



February 2, 2024

Submitted via Regulations.gov
Information Technology Laboratory
ATTN: AI E.O. RFI Comments
National Institute of Standards and Technology
100 Bureau Drive, Mail Stop 8900
Gaithersburg, MD 20899–8900

Re: Request for Information (RFI) Related to NIST’s Assignments Under Sections 4.1, 4.5 and 11 of the Executive Order Concerning Artificial Intelligence (Sections 4.1, 4.5, and 11)

Dear Dr. Prabhakar:

The Securities Industry and Financial Markets Association (“SIFMA”)¹ welcomes the opportunity to respond to the National Institute of Standards and Technology (“NIST”) request for public comment on artificial intelligence (the “Request”).² SIFMA recognizes that maintaining public trust in AI applications is essential to realizing the enormous benefits that AI has to offer, and that recent developments in AI across economic sectors support the establishment of certain controls.

SIFMA, on behalf of its members, encourages policymakers to consider the following key issues in any future regulatory policies involving AI:

- Policymakers should adopt a governance framework for AI that is risk-based and not overly prescriptive. An AI governance framework should treat AI models, algorithms, applications, and systems (collectively, “AI applications”) differently depending on the likelihood or

¹ SIFMA is the leading trade association for broker-dealers, investment banks, and asset managers operating in the U.S. and global capital markets. On behalf of our members, we advocate for legislation, regulation, and business policy affecting retail and institutional investors, equity and fixed income markets, and related products and services. We serve as an industry coordinating body to promote fair and orderly markets, informed regulatory compliance, and efficient market operations and resiliency. We also provide a forum for industry policy and professional development. SIFMA, with offices in New York and Washington, D.C., is the U.S. regional member of the Global Financial Markets Association (GFMA).

² *Request for Information (RFI) Related to NIST’s Assignments Under Sections 4.1, 4.5 and 11 of the Executive Order Concerning Artificial Intelligence* (Sections 4.1, 4.5, and 11) 88 Fed. Reg. 34194 (May 26, 2023), <https://www.govinfo.gov/content/pkg/FR-2023-05-26/pdf/2023-11346.pdf>



severity of the potential harm they might cause. Such a framework would subject low-risk AI applications to different compliance obligations than high-risk applications.

- An effective AI governance framework would cover certain foundational components, including scoping, inventory, risk assessments, training, documentation, and third-party risk management. Companies should have flexibility on how best to integrate these components with existing policies and functions, including enterprise risk governance programs, model risk, data governance, privacy, cybersecurity, vendor management, and product development, as well as auditing and third-party risk management practices.
- A framework that is overly prescriptive would subject every AI application to onerous risk assessments and audits that are unworkable, would waste resources on low-risk applications at the expense of effectively mitigating high-risk applications, and would potentially prevent AI applications from benefitting consumers and businesses.

I. Policymakers Should Provide for a Risk-Based AI Governance Framework

A. A risk-based approach provides accountability by balancing upside potential with downside risks

SIFMA believes that a risk-based approach to AI governance provides the necessary flexibility to balance the many potential risks with the many potential benefits and opportunities in deploying AI.

There are several benefits associated with a risk-based governance framework for AI, including (1) identification of specific risks a company should consider when deciding which AI applications are high risk, (2) consideration of how best to mitigate the risks associated with high-risk AI applications, and (3) identification of which AI applications carry unacceptable risks and should not be pursued.

The granular determinations of such considerations are best made by the company's management, with guidance from its applicable sectoral regulators. Regulators therefore should primarily guide companies to focus their efforts on identifying their highest-risk AI uses, and on mitigating the risks they present. Such risks may include legal and regulatory risk, which may also include operational, reputational, contractual, discrimination, cybersecurity, privacy, consumer harm, lack of transparency, and confidentiality risks. This flexible framework would provide accountability mechanisms to reduce risk, while also providing companies with space to innovate in collaboration with applicable sectoral regulators.

For example, in the financial services sector, many firms have already heavily invested in developing sophisticated frameworks for model risk management after considering risks associated with such models. and after financial regulators issued guidance, such as the *Supervisory Guidance on Model Risk Management*.³ The existing collaboration between financial institutions and their regulators on model

³ See, e.g., Federal Reserve SR Letter 1107; OCC Bulletin 2011-12; and FDIC FIL 22-2017.



risk management illustrates that a tailored yet flexible approach provides strong accountability measures that also allow industry to innovate.

B. AI guidelines and regulations should reflect the risk-based requirements in cybersecurity but avoid being overly prescriptive

Many effective cybersecurity guidelines and regulations provide companies with flexibility to implement risk-based requirements for policies and governance. However, when cybersecurity regulation is too prescriptive, compliance suffers and risk mitigation is less efficient and successful. So, although policymakers can take valuable lessons from cybersecurity in considering how to approach AI accountability, the applicability of those lessons is somewhat limited because cybersecurity risks and mitigation tend to be more universally applicable across organizations and industries. AI applications, on the other hand, vary significantly both within and across organizations and industries, and therefore present an extremely broad range of risks and mitigation options that are often not similar from company to company. As a result, general AI guidelines and regulations must offer even more flexibility and must be less prescriptive than cybersecurity guidance and regulations to be broadly effective.

Accordingly, a regulatory regime that subjects all AI applications to the same compliance obligations and audits would be a significant impediment to developing an effective AI accountability ecosystem, particularly within the already heavily regulated financial services industry. Many companies have, or will soon have, numerous AI applications to better serve clients and increase efficiency. Putting each application through a complicated, expensive, and time-consuming compliance process would waste resources on low-risk applications and be an inefficient way to address the real concern – the mitigation of risks associated with high-risk applications. Such a cost-heavy approach would also run the risk of centralizing the use of AI applications among large firms and limit the ability of startups to participate.

C. Adopting a risk-based approach reduces the need to define AI

Adopting a risk-based approach will have the additional benefit of reducing the importance of crafting a precise definition for AI which, even if achieved, is likely to become outdated in the near term and require additional resources to update. Risk rating of AI applications should focus on high-risk uses, rather than assessing the risk of the AI tool itself, because the same AI tool can be used to produce vastly different risk profiles depending on the manner and context of its use.

II. Risk-Based AI Governance Components

A risk-based AI governance framework could include the following foundational components:

1. **Scoping.** Companies need to be able to easily determine which AI applications fall within the scope of the framework as part of building their governance programs.
2. **Inventory.** Companies should prepare and maintain an inventory of their AI applications with sufficient detail to allow for those AI applications to be risk rated.



3. **Risk Rating.** Companies should have a process through which they can identify their highest-risk AI applications. The risks considered in this process should include legal and regulatory risks, which may also include operational, reputational, contractual, discrimination, cybersecurity, privacy, consumer harm, lack of transparency, and confidentiality risks.
4. **Responsible Person(s) or Committee(s).** Companies should designate one or more individuals or committees who are responsible for identifying and assessing their highest-risk AI applications and either accepting those risks, mitigating them, or determining that the risks are too high, and that the particular AI application should be abandoned.
5. **Training.** Companies should develop training programs to ensure that stakeholders are able to identify the various risks associated with their use of AI and the various options for reducing risk.
6. **Documentation.** Companies should maintain documentation sufficient to support an audit of the risk assessment program.
7. **Vendor Management.** Because companies are increasingly using AI applications that are provided in whole or in part by third parties, they should be required to develop reasonable third-party AI risk management policies that acknowledge the leverage third-party providers hold in negotiations and their valid concerns regarding disclosure of the inner workings of their intellectual property.

To the extent regulation or guidance on the above components is issued, it should be sufficiently flexible to allow companies to incorporate these components in a manner that best fits their existing governance and compliance in related areas such as model risk, data governance, privacy, cybersecurity, vendor management, and product development, with further guidance from their applicable sectoral regulators as needed.

A. Policymakers should not dictate how and when companies audit AI applications

In addition to the foundational components described above, auditing is a crucial third-line control that companies using AI should include in their AI governance. Companies should be permitted to determine the most effective manner of auditing their AI applications, including how and when to conduct periodic audits and who can conduct them. As many financial services companies have already implemented a three-lines-of-defense model, internal auditors who are already familiar with their company's AI applications and corresponding procedures and processes could efficiently and effectively perform audits. Such an approach would enhance the development of in-house AI expertise while avoiding the upskill time often necessary when relying on external auditors.

AI audits should focus on risk assessment programs rather than individual AI applications. Auditing each application will result in a misallocation of resources, with too much effort spent on low-risk AI (e.g., spam filters, graphics generation for games, inventory management, and cybersecurity monitoring) and not enough effort spent on high-risk AI (e.g., hiring, lending, insurance underwriting, and education admissions). Moreover, a cost-heavy universal audit requirement may



result in AI usage centralizing among large firms and the exclusion of smaller firms and startups from participating, simply because the costs of auditing each AI application would be too great.

B. Third-party risk management for AI applications should also be risk-based

AI applications that are provided by or for third parties constitute what the Request identifies as the “AI value chain” and add a layer of complexity to accountability. Companies should use the same principles applied to AI applications that are developed in-house for identifying risks associated with third-party AI applications and mitigate those risks through commercially reasonable diligence, audits, and contractual terms.

SIFMA notes that there are many parallels between the third-party risks for AI applications and cybersecurity, and that regulatory requirements for third-party cybersecurity risk mitigation may be instructive for AI applications. For example, NYDFS Part 500.11(a) requires covered entities to implement written policies and procedures that are designed to ensure that information systems and confidential data that are accessible to, or held by, third-party service providers are secure. SIFMA supports a similar, flexible, and principles-based approach to third-party AI risk management, with the applicable sectoral regulators providing additional specific requirements as needed, so long as these requirements recognize that third parties will be reluctant to disclose competitively sensitive aspects of their offerings.

III. Conclusion

The risk-based approach described above will provide a valuable, flexible framework through which companies and their applicable sectoral regulators can build more tailored AI governance and compliance programs that will ensure accountability and trust. It also will not stifle innovation or waste resources on low-risk AI applications at the expense of the important work that needs to be done to ensure that high-risk AI is meaningfully reviewed and effectively mitigated.

* * *

SIFMA greatly appreciates NIST’s consideration of these comments and would be pleased to discuss any of these views in greater detail if that would assist NIST’s deliberations on this issue. SIFMA would welcome the opportunity to continue to participate in this valuable process. Please feel free to contact me at mmacgregor@sifma.org if you would like to discuss these issues further.

Sincerely,

Melissa MacGregor

Melissa MacGregor
Deputy General Counsel & Corporate Secretary