoing monitoring and lifecycle governance. The report recognizes that responsible AI governance extends beyond initial deployment through ongoing monitoring, periodic review, and continued risk management. This provides an important foundation for governing AI systems whose capabilities and operating environments may change over time.

Areas for Clarification:

1. Provide additional implementation illustrations for reassessment following material changes. The report already recognizes the importance of ongoing monitoring and reassessment. It could further illustrate circumstances institutions may consider when determining whether renewed assessment is appropriate, such as significant changes to AI capabilities, providers, permissions, data sources, business use cases, or operating environments. This may promote more consistent implementation while preserving institutional judgment.
2. Further illustrate governance of evolving AI-enabled business processes. As AI capabilities expand through new tools, agents, external services, or business workflows, institutions may benefit from periodically reassessing the complete AI-enabled business process rather than only individual AI components. Additional implementation examples could help institutions identify emerging interaction, orchestration, and dependency risks as AI ecosystems evolve.

3. Provide additional implementation considerations for changing third-party dependencies. AI services may evolve through model updates, newly embedded functionality, changes in subproviders, or revised orchestration arrangements. Additional illustrations could help institutions consider how existing third-party governance and change-management processes may respond proportionately to these developments while preserving the report's principles-based approach.

4. Encourage continued sharing of implementation experience. As AI technologies and governance practices mature, periodic updates to illustrative case studies, implementation examples, or observations drawn from supervisory and industry experience could support greater consistency across jurisdictions without introducing prescriptive expectations or reducing institutional flexibility.

Taken together, the Sound Practices appear well designed to remain relevant as AI continues to evolve. The additional illustrations suggested above would help institutions better recognize when changes in capabilities, dependencies, or deployment context warrant renewed attention, without reducing the framework's flexibility or introducing prescriptive expectations.

5. Do the case studies in this report sufficiently highlight how different types of financial institutions can benefit from responsible AI adoption? Are there additional case studies for inclusion in the report? If so, please provide such case studies, particularly for nonbanks.

Overall, the case studies provide practical and relevant illustrations of how financial institutions can apply the Sound Practices to support the responsible adoption of AI across a range of business activities and risk areas.

Strengths:

1. Broad coverage across AI use cases and governance topics. The case studies include examples involving credit risk, fraud detection, customer relationship management, insurance operations, trading, cybersecurity and operational resilience, AI inventories, explainability, testing, and third-party risk management. This demonstrates the relevance of the Sound Practices across a diverse range of financial activities.

2. Practical connection between principles and implementation. The examples effectively illustrate how high-level governance principles may be translated into risk assessments, controls, monitoring arrangements, and organizational practices.

3. Recognition of proportionate implementation. Several case studies appropriately reflect that responsible AI adoption may vary according to an institution's size, business model, risk profile, available resources, and level of AI maturity.

Potential Improvements:

1. Broaden representation across additional nonbank financial-sector participants. While the report discusses AI adoption across several financial sectors, many of the detailed institutional case studies focus on large or internationally active banks. Additional case studies involving asset managers, securities firms, payment providers, regulated fintechs,

insurers, and financial market infrastructures could further demonstrate how different types of financial institutions may benefit from applying the Sound Practices.

2. Include more examples of proportionate implementation by smaller institutions. Illustrations showing how smaller institutions may use narrower deployments, standardized evidence, shared assurance resources, proportionate vendor controls, or industry resources could demonstrate how the Sound Practices remain practical without requiring the governance infrastructure of larger organizations.

3. Include examples of challenging governance decisions and non-deployment outcomes. Case studies involving paused deployments, rejected vendors, significant control concerns, near misses, or decisions not to proceed could provide valuable insights into how responsible AI adoption also includes recognizing when risks outweigh expected benefits.

4. Expand illustrations across the spectrum of third-party AI arrangements. The existing discussion of customized vendor solutions, software-as-a-service, embedded AI, and open-source arrangements is valuable. Additional examples involving external foundation-model APIs, nested AI services, or less transparent subprovider dependencies could further illustrate how governance and assurance approaches may vary according to an institution's level of visibility, control, substitutability, and operational dependence.

5. Illustrate enterprise-wide governance in practice. At least one case study could demonstrate how business owners, technology teams, model risk management, compliance, legal, cyber, data, operational risk, and third-party risk functions coordinated their respective responsibilities and resolved differing risk perspectives. This could make the report's enterprise-wide governance expectations more tangible while preserving flexibility in institutional operating models.

Overall, the case studies already provide meaningful practical value by demonstrating how the Sound Practices can be applied across a variety of AI use cases. Broader representation of nonbank and smaller institutions, together with examples of more challenging governance outcomes, could further strengthen their relevance across the financial sector.

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

Overall, the case studies provide practical and actionable insights that help translate the Sound Practices into real-world governance and risk management activities. They complement the principles-based framework by illustrating how institutions may approach responsible AI adoption without prescribing a single implementation model.

Strengths:

1. Practical illustrations that support implementation. The case studies effectively demonstrate how the Sound Practices may be translated into governance processes, risk assessments, controls, and ongoing monitoring.

2. Appropriate balance between flexibility and practical guidance. The examples provide useful implementation context while preserving the principles-based and proportionate

nature of the framework, allowing institutions to adapt the practices to their own operating models and risk profiles.

3. Coverage of multiple governance dimensions. Collectively, the case studies demonstrate that responsible AI adoption extends beyond technical performance to include governance, third-party risk, human oversight, operational resilience, and organizational accountability.

Potential Improvements:

1. Consider using a more consistent structure across case studies. An illustrative common format covering the business objective, AI use case, materiality considerations, governance approach, key controls, monitoring activities, and implementation lessons could make the examples easier to compare and adapt while remaining non-prescriptive.

2. Provide additional context on why particular governance approaches were appropriate. Briefly explaining the factors that influenced governance decisions, such as business criticality, customer impact, autonomy, third-party dependence, or data sensitivity, could help institutions better understand how proportionality is applied in practice.

3. Include additional examples of ongoing monitoring and governance over time. Illustrations showing how institutions respond to material changes in AI capabilities, third-party services, business use cases, or operating environments could reinforce that responsible AI adoption is an ongoing governance process rather than a one-time assessment.

4. Expand practical examples of enterprise-wide governance. Case studies could further illustrate how business owners and existing governance functions, including technology, model risk management, compliance, legal, cyber, data, operational risk, third-party risk, and internal audit, contribute complementary perspectives throughout the AI lifecycle. This could help institutions translate the report's governance principles into effective operating practices without implying a single organizational model.

5. Include examples of management reporting and governance escalation. Brief illustrations of how material AI risks, significant incidents, emerging concentrations, or material changes in AI capabilities may be communicated to senior management and the board could help institutions apply the governance principles in a proportionate and decision-useful manner.

Taken together, the case studies offer a useful bridge between the report's high-level principles and institutional implementation. The improvements suggested above could strengthen that bridge by helping institutions understand not only what sound governance may look like, but also why particular approaches may be appropriate in different circumstances.

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

Overall, the glossary provides a clear and useful foundation for establishing a common understanding of the key AI concepts used throughout the report. The definitions are generally consistent with current industry terminology and appropriately support the report's principles-based and technology-neutral approach.

Strengths:

1. Clear and accessible terminology. The glossary defines key AI concepts in a manner that is understandable to a broad audience while remaining sufficiently precise for financial-sector use.
2. Alignment with current industry practice. The terminology is generally consistent with widely accepted AI governance frameworks and reflects contemporary AI capabilities without being tied to specific technologies or vendors.
3. Appropriate support for the report. The glossary provides sufficient context to support consistent interpretation of the Sound Practices without becoming overly technical or prescriptive.

Potential Improvements:

1. Expand terminology related to evolving AI architectures. The glossary already provides a useful definition of agentic AI. Future editions could consider distinguishing an individual AI agent from the broader concept of agentic AI and adding related terms such as multi-agent systems, AI orchestration, and AI-enabled business processes. These concepts are increasingly relevant as institutions deploy more complex AI solutions, and greater consistency in their definitions could help reduce fragmentation in industry and supervisory interpretations across jurisdictions.
2. Clarify terminology describing third-party AI arrangements. Additional definitions or examples distinguishing externally hosted AI services, embedded AI functionality, foundation-model APIs, and open-source AI deployments could support more consistent discussions of third-party governance and risk management.
3. Continue periodic review of the glossary. Given the pace of AI innovation, periodically reviewing and updating the glossary to reflect emerging terminology and evolving industry practice would help maintain its relevance while preserving the report's technology-neutral approach.

Overall, the glossary is well aligned with current industry practice and provides a strong common vocabulary for applying the Sound Practices. Selective additions over time could help it remain useful as AI architectures, deployment arrangements, and associated governance practices continue to develop.

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

Overall, the proposed Sound Practices provide a timely, balanced, and practical framework for supporting the responsible adoption of AI across the financial sector. The report appropriately recognizes both the opportunities and risks associated with AI while preserving flexibility for institutions to tailor their governance arrangements to their size, business model, risk profile, and level of AI maturity.

Strengths:

1) Balanced treatment of innovation and risk. The report recognizes the potential benefits of AI while emphasizing effective governance, accountability, human oversight, third-party risk management, and ongoing monitoring.

2) Strong principles-based and proportionate approach. By focusing on governance outcomes rather than prescribing specific organizational structures or technical methods, the Sound Practices can be applied across institutions with different operating models and levels of complexity.

3) Valuable contribution to international consistency. The report provides a common foundation that can support greater alignment across jurisdictions while remaining adaptable to local legal and supervisory frameworks.

Potential Improvements:

1) Continue developing implementation support over time. Periodic publication of additional case studies, implementation observations, or thematic examples could help institutions interpret and apply the Sound Practices as AI capabilities and deployment approaches evolve, without requiring frequent revisions to the core framework.

2) Place additional emphasis on evidence of governance and control effectiveness. Institutions may benefit from demonstrating not only that governance frameworks exist, but also that monitoring, control testing, independent challenge, escalation, and oversight processes are operating effectively in practice. Illustrative, non-prescriptive examples could support this objective while preserving institutional flexibility.

3) Further illustrate the distinction between institution-specific and shared ecosystem risks. Institutions can directly manage many risks through their own governance and controls. However, dependencies involving common cloud services, foundation models, data providers, or critical subproviders may create risks that cannot be fully addressed by an individual institution acting alone and may also benefit from industry collaboration, supervisory dialogue, and information sharing.

4) Reinforce governance of complete AI-enabled business processes. As institutions combine models, agents, tools, data, permissions, third-party services, business rules, and human decision-making, governance may need to consider interactions and cumulative risks across the complete process rather than assessing each component only in isolation.

5) Recognize the potential role of complementary assurance mechanisms. For high-consequence uses involving probabilistic AI, institutions may benefit from considering mechanisms that support policy enforcement, independent verification, traceability, or reproducible evidence. Non-exhaustive examples could illustrate how such mechanisms may strengthen assurance without favoring a particular technology or prescribing a specific control approach.

Taken as a whole, the report offers a strong foundation for responsible AI adoption in finance and should contribute meaningfully to greater consistency across institutions and jurisdictions. Its long-term value could be further strengthened through continued

implementation dialogue and periodically refreshed examples, while retaining the flexibility, proportionality, and technology neutrality that are central to the framework.