



22 West Washington Street
Chicago
Illinois 60602

Telephone: +1 312 696-6000
Facsimile: +1 312 696-6001

Submitted Electronically

Aug. 12, 2024

US Department of the Treasury

RE: Request for Information on Uses, Opportunities, and Risks of Artificial Intelligence in the Financial Services Sector

Ladies and Gentlemen:

Morningstar welcomes the opportunity to comment on the “Request for Information on Uses, Opportunities, and Risks of Artificial Intelligence (AI) in the Financial Services Sector,” recently published by the Department of the Treasury.¹ Morningstar is a leading provider of independent investment research and has a long history of advocating for transparency in global markets. Morningstar brings several perspectives on the uses of artificial intelligence in the financial services sector. Morningstar’s mission is to empower investors to reach their financial goals, and our comments reflect this.

This letter contains: 1) a summary of our views and 2) detailed answers to selected questions posed in the Request for Information (RFI), attached as Appendix A.

Executive Summary:

To assist the Treasury in achieving its objectives, we offer the following comments and recommendations:

1. The definition of AI in the Request for Information is overly broad and inclusive; its breadth will overregulate long-established and well-regulated technology such as deterministic models, and any new regulation should instead focus on generative models and large language models.
2. Deterministic models have many benefits in the financial services industry. Deterministic tools have provided low-cost investment advice and investment allocation technology to millions of retail investors at a low cost. They also serve regulators in their detection of fraud and other misconduct.
3. AI has benefits for compliance and risk monitoring, allowing personnel to better support various regulatory obligations.
4. To properly regulate AI, regulators should take a principles-based approach, making AI regulation commensurate with its risks. As such, deterministic algorithms present far less risk than generative AI in the forms of large language models or chatbots taking input that is unconstrained and with unpredictable outputs. Deterministic algorithms are generally well-regulated already in the financial services industry.
5. Regulation should also account for consistency and uniformity across jurisdictions to be cost-efficient and effective.

I. Definition of AI Is Overly Broad

¹ Treasury. 2024. Request for Information on Uses, Opportunities, and Risks of Artificial Intelligence in the Financial Sector. <https://www.federalregister.gov/documents/2024/06/12/2024-12336/request-for-information-on-uses-opportunities-and-risks-of-artificial-intelligence-in-the-financial> (Request for Information).

Morningstar believes that the definition of AI used in this RFI is too broad. It incorporates both deterministic models and generative AI when it should only focus on generative AI. In this comment letter, we will distinguish between generative AI, which presents some risks that should be managed through some new regulations, and deterministic models, which we think are already well-regulated and do not present the same risks. Overregulating deterministic models would result in burdensome testing, stifling innovation, and unintentionally restricting deterministic technologies that have long served investors.

The RFI defines AI as:

“A machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human-based inputs to perceive real and virtual environments; abstract such perceptions into models through analysis in an automated manner; and use model inference to formulate options for information or action.”

This definition makes no distinction between deterministic models or generative AI. When a model is deterministic, its responses are determined beforehand, whether they be specific or encompass a range, and users are able to anticipate outputs. Such models are typically lower risk than generative AI models as outputs are deterministic and should follow logically from inputs. Some examples include basic Excel algorithms, autocorrect, and simple research algorithms. Generative AI models, on the other hand, are more complex and less predictable; they are heavily reliant on the data upon which the model is trained, as well as any alignment processes incorporated as part of the model training. When fed biased data, the generative model could have biased outputs, and it is far less apparent when issues occur with outputs because users do not know what to expect nor where the data on which the model is trained comes from.

By blanketing both higher-risk and lower-risk technologies under the same definition, two problems emerge. First, this broad definition could stifle existing deterministic technologies that have long served investors. Basic Excel or simple smartphone and computer functions such as autocorrect could fall under this definition as those can be classified as “machine-based systems that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments.” Second, the lack of distinction could hinder the future progress of deterministic model advancement and increase compliance costs. To fully understand the problems a broad definition presents, it is important to first understand the benefits of AI in financial services.

A definition of AI should center on generative AI models, such as large language models, where output is not a predictable result of inputs. Such models present fundamentally different risks than deterministic ones.

Morningstar encourages regulators to take other jurisdictions into account when creating definitions or regulations. The recent EU Artificial Intelligence Act defines AI in a more restricted manner than the definition presented in the RFI. The EU AI Act defines AI as:

“ ‘AI system’ means a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.”²

² EU. 2024. Artificial Intelligence Act. Article 3.

<https://artificialintelligenceact.eu/article/3/#:~:text=%E2%80%98AI%20system%E2%80%99%20means,or%20virtual%20environments%3B>

Specifically, this definition highlights that AI systems are models that have some degree of unpredictability by recognizing that they have varying levels of autonomy, acknowledging the distinction between deterministic and generative models. The EU AI Act definition considers what constitutes an AI model more than the RFI's definition.

A definition that is more in line with the EU AI Act would not only free low-risk, deterministic models from overregulation but would also allow cross-jurisdictional financial institutions to take a more consistent approach to compliance between regions.

II. AI Has Many Benefits in the Financial Industry

Deterministic models and generative AI both have many benefits, from providing low-cost financial services to individuals to improving accessibility and reducing operational and compliance costs. Models can increase productivity by as much as 14% on average and even more in some instances, measured by the number of issues resolved per hour, with even greater gains for less experienced workers.³ Models can also upskill junior workers, improving their efficiency and quality of work. An MIT study found that access to AI enabled less-experienced workers to produce higher-quality work faster.⁴ Further, in a survey of marketing and compliance leaders in financial institutions, 85% of respondents said the implementation of AI would save them money.⁵ Compliance functions can do much more with what they are allocated if given access to AI. Both generative and deterministic models can also increase investor accessibility for investors and employees of financial institutions through text-to-speech and speech-to-text software for those who cannot see or type. Some text-to-speech programs have started incorporating AI that can analyze and describe images.

Deterministic models can also reduce human errors and be used for anomaly detection, allowing financial institutions and investors to save time in data analysis. Handling large amounts of data is a difficult and complex task for people to do on their own, but through fine-tuned algorithms, deterministic models can sort through, categorize, and analyze that data in much less time with far more accuracy. Further, deterministic models can operate at all times, allowing people to get important information when they need it. With deterministic models, people can work through more data, save time, and, by extension, streamline processes, leading to a decrease in labor costs. Financial institutions can accomplish a lot more, faster, with a lot less.

Financial institutions use deterministic models such as managed accounts and robo-advisors as low-cost tools to help investors manage their finances. Robo-advisors have a much cheaper barrier to entry with low-cost options and low AUM requirements relative to higher fees and minimums with traditional financial advisors.⁶ Robo-advisors greatly increase access to investing to lower-income groups; a study found that 80% of robo-advisor users have a lower income and net worth than those who invest with traditional financial advisors.⁷

³ NBER. 2023. "Generative AI at Work." <https://www.nber.org/papers/w31161>.

⁴ MIT. 2023. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." https://economics.mit.edu/sites/default/files/inline-files/Noy_Zhang_1.pdf.

⁵ Saifr. 2024. "AI insights survey: Adopters, skeptics, and why it matters." <https://insights.saifr.ai/request-ai-insights-survey-cw>.

⁶ Forbes. 2023. Robo-Advisor vs. Financial Advisor: What's The Difference? <https://www.forbes.com/advisor/investing/financial-advisor/robo-advisor-vs-financial-advisor/>

⁷ University of Georgia. 2018. Financial Services Review: Who uses robo-advisory services, and who does not? https://www.researchgate.net/publication/375682484_Who_uses_robo-advisory_services_and_who_does_not

Morningstar's own ratings system uses deterministic algorithms to simultaneously expand and streamline research and access to site resources, utilizing both internal and third-party data. Much of that data would be difficult for individual people to manage, but with algorithms, we can streamline our own internal processes and provide useful information to investors more efficiently. Morningstar Investment Management also provides plan sponsors with managed accounts for retirement plans. These are deterministic models that help investors make investment allocation decisions. Despite their deterministic nature, both the ratings and managed accounts algorithms would fall under the RFI's definition of AI, even though they are both low-risk, have served investors for a long time, and are sufficiently regulated under current laws and regulations. These types of models should not be classified as AI in the same way that higher-risk, generative models are.

Additionally, deterministic models can decrease costs by reducing compliance burdens through AI compliance models trained to flag noncompliant content in a financial institution's materials. Compliance departments are actively utilizing tools falling under the Treasury definition of AI for monitoring personal trading activities by employees. By assisting in the compliance process, deterministic compliance models support more orderly and efficient compliance programs, minimizing menial tasks and allowing compliance personnel to focus their efforts on items requiring additional attention, such as complex scenarios or problem solving.

III. Consequences of a Broad Definition and Blunt Regulations

Taking the benefits into account, it becomes clear that if a broad definition is employed in AI regulation, then regulators will create "blunt" regulations that affect all models evenly regardless of risk level. Financial institutions that rely on these technologies will face higher compliance burdens than they do already. For smaller institutions with fewer resources, increased compliance burdens could deter them from using certain models altogether, thereby weakening their capacity and disproportionately harming them in comparison to larger institutions.

We believe blunt regulations would stifle the benefits of low-risk, deterministic models. Deterministic models serve both financial institutions and investors by improving efficiency, reducing costs of labor, and providing objective information such as quantitative investment ratings to investors. These low-risk models, by definition, do not pose much risk, so it is counterproductive to restrict them to the same degree as high-risk, generative models.

If deterministic models become less accessible, investors will be negatively affected. Many of these models simplify tasks that would normally be too cumbersome for individuals to do, such as analyzing large swaths of investment data and rating funds. Investors rely on deterministic models, such as robo-advisors, to support their financial literacy and conduct data analysis. Such algorithmic tools can increase services for investors and help them plan for their future. Broad definitions and blunt regulations should not be the path forward.

There are many ways of addressing AI risks. Strategies include but are not limited to:

- Enhanced disclosure requirements to make sure users are aware of the risks of generative models.
- More robust oversight mechanisms such as oversight committees, internal generative AI use policies, and human-in-the-loop models.
- Disclaimers on chatbots that ensure users are aware the bot is not intended to give advice on particular topics and warn users not to provide personal information.

IV. Risk-Based Approach

When it comes to AI regulation, Morningstar suggests that regulators should take a risk-based approach to AI, as not all models pose the same level of risk, and to not classify deterministic models as AI. By categorizing the threat of each model and applying stricter regulations only to those models that pose the greatest risk for harm, low-risk, deterministic models can remain widely available. Definitions and regulations that address AI models commensurate to the risk will allow for a more principles-based and effective approach to AI regulation that will allow AI to develop and not increase compliance costs and burdens for financial institutions disproportionately to the risk involved.

Morningstar believes that there are ultimately better outcomes for end users where there is a balance between compliance measures and the promotion of innovation in models employed by financial institutions to the benefit of investor access and advice quality. Specifically, AI benefits a financial institution's clients where innovation is balanced with compliance. An institution can provide smoother service with faster turnaround in a compliant manner using deterministic models. Many existing deterministic models, such as managed accounts and robo-advisors, already have robust compliance policies and processes under current regulatory requirements. Regulations such as Regulation Best Interest account for conflicts of interest concerns through deterministic model usage. Other examples of the use of deterministic models in compliance include fraud detection, marketing compliance, and the detection of money laundering. By focusing regulation on the generative models, regulators can avoid stifling the progress of and benefits provided by low-risk models, including in the area of compliance itself. Adding more testing and compliance is redundant and burdensome for low-risk models. With a principles-based approach to compliance commensurate with the risk, financial institutions will find that they will still be able to serve their clients as they always have.

Conclusion

In summary, we acknowledge that some regulatory action should be taken in regard to high-risk generative AI models. The definition adopted by the RFI, however, is broad and would require cumbersome testing on existing products and services that have been serving investors for a long time, stifling deterministic models. These tools have myriad benefits for financial institutions and investors such as increasing efficiency, reducing labor costs, and improving accessibility to objective information on investments, among many others. Broad definitions, blunt regulations, and inconsistency between jurisdictions can all result in increased compliance and financial burdens for financial institutions and stifle the benefits and growth, particularly of deterministic models. Furthermore, Morningstar encourages the Department of the Treasury to carefully consider the concept of making compliance commensurate with the risk posed by a particular investor interaction as a means of achieving its goals without discouraging beneficial services for investors.

We have summarized our views above and answered some specific questions from the RFI in Appendix A. Morningstar thanks the Department of the Treasury for the opportunity to comment on the RFI. We would be pleased to engage with the Treasury on an ongoing basis, leveraging our global organization of experts operating in multiple jurisdictions. Should you wish to discuss these and other comments, please do not hesitate to contact either of us as indicated below:

Jasmin Sethi at jasmin.sethi@morningstar.com or + 1 617 501-5446
Aron Szapiro at aron.szapiro@morningstar.com or + 1 732 322-2464

Sincerely,
Aron Szapiro
Head of Government Affairs

Morningstar, Inc./Morningstar Investment Management, LLC

Jasmin Sethi
Associate Director of Policy Research
Morningstar, Inc.

Appendix A: Answers to Selected Questions

General Use of AI in Financial Services

Question 1: Is the definition of AI used in this RFI appropriate for financial institutions? Should the definition be broader or narrower, given the uses of AI by financial institutions in different contexts? To the extent possible, please provide specific suggestions on the definitions of AI used in this RFI.

- AI is defined as “a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human-based inputs to perceive real and virtual environments; abstract such perceptions into models through analysis in an automated manner; and use model inference to formulate options for information or action.”

Answer: The definition used in this RFI is inappropriate for financial institutions because it is too broad and fails to account for advancements in AI. By broadly defining AI, this RFI will likely have unintended consequences, such as drastically increasing the compliance burdens on businesses using technology covered by this definition. Defining AI as machine-based systems that make recommendations to impact real or virtual environments is too vague because the scope of that definition covers many routine uses of technology that are not necessarily AI. The definition presented by the Treasury is not forward-thinking, particularly regarding what AI is now and what it will become. Deterministic algorithms such as autocorrect and spellcheck are not considered AI, yet they would fall under this definition. Additionally, basic deterministic algorithms in programs like Excel or even simple functions of devices like smartphones and computers often make recommendations, predictions, or decisions, yet are widely not considered AI. As written, this definition presents unintended consequences for low-risk technologies that serve both investors and financial institutions.

A definition of AI should focus on dynamic generative models, such as large language models, where output is not a predictable result of inputs. Such models present fundamentally different risks than deterministic ones. The definition could stifle deterministic technology that has been assisting investors for years. Deterministic technology is already subject to sufficiently appropriate regulation.

Question 2: What types of AI models and tools are financial institutions using? To what extent and how do financial institutions expect to use AI in the provision of products and services, risk management, capital markets, internal operations, customer services, regulatory compliance, and marketing?

Answer: Financial institutions provide solutions, including managed accounts, informational tools, and robo-advisors, to help investors manage their finances. These existing deterministic technologies already have robust compliance policies and procedures under the current regulatory regime to protect investors. These tools help lighten the burden on financial

institutions and streamline their work. They are also often low-cost tools that provide a cheaper, more accessible alternative to a traditional financial advisor. Deterministic models can streamline the work process in a few ways, including usage for internal research, and they can also improve investment analysis capabilities, which allow financial institutions to provide clients with more information about their investments. The increase in efficiency that these tools provide allows financial institutions to increase their capacity, take on more clients, and have a larger impact, all while helping a broad base of investors accomplish their goals.

Further, institutions may use deterministic or generative models in editorial production to streamline editing for public communications. Financial institutions can also use deterministic models for basic personal finance tracking services to help investors keep track of their expenses and investments through algorithms. Additional use cases include Retrieval Augmented Generation to provide chatbot services or summarize information across large document libraries or datasets, workflow automation and enablement (that is, chatbots that take actions using software or systems on the users' behalf or predictive actions/suggestions based on usage patterns and tendencies), as well as data extraction and processing from unstructured sources.

For Morningstar in particular, deterministic algorithms make it possible for investors to use Morningstar tools and analytics to research securities, assess their risk, and monitor their portfolios. When managing securities, there are too many variables and calculations for humans to do without the use of deterministic algorithms, some of which fall in scope of the proposed AI definition. Most of these forms, however, are not generative and, therefore, provide predictable, explainable outcomes.

Question 3: To what extent does the type of AI, the development of AI, or AI applied use cases differ within a financial institution? Please describe the various types of AI and their applied use cases within a financial institution.

Answer: There are several main types of AI use cases within an institution. Internal tools are meant to support the work of employees, and customer-facing tools are intended to assist customers with their needs. Third-party vendors may also facilitate the use of tools. Within these categories, of course, all kinds of AI tools exist – including ones that take input from predetermined research from users, customers, the internet, third-party sources, and so forth. The inputs for the generative AI model are extremely important, as AI accuracy bias depends on the data given and the efforts of the humans who create it.

Internal tools could include generative AI models to help with editing or research or simpler research tools trained to output predictable responses, among others. Internal tools may also further compliance and risk management.

Customer-facing tools could include deterministic informational tools to improve financial literacy and determine and manage investment allocation effectively and cheaply. For instance, managed accounts, long used in defined-contribution plans, allow retirement plan

providers to avoid conflicts of interest and other potential compliance risks and provide investment allocation support to their employees. The algorithms underpinning managed account services are deterministic, predictable, and subject to existing regulatory scrutiny.

Question 3b: Are there additional use cases for which financial institutions are applying AI or for which financial institutions are exploring the use of AI? Are there any related reputation risk concerns about using AI? If so, please provide specific examples.

Answer: Some financial institutions are also developing models that can help identify compliance risks within an organization, such as noncompliant marketing language. Such models can ease compliance burdens and help smaller financial institutions keep up by reducing compliance costs. Morningstar's Mo chatbot can also provide quick and easy answers drawn from Morningstar's extensive editorial, research, and data content, such as what the Morningstar Ratings are or how to think about saving for retirement, greatly streamlining research and access to resources on the site. This kind of consumer-facing model can greatly assist investors in their own research.

Financial institutions may face reputational risks if generative AI tools provide responses that are inaccurate or do not comply with applicable legal requirements. These AI tools, however, can be a tremendous resource for providing information to users in an easy and accessible way, allowing everyone, including seniors and people with disabilities, to obtain information that can benefit them easily and in a low-cost manner.

Question 4: Are there challenges or barriers to access for small financial institutions seeking to use AI? If so, why are these barriers present? Do these barriers introduce risks for small financial institutions? If so, how do financial institutions expect to mitigate those risks?

Answer: AI access is determined by how much data is available to a model. Generic and publicly available models may be available for free to institutions and consumers. If institutions are not able to make their own models, they are at the mercy of the costs set by third-party providers, though currently, commercial large language model pricing is highly competitive and affordable. Stricter regulation could affect pricing, but third-party models are a valid option in the present market. There are ways to adjust to account for high costs, but building their own AI could be cost-prohibitive for smaller financial institutions. One of the biggest challenges small financial institutions face when seeking to use AI is compliance burdens. Strict compliance burdens lead to an increase in excessive costs and administrative burdens. The cost of complying, in some cases, may be too great for small financial institutions to justify, thereby leading to decreased access to AI.

Actual and Potential Opportunities and Benefits

Question 5: What are the actual and expected benefits from the use of AI to any of the following stakeholders: financial institutions, financial regulators, consumers, researchers, advocacy groups, or others? Please describe specific benefits with supporting data and

examples. How has the use of AI provided specific benefits to low-to-moderate income consumers and/or underserved individuals and communities (e.g., communities of color, women, rural, tribal, or disadvantaged communities)?

Answer: AI's benefits are numerous, whether the model is deterministic or generative. Studies show that the use of generative AI can increase productivity, help professionals generate content faster in a more efficient manner, and accelerate “upskilling” junior workers.⁸ These efficiency gains can help institutions save money. A survey conducted by Saifr, a startup incubated within Fidelity, found that 85% of marketing and compliance leaders reported that AI would save them both time and money through efficiency gains and automated compliance assistance.⁹

While we believe that the definition of AI should be limited to generative AI, we think it is important to consider examples of deterministic technology that fit the definition in the RFI to better understand the risks and appropriate regulation.

For instance, robo-advisors are low-cost deterministic investment allocation models that can serve as a cheaper alternative to traditional financial advisors, making financial advice more available to lower-cost consumers. Some, like Acorns or Betterment, allow a customer to start with \$10 or less, whereas traditional financial advisors may require as much as \$50,000 or more to get started.¹⁰ Accordingly, a financial services review by the University of Georgia found that those who primarily use robo-advisors have lower incomes and net worths than those who have traditional financial advisors.¹¹

Deterministic models can serve as an accessibility tool for disabled individuals by allowing text-to-speech readers to analyze and describe images, allowing speech-to-text for those who cannot type or see, and captions for those who cannot hear – thereby benefiting both financial professionals and their clients by making information more widely accessible.

Regulators such as the SEC use AI analytics to predict risk and protect investors and consumers. The SEC created the Office of Data Science within the Division of Economic and Risk Analysis partly to utilize technology such as AI and machine learning. DERA staff use AI algorithms to identify possibly problematic language in regulatory filings that could result in enforcement. AI has also aided the SEC in more effectively predicting, managing, and allocating resources for investigations and enforcement actions against investment advisors and corporate issuers.¹² Another example of regulatory AI use is Finra's Rulebook Search

⁸ NBER. 2023. Generative AI at Work. <https://www.nber.org/papers/w31161>

⁹ Saifr. 2024. AI insights survey: Adopters, skeptics, and why it matters. <https://insights.saifr.ai/request-ai-insights-survey#top-of-page>

¹⁰ Forbes. 2023. Robo-Advisor vs. Financial Advisor: What's The Difference? <https://www.forbes.com/advisor/investing/financial-advisor/robo-advisor-vs-financial-advisor/>

¹¹ University of Georgia. 2018. Financial Services Review: Who uses robo-advisory services, and who does not? https://www.researchgate.net/publication/375682484_Who_uses_robo-advisory_services_and_who_does_not

¹² SEC. 2017. The Role of Big Data, Machine Learning, and AI in Assessing Risks: a Regulatory Perspective. <https://www.sec.gov/newsroom/speeches-statements/bauguess-big-data-ai>

Tool, which allows users, including retail investors, regulated entities, market participants, and computer systems, to search rules in an automated manner.¹³ Thus, AI – both generative and the deterministic models captured under the RFI’s definition – has myriad benefits ranging from financial and professional applications to accessibility and social applications.

Actual and Potential Risks and Risk Management

Question 6: To what extent are the AI models and tools used by financial institutions developed in-house, by third-parties, or based on open-source code? What are the benefits and risks of using AI models and tools developed in-house, by third-parties, or based on open-source code?

Answer: It is cheaper and easier for financial institutions to rely on a third-party generative AI model to build off of rather than construct their own from scratch, as that requires technical expertise and a lot of resources. However, there are both advantages and disadvantages to doing so. On the one hand, third-party models decrease the need for technical expertise in-house and free an institution from the burdens of maintaining a model to prevent model drift, which is the degradation of a model over time due to changes in data affecting input and output properties. On the other hand, depending on the model, there could be privacy or confidentiality concerns when entrusting internal data to an external tool. In-house AI models do not have the same concerns but come with increased costs and administrative/technical burdens on a financial institution. Further, third parties or models using open-source material sometimes lack the industry-specific or market-idiosyncratic contexts necessary to produce accurate results. Challenges exist around who is responsible for vetting open-source code in a variety of scenarios – the publisher may have tested in one set of scenarios, whereas the user may have a completely different use case in mind.

In-house models allow an institution complete control over how the model functions, which means an institution can tailor the model to its particular needs or the needs of its clients. Additionally, one large benefit to in-house models for financial institutions is that they can be tailored to help their clients access and operate the institution’s own financial tools by automating existing features and making them easier to use. This automation, in turn, makes financial technology more accessible to the average consumer and can benefit their financial planning in the long run.

Risks are also greatly influenced by the dataset and alignment processes used. If a dataset or alignment process has flaws and biases, then the resulting model and its outputs may be flawed and biased. It is important that AI oversight committees ensure that data is properly vetted before implementation into a model. Biases often come from human errors when data is not thoroughly checked for inaccuracies and incompleteness. Further, it is important that AI oversight committees are diverse so that diverse representation in data selection and oversight is accounted for to reduce biases.

¹³ Finra. 2024. Finra Rulebook Search Tool™ (FIRST™) Overview. <https://www.finra.org/rules-guidance/rulebooks/finra-rulebook-search-tool-first>

Question 6b: To what extent are a particular financial institution's AI models and tools connected to other financial institutions' models and tools? What are the benefits and risks to financial institutions and consumers when the AI models and tools are interconnected among financial institutions?

Answer: Interconnectedness depends on the model. Some models are developed entirely in-house from scratch, and other models use an external open-source model as a base and then build off of it. It may be the case that several models share the same core base, but they build off it in different ways, so they are unique. One model building on another brings in the risks from the underlying model. Contract terms around use cases can mitigate misuse of the model to manage risk but may not be able to anticipate all situations. Testing is an important way to mitigate risks introduced by layering models on top of each other.

Open-source models also come with a variety of privacy risks. Any information put into an open-source model may no longer be private as it is possible for others to retrieve that information. Open-source models are not typically used for financial advice as they would likely give bad financial advice because they include too much general information in their training unrelated to finance. The best solution is for financial institutions to develop their own individual AI code of ethics for their employees to ensure that open-source models are used responsibly and effectively.

Question 7: How do financial institutions expect to apply risk management or other frameworks and guidance to the use of AI, and in particular, emerging AI technologies? Please describe the governance structure and risk management frameworks financial institutions expect to apply in connection with the development and deployment of AI. Please provide examples of policies and/or practices, to the extent applicable.

Answer: Financial institutions can and are creating AI governance committees that oversee the usage, development, and ongoing maintenance of AI within the organization. These committees can enforce internal AI policies and ensure that AI usage and output are on par with their own standards. Through an AI code of ethics, they can protect both employees and customers. For example, user directions provided alongside any generative AI system can help customers not only understand what not to input in order to keep their information safe but also help them learn to use the model in a manner that most benefit them. In terms of risk management, financial institutions benefit from implementing internal AI policies specific to generative AI that ensure that employees understand how to use generative AI in a safe and responsible manner without introducing data privacy, intellectual property, ethics, or other risks.

Question 7b: What types of testing methods are financial institutions utilizing in connection with the development and deployment of AI models and tools? Please describe the testing purpose and the specific testing methods utilized to the extent applicable.

Answer: Testing for AI needs to be commensurate with the risk of the model. Higher-risk models should be subject to more rigorous testing.

Question 7c: To what extent are financial institutions evaluating and addressing potential gaps in human capital to ensure that staff can effectively manage the development and validation practices of AI models and tools?

Answer: Monitoring AI requires technical and human oversight. One necessary practice is using human-in-the-loop models. Human-in-the-loop models require human oversight, mitigating the drift of AI models. A human-in-the-loop model does not require a large number of people in order to function, so it is not demanding on human capital. With this check in place, technical oversight through the use of advanced algorithms designed to monitor compliance risks in AI can help keep deployment of AI in check. AI oversight committees may require more human capital, but they are also effective in managing the development and validation practices of AI models. If the members of an oversight committee are diverse, then they can better mitigate biases.

Question 7d: What challenges exist for addressing risks related to AI explainability? What methodologies are being deployed to enhance explainability and protect against potential bias risk?

Answer: Model transparency is vital for AI explainability. If all parts of a model's decision-making process are known, as well as the data it is trained in, then it becomes much simpler to explain the model. Any financial institution should be wary of third-party models with little transparency, and model transparency should be a best practice for both internal and external models. If an institution does not fully understand the processes of its own model, then it cannot reliably trust model outputs. Modern LLM and generative AI models present unique risks with respect to explainability because of their nondeterministic nature. Generated responses can vary meaningfully for a static set of inputs or context, and LLMs can create responses that are not grounded in their training data or provided context. This nondeterministic behavior can potentially produce biased results even on models trained with unbiased data and subject to well-constructed alignment processes.

Question 8: What types of input data are financial institutions using for development of AI models and tools, particularly models and tools relying on emerging AI technologies? Please describe the data governance structure financial institutions expect to apply in confirming the quality and integrity of data. Are financial institutions using "non-traditional" forms of data? If so, what forms of "non-traditional" data are being used? Are financial institutions using alternative forms of data? If so, what forms of alternative data are being used?

Answer: The type of data used affects the risk of the model. Reliable data, such as internal data or well-vetted third-party data, typically presents less risk than AI-generated or non-

traditional data. Creating good AI-generated data requires the AI to create the data accurately so that any model using the AI-generated data is accurate. Raw data, if known to be accurate, provides the more reliable input for AI models. Raw data can be internal, public, or provided by a third party.

Fair Lending, Data Privacy, Fraud, Illicit Finance, and Insurance

Question 11: How are financial institutions addressing any increase in data privacy risk related to the use of AI models, particularly emerging AI technologies? Please provide examples of how financial institutions have assessed data privacy risk in their use of AI.

Answer: Data privacy risks, such as users of generative AI asking questions that include personal information, are a concern; however, models can often be adjusted to reject inputs that contain personal information. Some chatbots, such as Morningstar's Mo chatbot, may reject any questions that contain information that resembles phone numbers, credit card numbers, addresses, Social Security numbers, dates of birth, and so on in order to protect users. Models like Morningstar's Mo will also often display persistent disclaimers that state users should never give the model personal information or that provide useful advice on how to best use the model.

Third-Party Risks

Question 15: To the extent financial institutions are relying on third-parties to develop, deploy, or test the use of AI, and in particular, emerging AI technologies, how do financial institutions expect to manage third-party risks? How are financial institutions applying third-party risk management frameworks to the use of AI?

What challenges exist to mitigating third-party risks related to AI, and in particular, emerging AI technologies, for financial institutions? How have these challenges varied or affected the use of AI across financial institutions of various sizes and complexity?

Answer: It is important that third-party generative AI models are transparent. Transparent training and operational processes vastly improve the explainability and comprehension of generative AI models. Human-in-the-loop models are also effective as they can prevent an AI from drifting and ensure human oversight over the model. As long as a model is explainable and there is proper oversight, a financial institution can more comfortably utilize third-party models while avoiding some of the risks. Risks typically include biased models, models that do not operate correctly or are drifting because of little oversight, and data privacy or confidentiality concerns. Through an explainable model, all these concerns can be alleviated for any financial institution, large or small.

There is some risk posed by the industry trend of third-party generative AI model suppliers – which are typically large technology firms – unilaterally imposing terms regarding model training and data use. A supplier firm can unilaterally decide to train its model with customer data, which would allow the supplier to incorporate customer IP into their products, causing

significant business risks to customers. Regulators should require supplier firms to obtain consent from customers before any change in model training or data use practices.

Question 16: What specific concerns over data confidentiality does the use of third-party AI providers create? What additional enhancements to existing processes do financial institutions expect to make in conducting due diligence prior to using a third-party provider of AI technologies?

Answer: The use of third-party generative AI models does present some concern over data confidentiality. If a financial institution feeds sensitive information to a generative AI model, there is the worry that this information can be retrieved by anyone else with access to the same model. To address this risk, financial institutions can establish generative AI governance to ensure that users do not enter sensitive information to external models. Further, financial institutions can seek contractual commitments with third-party AI providers to ensure that any data fed to the AI will remain private and safe. Transparency between financial institutions and third-party providers is key.

Question 16b: What additional enhancements to existing processes do financial institutions expect to make in monitoring an ongoing third-party relationship, given the advances in AI technologies? How do financial institutions manage supply chain risks related to AI?

Answer: The establishment of AI governance oversight committees can help with monitoring the use of third-party generative AI. With a diverse and experienced committee, a financial institution can prevent supply chain risks by closely monitoring and controlling AI model usage throughout the company, as well as working directly with third-party providers to resolve any issues.

Question 17: How are financial institutions applying operational risk management frameworks to the use of AI? What, if any, emerging risks have not been addressed in financial institutions' existing operational risk management frameworks?

Answer: Chatbots can be a concern if customers are able to use questions to somehow generate investment advice. Morningstar and other financial institutions are able to implement protections on the back-end that prevent investment advice from being given via a chatbot. Nevertheless, on the off chance that someone is able to generate advice, persistent disclaimers that the chatbot is not an advisor and that users should not take advice from it can protect users. Morningstar is able to restrict what kinds of questions users are allowed to ask, and others can be rejected. An additional risk is if an unauthorized user gains access to a model's data repository and uploads bad data, tainting the data pool. It is important that financial institutions keep their systems secure for this reason, though this data security risk is not a new concern.

Question 17b: How are financial institutions ensuring their operations are resilient to disruptions in the integrity, availability, and use of AI? Are financial institutions using AI to

preserve the continuity of other core functions? If so, please provide examples.

Answer: In some instances, algorithms that are not AI may be employed to streamline calculations or access information and data, but AI can be used to supplement those algorithms purely for the purposes of providing context for data. Deterministic algorithms supplemented by contextual generative AI models can work in tandem to provide users with the data and information they need. This is just one way in which AI can preserve the continuity of other core functions within a financial institution.

Further Actions

Question 18: What actions are necessary to promote responsible innovation and competition with respect to the use of AI in financial services? What actions do you recommend Treasury take, and what actions do you recommend others take? What, if any, further actions are needed to protect impacted entities, including consumers, from potential risks and harms?

Answer: Morningstar recommends that financial regulators take a principles-based approach commensurate with the risk when it comes to AI, as not all AI poses the same level of risk. Some low-risk models are already covered by the existing technological regulatory landscape. Regulations such as Regulation Best Interest, for example, already account for conflicts of interest concerns within deterministic model usage. By focusing regulation on the highest-risk models, regulators can avoid stifling the progress of and benefits provided by low-risk, deterministic models. If further regulations are applied to deterministic models, especially with a broad definition of AI, regulators will hinder financial institutions in their ability to help their clients. Further, these regulations could disproportionately impact smaller institutions and low-income customers, who may stand the most to gain from deterministic technologies.

Question 18b: Please provide specific feedback on legislative, regulatory, or supervisory enhancements related to the use of AI that would promote a financial system that delivers inclusive and equitable access to financial services that meet the needs of consumers and businesses while maintaining stability and integrity, protecting critical financial sector infrastructure, and combating illicit finance and national security threats. What enhancements, if any, do you recommend be made to existing governance structures, oversight requirements, or risk management practices as they relate to the use of AI, and in particular, emerging AI technologies?

Answer: To address the risks associated with AI usage in financial institutions, regulations could include risk-based enhanced disclosure requirements, testing, more robust oversight mechanisms, and industrywide best practices. These measures can complement existing rules around the mitigation of conflicts and data privacy and offer a multifaceted approach to safeguarding investor interests. Further, reinforcing existing regulations such as Reg BI, rather than inventing new ones, will go a long way in regulating AI without increasing compliance burdens.

Question 19: To what extent do differences in jurisdictional approaches inside and outside the United States pose concerns for the management of AI-related risks on an enterprise-wide basis? To what extent do such differences have an impact on the development of products, competition, or other commercial matters? To what extent do such differences have an impact on consumer protection or availability?

Answer: US regulations should take an approach that is consistent with other jurisdictions. The EU AI Act's definition, as mentioned in the summary section of this comment letter, takes an approach that defines AI models particularly as unpredictable and autonomous. This makes their definition more specific than the one presented by the RFI. If the US adopted a definition of AI that was more consistent with the EU AI Act, then cross-jurisdictional financial institutions would have fewer compliance concerns to worry about, and it would both allow and encourage institutions to offer the same services both domestically and abroad – benefiting both the institution and their consumers.

Regulation of AI that varies state by state is also not conducive to conducting business nationwide and globally. More consistent and streamlined regulations would be best for effective regulation and allowing AI to continue to benefit consumers.