



September 3, 2020

Marcia E. Asquith
Office of the Corporate Secretary
Financial Industry Regulatory Authority
1735 K Street, NW
Washington, D.C. 20006-1506

Re: Response to Request for Comments on FINRA Report, "Artificial Intelligence (AI) in the Securities Industry"

Dear Ms. Asquith:

The Securities Industry and Financial Markets Association ("SIFMA")¹ welcomes the opportunity to present its views to the Financial Industry Regulatory Authority Inc. ("FINRA") in response to FINRA's request for public comment on its recent report discussing the deployment of artificial intelligence ("AI") in the securities industry (the "AI White Paper").²

SIFMA believes that the AI White Paper is a thoughtful, detailed, and productive step forward in an ongoing process of developing the industry's understanding of the opportunities and challenges posed by AI technologies. The paper is especially valuable in describing and emphasizing the fast pace of ongoing innovation, the significant consumer benefits associated with AI deployment, and the extensive investment and attention that firms are already devoting to managing the risks associated with these new technologies and capabilities.

Part I of the comments that follow addresses the need for an appropriate approach to regulation of issues related to AI technologies and processes. It specifically addresses FINRA's request for comment (p. 20) regarding "areas where guidance or modifications to FINRA rules may be desired to support adoption of AI applications while maintaining investor protection and market integrity." In short, for the reasons set out in that section, issuing guidance or modifying FINRA rules at this juncture is not necessary. Firms operating in the securities industry are already subject to complex and often overlapping regulatory frameworks, and adding additional regulations would risk curtailing the innovation, consumer benefits, and compliance and risk management capabilities provided by AI—the very benefits that the AI White Paper itself identifies.

Part II of these comments addresses particular issues relating to points made by FINRA in the AI White Paper. These include the need for a common industry understanding of and lexicon for AI and the risks it

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 Artificial Intelligence (AI) in the Securities Industry (June 10, 2020), <https://www.finra.org/rules-guidance/key-topics/fintech/report/artificial-intelligence-in-the-securities-industry>.

may pose and benefits it will provide; the challenges that arise from applying existing rules and regulations to AI technologies; and questions of data governance and model risk management. Part II points to important considerations for FINRA as it develops its supervisory model; suggests initial solutions to the problems that firms have been facing; and, when necessary, requests clarification on individual sections of the AI White Paper.

Part III concludes with some thoughts about the most productive ways to continue the very useful process of sharing information and insights that FINRA has initiated.

I. General Approach to Regulation and Fostering Innovation

SIFMA commends FINRA's efforts to develop a deeper, common understanding of how AI is used by securities firms. FINRA's exploratory approach to the AI White Paper—requesting public comment, being transparent, and engaging member firms in the information gathering and development process—is an astute approach to gathering information about this rapidly evolving space. SIFMA encourages FINRA to continue engaging firms in this manner. Transparency and an ongoing dialogue between firms and FINRA will lead to a dynamic feedback process, which, in turn, will allow FINRA to more effectively and efficiently assist industry participants as they address opportunities and risks associated with the use of AI in securities business processes.

The AI White Paper also identifies several crucial characteristics of the current state of firms' efforts to develop AI. Perhaps most importantly, the paper appropriately identifies certain substantial benefits associated with AI (p. 11), including how "benefits include enhanced access to customized products and services, lower costs, access to a broader range of products, better customer service, and improved compliance efforts leading to safer markets." These benefits are, we believe, very much real and currently being realized, not just "potential."

The paper also appropriately recognizes that industry participants are engaged in a very extensive process of innovation to develop responsible AI capabilities that benefit investors, other consumers, and market integrity, and to explore and develop the oversight and other processes needed to reduce risks that can be associated with AI capabilities. Firms are, as the paper summarizes (p. 1), "allocating significant resources to exploring, developing, and deploying AI-based applications to offer innovative new products." The AI White Paper correctly recognizes (p. 2) that "AI is used as an umbrella term that encompasses a broad spectrum of different technologies and applications," and that those different technologies and applications are context-specific and present varying levels and types of risks and benefits.

These crucial insights recognized in the AI White Paper should lead FINRA to continue to pursue efforts to learn, develop, and disseminate information about AI practices and challenges rather than move to issue rules or offer overly prescriptive guidance regarding particular recommended practices. As the AI White Paper repeatedly notes, multiple legal requirements and regulatory regimes already exist to address the risks posed by AI, providing reason to refrain from adding additional regulatory requirements at this juncture. Very complex and often relatively open-ended legal requirements already apply to how firms ensure customer privacy, store information, defend against cyberattacks, inform customers and secure their consent where needed, and protect particularly sensitive types of information. Many of these existing regulations and enforcement processes apply specifically to firms. If FINRA were to contemplate imposing additional requirements or obligations, it would have to carefully assess how they may interact with existing requirements and the processes firms have developed to ensure compliance with them. More importantly, FINRA would have to identify some relevant regulatory gap in this overlapping patchwork of existing legal requirements. At this point, especially given the flexible and expanding nature of other legal and regulatory regimes, no such gap is apparent.

Regulation, either through rules or “soft” regulatory guidance, threatens to inhibit innovation and limit associated benefits for consumers and industry participants. No rules or even guidance could account for how AI is likely to develop or pinpoint the particular contexts that give rise to concern. Indeed, regulatory intervention at this juncture is likely to impose costs and inhibit innovation in unknown but potentially significant ways. These potential regulatory costs are especially significant because, as the AI White Paper identifies (pp. 8-10), innovative uses of AI may be especially valuable to firms, their customers, and markets by improving processes devoted to reducing risk and ensuring compliance.

Even issuing regulatory “guidance” would present the risks of impeding useful, efficient, and effective development of AI capabilities and imposing other costs associated with regulation. As the AI White Paper appears to recognize (for example, at p. 11), “guidance” can often harden into required practice or be invoked as a “best practice” when, in particular contexts, the guidance may be in fact inapplicable, overbroad, or ill-suited to address particular risks. As AI capabilities and risk identification so quickly change, these adverse consequences of inapt or prescriptive guidance could be especially pronounced and difficult to overcome.

The varied nature of AI capabilities and technologies and their varied uses by firms further counsel against the development of regulation or guidance, especially if addressed to particular technologies. As the AI White Paper notes, particular processes and deployments of AI capabilities can pose difficult issues of supervision and can lead to heightened risks for firms. At the same time, the issues and risks associated with the use of AI are context-specific—in many contexts and for many applications, potential risks are limited or are quite different from those posed by other uses of AI. Stark differences exist between, for example, algorithms that can autonomously assist a customer with trading and algorithms used in a website navigation function or a chatbot. Even for particular AI capabilities, there are many ways firms may deploy the capability and many context-based considerations that inform assessments of risks and benefits. It would be extremely difficult to issue useful guidance or rules that apply across all these varied circumstances and AI functions, and efforts to do so would likely be hopelessly general or affirmatively harmful when applied to incongruous circumstances or capabilities.

The absence of any compelling, current need for regulatory intervention or guidance is especially clear in light of firms’ incentives and capabilities to identify and address risks. As the AI White Paper repeatedly recognizes, firms understand that AI is integral to their core functions and are devoting considerable resources to using AI to improve compliance and risk management functions, including by “establish[ing] centers of excellence to review, share, and build expertise and create synergies related to the use of AI across their organizations” (p. 5). Firms are attentive to the importance of AI and to the financial, reputational, and legal risks posed by the use of AI technologies.

Many companies have or are developing board-level oversight of the policies and processes associated with the adoption and deployment of AI capabilities, and have control functions focused on ensuring that AI is developed with appropriate attention to existing regulations and risks to the firm and customers. Many firms use cross-functional review and monitoring processes to innovate in the development of AI capabilities, remain alert to any unexpected outcomes, and manage associated risks. Because risks associated with AI capabilities are so context-specific, and dependent on the distinct technology employed and the particular processes it facilitates, firms are uniquely capable of assessing and addressing those particular risks meaningfully.

Should regulation or guidance be deemed appropriate and necessary at some point, it will be crucially important to beneficial innovation that FINRA remain conscious of the adverse regulatory consequences described above. For example, these risks are likely to be mitigated by adopting a principles-based perspective on AI. In other words, FINRA should clearly identify the harms to be abated and let firms be responsible to find relevant solutions. Such an approach would focus on the risks and outcomes associated with securities business processes, rather than prescribing particular standards, mechanisms,

or processes. This technology-neutral and non-prescriptive approach would enable FINRA to remain nimble in its oversight of the risks associated with AI, while also allowing firms the flexibility to continue exploring beneficial uses of this emerging technology.

II. Specific Comments on the AI White Paper

A. Developing a common understanding of AI is critical to a productive working relationship among regulators, industry participants, investors, other consumers, and others.

The AI White Paper continues an important process of developing a widely shared and common understanding of AI. This process would ideally result in a common lexicon for discussing AI systems and their deployment in the securities industry and is critical for fostering a productive working relationship between FINRA and securities firms as they address challenges related to AI. Such an understanding can be achieved by synthesizing existing terms used to discuss AI by member firms and promoting harmonization between the various industry and regulatory bodies currently addressing AI with relevance to the securities and financial sectors.

Section I of the AI White Paper provides a useful outline of potentially relevant issues and definitions that continue this process of helping industry participants develop a common lexicon for discussing AI systems. This section begins (p. 2) by offering three broad definitions of artificial intelligence:

- “The term artificial intelligence broadly refers to applications of technology to perform tasks that resemble human cognitive function and is generally defined as the capability of a machine to imitate intelligent human behavior.”
- “AI typically involves the theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages.”
- “John McCarthy, one of the founders of AI research, once defined the field as getting a computer to do things which, when done by people, are said to involve intelligence.”

An important further step in developing a common understanding of AI would be to more narrowly define these broader concepts. For example, in its response to the European Commission’s White Paper on Artificial Intelligence,³ the Association for Financial Markets in Europe (“AFME”) defined artificial intelligence systems as follows:

Artificial intelligence (AI) systems are systems that act in the physical or digital world by perceiving their environment through data acquisition, interpreting the collected data, reasoning on the knowledge or processing the information derived from this data, and identifying the best action(s) to take to achieve the given goal. AI systems adapt

³ WHITE PAPER On Artificial Intelligence - A European approach to excellence and trust, European Commission (Feb. 19, 2020), https://ec.europa.eu/info/sites/info/files/commission-white-paper-artificial-intelligence-feb2020_en.pdf.

*themselves or their own algorithms by analyzing how the environment is affected by previous actions, knowledge, or data.*⁴

While AFME's particular definition of AI need not be adopted, having a similarly refined definition of the "umbrella" concept—one that focuses on the characteristics of artificial intelligence—would be helpful for fostering constructive dialogue among FINRA, firms, and other market participants.⁵

In addition, this process is ongoing and perspectives are continuing to evolve. Further work might include, for example, efforts by FINRA to develop a generally applicable understanding of AI with various domestic regulators—like the Office of the Comptroller of the Currency (OCC) and the Federal Reserve—and even potentially seeking to develop a common understanding that extends beyond regulators addressing the financial sector.

Agency harmonization might especially be advanced through FINRA's continued effort to draw out and find consensus regarding the broad principles and definitional assumptions that already underpin various regulatory frameworks. One way to develop a more meaningful, common understanding of AI is to draw on these already existing sets of principles and the sources of information. For example, a working group comprised of the various agencies already involved with AI supervision might promote harmonization among often overlapping policy regimes, and FINRA's efforts to build on its AI White Paper could provide a further, valuable contribution to that process.

International regulators, such as the UK's Financial Conduct Authority, the UK's Information Commissioner's Office, and the Monetary Authority of Singapore, would also likely have valuable lessons to share from the early implementation of their policy frameworks. This harmonization effort, building upon a continuation of the constructive dialogue between FINRA and financial services firms, could help in promoting more meaningful measures to reduce risks associated with AI technologies and facilitate potentially productive AI deployments, while also protecting and increasing the benefits that AI technologies provide to firms, investors, and other consumers.

B. FINRA might usefully clarify how existing laws—particularly recordkeeping regulations and general supervisory rules—apply to AI processes.

The AI White Paper very appropriately notes that "the use of AI tools...may create novel issues in the context of compliance with applicable recordkeeping requirements" (p. 19). The varied nature of AI processes deployed by financial industry firms and the varied outputs and risk levels associated with those processes make traditional recordkeeping requirements extremely difficult to apply in the context of these new technologies. For example, an overly rigid application of regulations like FINRA Rule 4510 could cause counterproductive uncertainty as well as practical compliance difficulties. This would impede the development and adoption of AI technologies while providing no additional protection to investors. Accordingly, FINRA might usefully work with firms to develop a workable and practical way to adapt recordkeeping and supervision obligations to apply if and as appropriate to the use of AI models, particularly in the context of AI-to-AI communications and human-to-AI communications.

⁴ *Consultation Response, Artificial Intelligence – A European Approach to Excellence and Trust*, Association for Financial Markets in Europe (June 12, 2020), https://www.afme.eu/Portals/0/DispatchFeaturedImages/20200612%20AFME%20EC%20AI%20CP%20Response%20-%20Final_.pdf.

⁵ With respect to the AI White Paper's definition of Unsupervised Machine Learning (p. 3), it may not always be true—depending on the use case—that humans verify or review the outputs of an AI system once the AI is in production. Accordingly, this definition might be updated as follows: "*Results of unsupervised machine learning models may or may not be interpreted by humans to determine if they are meaningful and relevant, depending on the use case.*"

The AI White Paper also discusses (pp. 12, 14) the application of rules that impose more general requirements on firms, such as FINRA Rule 2010 (firms must “in the conduct of their business... observe high standards of commercial honor and just and equitable principles of trade”) and FINRA Rule 3110 (firms must “establish and maintain a system to supervise the activities of its associated persons that is reasonably designed to achieve compliance with the applicable securities laws and regulations and FINRA rules”). As with existing recordkeeping rules, FINRA might usefully engage with firms to delineate the appropriate scope of rules like FINRA Rule 2010 and FINRA Rule 3110. This dialogue would facilitate reasonable supervisory efforts by FINRA that avoid imposing an overwhelming compliance burden on firms.

C. The need for explainability of an AI algorithm should depend on the features of that algorithm and the risks associated with its use.

Section III of the AI White Paper usefully highlights the importance of “an appropriate level of explainability” in AI functions, particularly in those functions that have “autonomous decision-making features” (p. 12). Explainability is indeed an important consideration for firms creating AI models. At the same time, the complexity and materiality of AI models often vary significantly based on each model’s parameters and sensitivities and each model’s business use. These variations may affect a firm’s ability to explain models that are especially complex.⁶

In addition, depending on the uses of a particular AI process, granular explainability may not always be relevant or worthwhile. The necessity of explaining an AI process depends on the risks associated with that process: as the risk increases, so does the need for explainability, and vice versa. Consider, for example, an algorithm that runs a chatbot that helps customers navigate a financial institution’s website. In this common scenario, there is almost no material risk associated with the chatbot making a mistake. As a result, the need for a detailed explanation of the algorithm is lessened. On the other hand, the deployment of a chatbot used as part of processes that can autonomously process an investment order or advise a client on investment options presents a higher level of risk.⁷ Being able to explain the basis for and operation of this latter algorithm is much more important. Given the wide variety of use cases and model types, and the varying materiality of associated risks, any potential guidance on AI explainability should allow flexibility for firms to tailor and adapt their explanations as appropriate. In other words, firms should be permitted to consider the appropriate level of and need for explainability in the context of the specific model at issue.

One focus of the ongoing dialogue between member firms and FINRA might usefully involve refinement of the circumstances that call for particular levels of explainability of AI processes. In order for FINRA to develop useful guidance, or for firms to assimilate that guidance, there should be clear definitions and identification of relevant risks. In particular, to avoid ineffective, onerous, or potentially unnecessary guidance or regulation, FINRA would benefit from understanding the models being used by firms and how those models impact the issues the agency seeks to address. Perhaps most important, FINRA should not

⁶ It is important to note that some complex AI models may pose little if any practical risk to investors or others (*i.e.*, a model may be particularly complex, but that does not necessarily make it particularly risky).

⁷ Firms recognize the risks associated with the limitations of human review and often seek to mitigate them by deploying certain controls or guardrails (*e.g.*, “kill switches” that trigger human review, caps, floors, alert systems, or performance testing and monitoring). As FINRA explores the risks associated with the deployment of AI processes in the securities industry and the need for explainability, it might usefully consider the varied nature of these processes and the differing levels of human involvement associated with their use. A reasonable practice might be for a firm itself to determine what degree of human oversight or other controls are necessary or appropriate for each AI application. Such judgments could be reassessed periodically based on actual experience and other relevant developments.

assume the existence of substantial risk simply because of the use of an AI-based model, but rather be rigorous and realistic in characterizing the harms that are important to prevent.

D. FINRA should take an output-driven, individualized approach to supervising firms' data governance.

The AI White Paper emphasizes that “the quality of the underlying dataset is of paramount importance in any AI application” and discusses the need for firms to maintain the integrity of their data used in this context (p. 13). This process also involves ensuring that the relevant underlying database is “sufficiently large, valid, and current” and considering certain data-related risks in their AI applications, including: data source verification (especially of externally-sourced or open-sourced data); effective data integration into their systems; appropriate entitlement, authentication, and access control procedures to ensure data security; and development of data quality benchmarks and metrics (pp. 14-15).

At this stage, concerns regarding data integrity are widely shared among firms, which are approaching the issue in different ways tailored to the particular context and purpose of the specific AI models at issue. Any broad regulations or supervision based on FINRA’s concerns about data integrity would be difficult to implement effectively. Firms in the securities industry are not identically situated, and they do not deal with identical constituents. The function of each AI model and the context in which it is operating are also important and varied, as the risks associated with different AI models often differ greatly. For these reasons, firms typically face different risks to the integrity of the data underlying their deployments of AI, which are not suited to a “one-size-fits-all” approach to supervising firms’ data governance.

The AI White Paper also encourages firms to review their data for potential bias. FINRA suggests that firms engage in testing, training for involved individuals, and using open-source tools to identify unwanted data biases. Limiting the effect of bias is indeed extremely important when deploying AI technologies. Many firms analyze, monitor, test, and re-evaluate data inputs to identify and correct bias-related errors and to determine whether any flaws are systematic or ad hoc. Even so, data sources and inputs are always potentially imperfect as a result of inaccuracy, statistical anomalies, and collection errors and it would be overly burdensome to require firms to test data integrity for every application regardless of the materiality and risk of the application. Consistent with the approach advanced in Section I, FINRA might productively focus on and explore further how to address unfairly biased or discriminatory outcomes, rather than supervising the myriad issues that arise with data inputs.

Finally, the AI White Paper’s Model Risk Management section (p. 11) refers to “data integrity.” However, that function may not fall within a firm’s model risk management function but may be covered by other controls or functions the firm has in place (e.g., data governance functions). As a result, further development of FINRA’s treatment of this issue might reflect different firms’ varied approaches to these risks.

E. Firms’ policies and procedures regarding customer privacy are unlikely to require additional guidance or regulation.

The AI White Paper discusses certain rules and regulations relating to customer privacy that apply to firms operating in the securities industry. These include but are not limited to: SEC Regulation S-P, NASD Notice to Members 05-49, and SEC Regulation S-ID (pp. 15-16). This discussion suggests that firms may need to update their policies and procedures regarding customer data privacy as they deploy more AI applications. However, as a general matter, member firms understand the importance of protecting customer data; are subject to various internal and external policies and requirements; often have data and customer privacy functions dedicated to updating and revising their policies to adapt to changing circumstances; and endeavor to comply with all applicable guidance from state and federal regulators. As

part of this effort, firms regularly review and update their policies and procedures relating to customer data, including data associated with AI technology and applications.

Because of this existing regulatory framework, and firms' significant, corresponding compliance capabilities, promulgating additional guidance relating to the way that firms handle personal information is likely to be unnecessary or duplicative. Any guidance that FINRA might still choose to issue should reflect the fact that not all firms are similarly situated in terms of the nature of or risks to the information they collect. Accordingly, any guidance would preferably provide individual firms with the flexibility to develop or adopt their own frameworks and processes for handling and protecting personal information relevant to their AI applications. This approach to guidance will allow firms to continue their compliance with existing rules like Regulations S-P and S-ID without undue burden and additional complexity.

As with the AI White Paper's treatment of "data integrity," noted above in Subsection D, the AI White Paper refers in the Model Risk Management section (p. 11) to "customer privacy," which may not fall within a firm's particular model risk management function but may be covered by other controls or functions the firm has in place (e.g., privacy, legal and compliance functions). Further work in this area might usefully take these variances among firms' practices into account.

F. FINRA should focus the application of its licensing and registration rules on business processes rather than specific AI technologies.

The AI White Paper discusses the evolving skillsets of securities industry personnel and suggests that "[f]irms may need to evaluate the roles of [technical and operational personnel] to ensure that they have the appropriate FINRA licenses and registration" (p. 17). FINRA notes, for example, that "Regulatory Notice 16-21...requires registration of associated persons involved in the design, development, or significant modification of algorithmic trading strategies." SIFMA agrees that, as with other technical and operational roles, firms should always evaluate the need for the registration of personnel involved in the development and implementation of AI technology.

III. Suggested Next Steps.

The AI White Paper is an important step in the ongoing process of developing a common industry and regulatory understanding of the opportunities and challenges posed by the deployment of AI technologies in the securities industry. FINRA should be commended both for taking this step and for the way in which it is actively engaging firms in its exploratory process.

It is important to continue to foster open, ongoing, and transparent communication between FINRA and member firms. Although the Coronavirus Pandemic complicates traditional methods of meeting (e.g., industry-sponsored conferences), FINRA and member firms can continue this productive dialogue in a variety of ways. For example:

- Once FINRA synthesizes the comments it receives from member firms and updates the AI White Paper accordingly, it might choose to solicit another round of comments from member firms and other industry participants.
- Member firms can sponsor virtual conferences that feature speakers from government regulators, including FINRA, that are actively involved in the securities industry and with AI.
- FINRA can work with one or more industry bodies (e.g., SIFMA itself, The Institute of International Finance ("IIF"), or others) to conduct deeper studies and possibly identify and develop industry best practices. IIF, for example, has already done work in this area.

- FINRA and other regulators could host virtual workshops for member firms, where they present discrete topics relating to AI and invite discussion.

SIFMA would welcome the opportunity to be able to continue to contribute to and participate in this valuable process.

* * *

SIFMA greatly appreciates FINRA's consideration of these comments and would be pleased to discuss any of these views in greater detail if that would assist FINRA's deliberations. Please feel free to contact me at (202) 962-7385 if you would like to discuss further.

Sincerely,

Melissa MacGregor
Managing Director and Associate General Counsel

cc: Haimera Workie, Senior Director and Head of Office of Financial Innovation, FINRA
Alan Charles Raul, Richard Klingler, and Christopher Joyce, Sidley Austin LLP