



State Street Corporation

State Street Financial Center
One Lincoln Street
Boston, MA 02111-2900

July 1, 2021

Office of the Chief Counsel
Office of the Comptroller of the Currency
400 7th Street, SW – Suite 3E-218
Washington, DC 20219

Comment Intake
Bureau of Consumer Financial Protection
1700 G Street, NW
Washington, DC 20552

Federal eRulemaking Portal: Regulations.gov
Docket ID: OCC-2020-0049

E-mail: 2021-RFI-AI@cfpb.gov
Docket Number: CFPB-2021-0004

Ann E. Misback, Secretary
Board of Governors of the Federal Reserve
System
20th Street and Constitution Avenue, NW
Washington, DC 20551

Melane Conyers-Ausbrooks, Secretary of the
Board
National Credit Union Administration
1775 Duke Street
Alexandria, VA 22314

E-mail: regs.comments@federalreserve.gov
Docket Number: OP-1743

Federal eRulemaking Portal: Regulations.gov
Docket Number: NCUA-2021-0023

James P. Sheesley, Assistant Executive Secretary
Federal Deposit Insurance Corporation
550 17th Street, NW
Washington, DC 20429

E-mail: comments@fdic.gov
RIN Number: 3064-ZA24

Re: Request for Information and Comment on Financial Institutions' Use of Artificial Intelligence, Including Machine Learning

Dear Sir/ Madam:

State Street Corporation (“State Street”) welcomes the opportunity to respond to the Request for Information and Comment (“RFI”) issued by the Office of the Comptroller of the Currency, the Board of Governors of the Federal Reserve System, the Federal Deposit Insurance Corporation, the Bureau of Consumer Financial Protection and the National Credit Union

Association (collectively the “Agencies”), on financial institutions’ use of artificial intelligence (“AI”), including machine learning (“ML”) (collectively “AI/ML”), in the provision of services to their clients and in the conduct of their operations. The RFI seeks input on current industry risk management practices for AI/ML and insight on the challenges that financial institutions face in developing, adopting and managing AI/ML solutions.

The Agencies intend to use responses to the RFI to help guide their assessment of whether any clarifications are needed to help promote the appropriate use of AI/ML by the industry in a manner consistent with safety and soundness considerations and existing laws and regulations. We broadly support the Agencies’ effort to explore trends in the use of AI/ML by financial institutions and believe that the existing model risk management framework for prudentially regulated entities, when combined with existing information technology (“IT”) risk management and third-party risk management standards, is sufficiently robust to address potential risk related to the deployment and use of AI/ML.

Headquartered in Boston, Massachusetts, State Street is a global custody bank which specializes in the provision of financial services to institutional investor clients. This includes investment servicing, investment management, data and analytics, and investment research and trading. With \$40.3 trillion in assets under custody and administration and \$3.6 trillion in assets under management, State Street operates in more than 100 geographic markets globally.¹ State Street is organized as a United States (“US”) bank holding company, with operations conducted through several entities, primarily its wholly-owned Massachusetts state-chartered insured depository institution subsidiary, State Street Bank and Trust Company. Our primary prudential regulators are therefore the Massachusetts Division of Banks and the US Federal Reserve System.

The RFI seeks input on a series of questions within several broad categories of focus relative to the use of AI/ML, including: the ‘explainability’ of models; the implications of broader and more intensive data usage and processing; ‘overfitting’; cybersecurity risks; dynamic updating; the use of AI/ML by community financial institutions; oversight of third-party vendors; and fair lending considerations. We appreciate the opportunity to offer insight on the use of AI/ML by prudentially-regulated financial institutions, informed by our role as global custody bank, a role that is widely understood by the market and by the regulatory community as providing important benefits for the safety of client assets and the seamless day-to-day operation of the financial system.

Below we offer: (i) certain high-level observations regarding the use of AI/ML by prudentially-regulated financial institutions; (ii) considerations regarding the proper definition of AI/ML; (iii) an overview of current select State Street use cases; and (iv) responses to the several questions posed within the RFI that are relevant to our business model and its specialized focus on meeting the financial services needs of institutional investors.

¹ As of March 31, 2021.

GENERAL OBSERVATIONS

As emphasized in our introductory comments, we believe that the existing model risk management framework for financial institutions subject to prudential regulation, when combined with existing IT risk management and third-party risk management standards, is sufficiently comprehensive to address potential risks related to the deployment and use of AI/ML solutions. This includes well-established expectations for model identification, model development, model implementation, model validation, and ongoing monitoring and use, which properly accommodate differences in underlying use cases and business models. As such, rather than pursuing potential new regulations for the use of AI/ML by banks, we recommend that the Agencies rely on existing guidance, such as SR 11-7 - Guidance on Model Risk Management, for the oversight of AI/ML usage, providing incremental additional clarification where appropriate.²

In general, we do not face major challenges in the use of AI/ML models in our operations, since existing risk management expectations are well-suited to identify and address potential risks. Furthermore, no immediate gaps have been identified in applying existing model review standards (derived from SR 11-7) for assessing the suitability of AI/ML models. Still, given the growing use of AI/ML within the industry and the expanding range of use cases involved (across both financial and non-financial models), we recognize the importance of industry focus on the particular challenges that AI/ML models may present. Key practices, in this regard, includes the implementation by financial institutions of enhanced and scalable IT infrastructure, the ongoing monitoring of model outputs, the use of targeted internal testing strategies and the implementation of adequate data governance oversight and controls. Furthermore, we believe that the financial industry would benefit from greater clarity on the use of third-party vendor products for AI/ML, given the increasing prevalence of these tools and applications in the market.

DEFINITION OF AI/ ML

According to a November 2017 publication by the Financial Stability Board, AI is defined as *“the theory and development of computer systems able to perform tasks that traditionally have required human intelligence. AI is a broad field, of which ‘machine learning’ is a sub-category. ML may be defined as a method of designing a sequence of actions to solve a problem, known as algorithms, which optimize automatically through experience and with limited or no human intervention.”*³

² ‘SR 11-7: Guidance on Model Risk Management’, Board of Governors of the Federal Reserve System – Division of Banking Supervision and Regulation (April 4, 2011).

³ ‘Artificial Intelligence and Machine Learning in Financial Services: Market Developments and Financial Stability Implications’, Financial Stability Board (November 1, 2017).

Broadly speaking, we agree that AI is simply intelligence exhibited by machines or systems to perform tasks normally requiring human intelligence. Under the umbrella of AI, there are two primary categories of activities requiring different internal risk governance approaches:

- The first category involves the use of ML, which leverages techniques or algorithms to identify patterns in large amounts of data from increasingly diverse and innovative sources to make inferences. ML is used to 'enable' the subjective attributes of AI, including the ability of systems to emulate visual perception, decision making and natural language recognition. ML employs statistical methods or subjective judgment to model inputs, which are then used to derive quantitative outputs. Within this category, there are different types of potential algorithms used based on the level of human involvement required in labelling the data:

Supervised learning is where the algorithm will detect certain patterns (*e.g.* linear relationships) from the training data that are labelled (*e.g.* a defaulted loan is labelled as 'defaulted'), patterns that can then be used to predict labels for the validation data. Supervised learning includes parametric/non-parametric algorithms, support vector machines, kernels and neural networks. Linear regression is a type of supervised learning, but AI/ML algorithms are not constrained to the linear relationship.

Unsupervised learning is where the algorithm is trained on data without any label. The algorithm detects patterns in the data through grouping observations with similar underlying characteristics. Unsupervised learnings include clustering, dimensionality reduction (such as principal component analysis) and recommender systems.

In Reinforcement learning, data is not pre-defined. The goal is to explore or learn a series of actions, rather than predict an output or find a pattern. While supervised learning maps labelled data to a known output, and unsupervised learning finds the pattern and predicts the output, reinforcement learning follows a trial and error method. It learns through delayed feedback by interacting with its environment. In other words, the algorithm is fed an un-labelled set of data, chooses an action for each data point, and receives feedback (including from a human) that helps the algorithm learn. Reinforcement learning can be used in robotics, game theory and self-driving cars.

Deep learning uses algorithms that work in 'layers' (called artificial neural networks), inspired by the structure and function of the brain. Deep learning algorithms can be used for supervised, unsupervised or reinforcement learning. Recently, deep learning has led to remarkable results in diverse fields, such as image recognition and natural language processing ("NLP"). NLP allows computers to 'read' and produce written text.

- The second category of AI relates to classic process automation, such as digital process automation, robotic process automation and business process management improvements. This category employs prescriptive, rule-based logic flow approaches that perform or simulate certain deterministic human tasks without subjective

judgement or forecasting elements. Internally, we subject these AI approaches to a separate technology risk management and governance management framework applicable to traditional software development.

Our observations in this comment letter are focused on the first category of AI/ML activities and their implications for our business model and its specialized focus on meeting the financial services needs of institutional investors.

STATE STREET USE CASES

To help illustrate the importance of a risk management framework for AI/ML that properly accommodates differences in use cases and underlying industry business model, below we provide certain examples of AI/ML-based solutions currently deployed by State Street to help meet our clients' needs and enhance the overall efficiency of our operations. These use cases address several business imperatives in our role as a global custody bank, including the achievement of greater operational efficiencies, improved client service and asset management capabilities, and strengthened risk management and compliance controls.

- Transaction Processing: As part of our trade settlement processing function, we're leveraging deep learning-based computer vision models to locate and identify signatures on manual, faxed trade instructions. This process also leverages an optical character recognition and natural language processing-based machine reading service to extract key information from the unstructured manual trade ticket to automate data entry, thereby improving operational efficiencies and reducing potential risk.
- Client Inquiry Management: During the course of each trading day, we receive a large volume of inquiries from our clients. To accelerate and improve our client service function, we've deployed natural language processing and ML to automatically read client inquiries, classify them by type for efficient routing, extract the relevant information from the inquiry and automate the research and retrieval of additional data from our systems to expedite resolution.
- Regulatory Oversight: As part of our internal risk management controls, we maintain a current and comprehensive catalog of regulatory obligations which are mapped to the required compliance function. We've developed an internal solution to help automate this process, leveraging ML and deep learning trained models to ingest regulatory documents and examination handbooks to identify and classify regulatory obligations by business control type.
- Client Contracts: Our internal contract management system provides for the digitization of client contracts. This solution is foundational to satisfy our 'know your customer' obligations and the appropriate management of our contractual risk. We have recently

enhanced this platform with an AI solution, using ML and natural language processing to classify contracts by type, extract key entity data and type of contract language (*e.g.* governing law, use of data consent, LIBOR clauses, proxy voting, etc.). We also provide an advanced AI-based search capability against the full repository of ingested contracts.

- Fund Compliance: To help support investment fund compliance with prospectus terms, including pre and post-trade obligations, we're deploying an AI-based solution to harvest compliance rules from complex fund documentation. Using deep learning based natural language processing, we are able to automatically identify compliance rule language and classify each rule for subsequent implementation within our rules management system. This solution also automates the discovery of new, modified or deleted rules between different versions of a fund's prospectus documentation.
- Fund Risk Assessment: The process of researching and assigning a risk rating for each of our investment fund clients is a highly manual and labor intensive undertaking, involving the review of hundreds of pages of fund, and in some cases multi-fund, prospectus documents in order to extract counterparty, domicile and other core information, answer a number of risk rating questions, and then feed the data into a risk rating model. Recently, we've deployed a solution leveraging ML and natural language processing to automate the information extraction process and help predict the answers to risk rating questions, thereby greatly improving internal efficiencies and reducing potential risk.
- NAV Quality: State Street is responsible for the calculation and generation of net asset values ("NAV") for tens of thousands of investment funds daily. We've deployed ML solutions to derive a custom benchmark for each investment fund, including what combination and weighting of market indices most closely represents that fund's historical performance. The customized benchmark is then used during the course of the daily NAV calculations to flag anomalies and areas requiring more focused investigation.
- Inadvertent Data Disclosure Prevention: As an additional safeguard to protect our client's data, we are deploying several AI-based checks on reports delivered to clients via e-mail. These checks automatically read the email and any attachments extracting sensitive information, such as client name, fund identification number and account number, and validates the data against a master record for the targeted recipient of the e-mail. The AI solution also detects behavioral differences and anomalies in the communication based on learning from the history of past communications between State Street and each client.
- Anti-Money Laundering Compliance: Like all banks, we have a responsibility to monitor our customers and their transactions that could be related to money laundering and/or terrorist financing and report any unusual activity to the Financial Crimes Enforcement

Network. We have developed new models, architected as a collection of both supervised and unsupervised ML-based estimators that detect potential suspicious activity while reducing the 'noise' and inefficiencies inherent in legacy 'rules-based' logic models.

RESPONSE TO SELECT RFI QUESTIONS

Explainability

Question #1 - How do financial institutions identify and manage risks relating to AI explainability? What barriers or challenges for explainability exist for developing, adopting, and managing AI?

As a general matter, explainability refers to whether an AI/ML model and its output can be understood by humans at an acceptable level. The lack of explainability tends to go hand in hand with the model's complexity. Traditional models (*e.g.* supervised learning models) tend to be more readily explainable, given the use of labelled data and the variables involved in the model development process. Some AI/ ML models, particularly deep learning models, are harder to explain. This is also true of third-party vendor models, given the general lack of access to information needed for ongoing validation.

As part of our existing model risk management framework, we identify and manage potential explainability risk of our AI/ML models through a tiering-based approach. This involves the assignment of our AI/ML solutions to one of three categories (tier 1, 2, or 3) based on the inherent or potential risk that each model may pose. This assessment is based on several factors including:

- Model application/ use;
- Size impact, materiality and volatility; and
- Intrinsic model factors

A model's classification determines the scope, intensity, and frequency of our model risk management activities. Models considered to have the highest inherent risk or potential impact (tier 1 models) receive the greatest level of oversight and governance. Similarly, AI/ML models with a regulatory, risk management or critical business purpose are assigned to higher use priority in the tiering framework. By comparison, models that provide research, informational or internal reporting uses, or workflow enhancements of non-critical functions, are assigned a lower use priority.

Non-linear regression techniques and the use of ML and deep learning approaches on non-structured data and processes tends to generate a higher level of risk than NLP, pattern detection and computer vision approaches on more standardized data sets and activities.

These and other intrinsic risk factors are considered in combination with both use and size impacts and will tend to trigger the 'up-tiering' of the model in our categorization hierarchy. The assessment and tiering of AI/ML models may also take other qualitative or expert judgment-based criteria into consideration, including inventory consistency, portfolio risk considerations, the impact of model linkages to downstream models and reputational risk.

In our experience, the lack of explainability in AI models is especially challenging for third-party vendor products, due to the general lack of access to underlying data. This limitation could be substantively mitigated via regulatory guidance that defines expectations for information sharing by third party vendors of AI/ML solutions.

Question #2 - How do financial institutions use post-hoc methods to assist in evaluating conceptual soundness? How common are these methods? Are there limitations of these methods? If so, please provide details on such limitations.

Consistent with SR 11-7, our internal model risk management framework offers guidance on how to assess and address the conceptual soundness of AI/ML models. Specifically, our model risk management standards:

- Clarify that the model validator 'must determine if the model results are sound, stable, robust and properly interpreted.' A model that does not produce intuitive outcomes when tested over a wide range of inputs and parametric values, is assigned to a high-risk score when assessing robustness.
- Incorporate a detailed framework for the validation of AI/ML models, with preferred methods and statistical testing requirements clearly explained. This includes the use of multiple techniques to diagnose multicollinearity and ways for model developers to resolve this issue.
- Ensure the maintenance of human oversight when modeling choices cannot be fully supported via empirical evidence.

There are various post-hoc methods that can be used to assist in evaluating the conceptual soundness of AI/ML models. This includes Shapley Additive exPlanations (SHAP), Local Interpretable Model-agnostic Explanations (LIME), Partial Dependence Plots (PDP), Accumulated Local Effects (ALE), and Anchors, with the choice of method driven by the underlying facts and circumstances of the use case. Still, it is important to recognize that there is no one-size-fits-all method for the various AI/ML models in use today and that it may be difficult to remediate deficiencies via model redevelopment or recalibration.

Question #3 - For which uses of AI is lack of explainability more of a challenge? Please describe those challenges in detail. How do financial institutions account for and manage the varied challenges and risks posed by different uses?

In our view, concerns regarding the potential lack of explainability in AI/ML models should be defined and assessed on the basis of the underlying use case. This includes the extent to which

the AI/ML model is used to meet either a regulatory obligation or financial purpose. Broadly speaking, use cases can be categorized as follows:

Category 1: Regulatory Models

Regulatory models (such as Comprehensive Capital Analysis and Review models) are expected to allow users to understand the major drivers of the projections under various scenarios. Per SR 15-18 – Supervisory Assessment of Capital Planning and Positions for Firms Subject to Category I Standards, ‘A firm should estimate losses, revenues, expenses, and capital using a sound method that relates macroeconomic and other risk drivers to its estimates.’⁴ Given the challenges involved, the use of AI/ML in regulatory models remains limited, with the exception of anti-money laundering capabilities, where the potential benefits of addressing and eliminating large volumes of ‘false positives’ through the use of AI-based solutions far outweighs potential risks.

Category 2: Decision Making Models

Certain AI/ML models are involved in critical business decisions, such as front office models (*e.g.* trading and asset management) or middle office models (*e.g.* risk monitoring, fraud detection and risk decision making). These models typically augment human functions and are developed by combining human insights, data and statistical methods. The impact of the lack of AI explainability on decision-making models can be significant, depending on a range of factors.

As part of our model risk management framework, each decision-making model goes through a rigorous testing and approval process before implementation. Appropriate investment research and risk professionals are closely involved in the process. Data inputs are well defined and tested independently to ensure appropriate data sets are used. Periodic reviews are conducted to validate model robustness. Even with the above mitigants, decision-making staff is expected to combine the data driven output of AI/ML models, with existing financial theories, market expectation or business intuition, in order to mitigate potential risk.

Category 3: Operational Efficiency

The majority of the AI/ ML models used today by State Street fall within the operational efficiency improvement category. In a limited number of cases, these models might also be used indirectly to support production processes in the form of anomaly or error detection and/or quality control. The impact of the lack of explainability for operational efficiency models is generally limited and therefore the risk may be reasonably accepted by the financial institution.

⁴ ‘SR 15-18: Federal Reserve Supervisory Assessment of Capital Planning and Positions for Firms Subject to Category I Standards’, Board of Governors of the Federal Reserve System – Division of Banking Supervision and Regulation (December 18, 2015; revised January 15, 2021).

Category 4: Research / Benchmarking Models

A number of additional AI/ML models used by State Street fall within the benchmarking category. The impact from the lack of AI explainability on this category of AI/ML models is also limited and the risk can, in most cases, be reasonably accepted by the financial institution.

To help manage underlying challenges and risks, it is common for firms to benchmark complex AI/ML models to simpler models that offer greater transparency and explainability. Also, for some more complex AI models (e.g. deep learning, neural network and image processing models), *post-hoc* methods are often deployed to mitigate or manage the potential risk. This includes:

- Defining quantifiable performance metrics (e.g. back testing of model outcomes) for both testing and monitoring purposes;
- Rigorous testing to assess whether models produce intuitive results across a wide range of inputs, such as a sensitivity test or stability test;
- Construction of simpler models to benchmark and gain explainability;
- Ongoing monitoring using long duration performance metrics (e.g. PAI / R^2 / MSE / Recall);
- Marrying feature importance and the decomposition of model contribution at the feature level, heat map or activation function, with fundamental analysis; and
- Reliance on human intervention and oversight.

Data Processing and Usage

Question #4 - How do financial institutions using AI manage risks related to data quality and data processing? How, if at all, have control processes or automated data quality routines changed to address the data quality needs of AI? How does risk management for alternative data compare to that of traditional data? Are there any barriers or challenges that data quality and data processing pose for developing, adopting, and managing AI? If so, please provide details on those barriers or challenges.

The importance of data quality is critical for prudentially regulated financial institutions and therefore features prominently in our internal risk management processes. Broadly speaking, we rely on two pillars to evaluate and ensure data quality. The first pillar encompasses data validation, where data quality and its implication for model outputs are assessed. Gaps in the quality of data are reflected in the internal risk rating of a model, typically accompanied by remediation actions or compensating controls that place restrictions on model usage. The second pillar encompasses the ongoing monitoring of data, where model owners are required to establish data quality metrics and to monitor these metrics against pre-defined thresholds.

In our experience, AI/ML models may face several data quality challenges, including:

- Bias introduced as a result of data limitations (*e.g.* insufficiently labeled data, or mislabeled data);
- Embedded bias within the data residing in third party vendor products;
- Lack of complete or comprehensive enough data to meet the use case; and
- Limitations in data quantity for the modeling technique chosen.

Alternative forms of data (*e.g.* chat messages) may need extensive data cleaning and processing before use. As such, certain control processes or automated data quality routines, are recommended to help address data quality issues before use.

Question #5 - Are there specific uses of AI for which alternative data are particularly effective?

AI/ML techniques are sometimes utilized to detect anomalies, errors or outliers when a large amount of traditional data or alternative data is used. This includes benchmarking to identify data quality issues, and the use of AI/ML algorithms to detect and manage inadvertent data disclosure. The use of AI/ML for these and other similar purposes is subject to model risk management oversight and control.

Overfitting

Question #6 - How do financial institutions manage AI risks relating to overfitting? What barriers or challenges, if any, does overfitting pose for developing, adopting, and managing AI? How do financial institutions develop their AI so that it will adapt to new and potentially different populations (outside of the test and training data)?

Improperly built AI/ML models are susceptible to overfitting as models built on the training dataset may underperform relative to different datasets. Overfitting risk is particularly hard to manage for ‘off-the-self’ packages (third-party vendor models and open source codes) that may not have been vetted sufficiently to mitigate risk.

We rely on our established model risk management framework to balance variance / bias trade-offs in the use of models. For AI/ ML models specifically, we consider whether the following steps have been properly performed:

- Whether model developers have split the development data across training, testing and validation data sets to balance in-sample and out-of-sample performance;
- Whether model developers have utilized certain algorithms to mitigate overfitting issues;
- Whether model developers have assessed performance of multiple models for comparison purposes (*e.g.*, production or candidate models);

- Whether a model has been retrained to higher layers of the model on a smaller data set that is relevant to the model's use;
- Whether the model was built through "transfer learning", *i.e.* starting from an existing model trained by others using a large amount of new data; and
- For certain models, whether model developers have assessed and continue to assess model performance through real-time parallel runs, *i.e.* assess performance of the model in real-time to see how model outputs adjust to changes in data.

Cybersecurity Risk

Question 7 - Have financial institutions identified particular cybersecurity risks or experienced such incidents with respect to AI? If so, what practices are financial institutions using to manage cybersecurity risks related to AI? Please describe any barriers or challenges to the use of AI associated with cybersecurity risks. Are there specific information security or cybersecurity controls that can be applied to AI?

To date, we have not experienced particular cybersecurity risks or incidents involving our use of AI/ML models. This is partially a function of the underlying use cases, which as previously noted, focus on enhancements to internal operational efficiencies, improved client service and asset management capabilities, and strengthened risk management and compliance controls.

As a general matter, the greatest potential cyber-vulnerability that AI/ML models face involves the penetration, capture and/or compromise of the underlying data. This includes, for instance;

- Poisoning of training data intended to influence or taint model outputs;
- Manipulation of data inputs in order to confuse and/ or disrupt model outputs;
- Reverse engineering, or inferring potentially sensitive data (data leakage) that was used to train a model; and
- Loss or theft of a trained model; both the IP of the model architecture and the learned weights.

These and other similar risks are particularly elevated when using open source frameworks, libraries and model architectures for the development and deployment of AI/ML models. Moreover, these risks are compounded by other more general information technology considerations, such as the use of public clouds, data security, access credentials and secrets management.

As with any information system or application, potential cybersecurity risk in AI/ML models is managed by banks, such as State Street, via the deployment of segmented network controls, highly integrated and sophisticated information security instrumentation and least privilege control programs. This includes the initial implementation of advanced security architecture and risk mitigation controls, the active 24/7 monitoring of core systems and their use, and the

deployment of penetration testing and other security protocols to identify and remediate potential vulnerabilities.

These programs reflect regulatory standards adopted by the Agencies to promote the use of appropriate industry-wide capabilities, including the Interagency Guidelines Establishing Information Security Standards for the safeguarding of customer records and information under Section 501(b) of the Gramm-Leach-Bliley Act.⁵

In our experience, the most effective approach for addressing potential cybersecurity risk in AI/ML models involves the deployment of comprehensive identity and access control procedures to manage access to key functions, with progressively more restrictive entitlements based upon the criticality of the underlying data. This includes:

- Explicit application execution authorization for the control of model training/retraining;
- Explicit application execution authorization for the submission of inputs for model inference and prediction;
- Data access entitlement to access training data;
- Data access entitlement to access or deploy a trained model;
- Data access entitlement to access data inputs submitted to a model for inference;
- Use of audit logs to capture who trained the model, when, and with what data; and
- Monitoring of model use, including input feature drift detection and alerting.

More broadly, AI/ML model use should follow the normal software development lifecycle process, including source reviews, vulnerability management and approvals. Also relevant is the use of firm-wide IT resiliency, disaster recovery and business continuity practices.

While we do not believe that prudentially-regulated financial institutions currently face challenges in the management of cybersecurity risk for AI/ML models that cannot be addressed based on existing supervisory guidance and regulation, we would urge the Agencies to deepen their existing collaboration with the industry, including the sharing information on specific threats and risks, to help identify emerging vulnerabilities that can inform the design of effective control systems and procedures.

Dynamic Updating

Question #8 - How do financial institutions manage AI risks relating to dynamic updating? Describe any barriers or challenges that may impede the use of AI that involve dynamic updating. How do financial institutions gain an understanding of whether AI approaches producing different outputs over time based on the same inputs are operating as intended?

⁵ 'Interagency Guidelines Establishing Information Security Standards', Board of Governors of the Federal Reserve System, Federal Register Volume 79, Number 126 (July 1, 2014), pages 37166-37167.

Dynamic updating is not unique to AI/ ML and can be present in other more traditional models subject to data-driven model selection algorithms or frequent parameter recalibration. As such, our existing model risk management framework, comprised of prescribed review standards, model change control mandates and ongoing monitoring requirements, provides a robust framework for assessing and managing the risk of dynamic updating in model usage generally. This includes both the initial ‘point-in-time’ validation of the model and the revalidation of the model over time.

In accordance with our existing model risk management framework, all models including AI/ML models are required to incorporate the frequent monitoring of input, processing and output components. Model owners are required to report any significant changes in the data, methodology, model use or external environment. Any significant changes to the model environment may, in turn, trigger an independent validation. Models subject to dynamic updating are required to monitor impacts against appropriate thresholds and to report material changes as they occur. In many cases, we employ champion/challenger model comparisons, which enable the model user to observe the performance of updated models compared to the existing incumbent model. For third party vendor products, it may be difficult to enforce such enhanced due diligence, since changes are not necessarily disclosed to model users.

Given the potential implication of dynamic updating for AI/ML models, it is important for firms to ensure the monitoring of certain core processing components (*e.g.* features, weighting, feature contributions) and output components (*e.g.* back-testing, use of benchmarking models, correlation or autocorrelation of outputs). This monitoring is particularly helpful in identifying whether the AI/ML model has changed substantively over time. The empirical assessment of the observation with human oversight plays an important role in mitigating this potential risk.

Oversight of Third Parties

Question #10 - Please describe any particular challenges or impediments financial institutions face in using AI developed or provided by third parties and a description of how financial institutions manage the associated risks. Please provide detail on any challenges or impediments. How do those challenges or impediments vary by financial institution size and complexity?

In general, we see greater challenges for prudentially regulated financial institutions when using third party vendor-based AI/ML models. While banks, such as State Street, are able to conduct due diligence on vendor products at the time of purchase and fulfill risk oversight and control on a periodic basis, the “black box” nature of third-party vendor products (*e.g.* no disclosure of the model construction) and the impact of dynamic updating may expose financial institutions to significant risk if legal contracts do not adequately limit its potential liability. Also, third party vendor products shared by major financial institutions may pose systematic risk due to the scope and ubiquity of use. As such, we recommend that the Agencies consider

greater regulatory guidance on this topic. This includes the level of input, processing and output monitoring that is required of third-party vendor products, requirements for independent validation, the frequency of review and testing, and the enforceability of an industry-wide standard for model due diligence.

CONCLUSION

Thank you once again for the opportunity to respond to the RFI and the key matters raised therein. In general, we believe that prudentially regulated financial institutions are well equipped to address challenges related to the increased use of AI/ML models in their products, services and operations, and that existing expectations regarding the management of model risk are sufficiently comprehensive and robust to address potential risks.

The model validation activities of banks are performed according to model review standards that are comprehensive, covering all aspects of the model (end-to-end), assessing the appropriateness and integrity of data, ensuing conceptual soundness, computational accuracy, performance and stability, and the deployment of appropriate ongoing monitoring activities. Furthermore, we see no immediate gaps in the use of the existing model review standards (derived from SR 11-7) for assessing AI/ML models.

Still, given the financial industry's increasing use of AI/ML and expanded business use cases on the horizon, we recognize the importance of identifying and addressing the particular challenges that AI/ML models may pose over time. In our view, these challenges can be mitigated by banks and other financial institutions by:

- o Enhancing the ongoing monitoring framework, notably in the areas of data quality, model updating and model performance;
- o Exploring opportunities for enhancement in the validation process, including the development of targeted testing strategies to assess model risks;
- o Implementation of enhanced IT-infrastructure, data governance oversights and controls;
- and
- o Creation of industry working groups to share AI/ML knowledge and business use cases.

Finally, we note the particular challenges raised by the use of third-party vendor-based AI solutions and believe that prudentially regulated financial institutions would greatly benefit from additional clarification or guidance from the Agencies on expectations for their use, management and oversight.

Please feel free to contact us should you wish to discuss the contents of this submission in greater detail.

Sincerely,



Randy Swanberg
Senior Vice President
Global Head of Automation and Artificial
Intelligence
rswanberg@StateStreet.com



Julia Hua Li
Managing Director
Global Head of Model Risk Management
hjli@StateStreet.com