July 1, 2021

Chief Counsel's Office
Attention: Comment Processing
Office of the Comptroller of the Currency
400 7th Street SW
Suite 3E-218
Washington, DC 20219

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

James P. Sheesley
Assistant Executive Secretary
Attention: Comments-RIN 3064-ZA24
Federal Deposit Insurance Corporation
550 17th Street NW
Washington, DC 20429

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

Melane Conyers-Ausbrooks
Secretary of the Board
National Credit Union Administration
1775 Duke Street
Alexandria, VA 22314-3428

Re: Request for Information and Comment on Financial Institutions' Use of Artificial Intelligence, Including Machine Learning (Docket No. OCC 2020-0049; OP-1743; RIN 3064-ZA24; CFPB 2021-0004; NCUA 2021-0023)

Ladies and Gentleman:

The Consumer Bankers Association[1] ("CBA") appreciates the opportunity to respond and comment on the Request for Information ("RFI") issued by the Board of Governors of the Federal Reserve System, Bureau of Consumer Financial Protection, Federal Deposit Insurance Corporation, National Credit Union Administration and the Office of the Comptroller of the Currency (collectively "the Agencies"). Artificial

[1] The Consumer Bankers Association partners with the nation's leading retail banks to promote sound policy, prepare the next generation of bankers, and finance the dreams of consumers and small businesses. The nation's largest financial institutions, as well as many regional banks, are CBA corporate members, collectively holding two thirds of the industry's total assets.

intelligence ("AI") and machine learning ("ML") models provide efficient, accurate, and predictive decisioning, which in turn expands access to financial services for all consumers. The banking industry is constantly innovating to meet the expectations of their customers. We have witnessed the benefits AI and ML have provided the banking industry, and we look forward to new ways to thoughtfully incorporate technology in our products, services and operations in the future. Though the benefits of AI and ML are many, CBA supports the Agencies' desire to ensure continued responsible innovation.

While AI and ML models may present different risks than traditional models, banks are already subject to a thorough regulatory framework designed to mitigate any increased risks and they are constantly improving these processes. Non-banks do not face similar regulation. Understanding the conceptual soundness of any model, tool, application or system aids in managing its risks, including those related to lack of explainability. The importance of conceptual soundness is described in existing agency guidance and is well established in industry practice. In turn, AI and ML models provide more efficient, accurate, and predictive decisioning, which ultimately expands access to financial services for all consumers.

CBA believes consumers need equal and consistent protections, regardless of where they secure credit. Consistent Model Risk Management (MRM) standards are needed among all financial regulators. A lack of consistent standards between regulators and across institutions leaves some consumers at risk, especially from entities providing financial services outside the well-regulated banking industry, by not affording equal protections to all consumers.

This letter details the intersection of AI and ML with: explainability; data quality and processing; overfitting and dynamic updating; and cyber security risk and third-party oversight. Specific examples from CBA member banks are incorporated throughout to demonstrate how banks continue to manage this powerful, ever-evolving technology.

Explainability

Banks are Thoughtfully Integrating AI and ML in their Banking Organizations

Banks have substantial experience identifying and mitigating those risks which may present themselves more frequently when using AI and ML by leveraging their existing risk management frameworks and with the involvement of multiple functions throughout the model development lifecycle, including independent model risk management teams, legal, risk and compliance, among others. By applying the same principles of model governance to AI models that banks use with traditional models, banks ensure risks related to these models are appropriately managed. In addition, explainability risks are not limited to AI and ML models. Humans frequently are unable to explain their decisions accurately, especially when they rely on subjective reasoning and judgement. Bankers have discovered AI techniques provide the most value when exposed to large amounts of data, both traditional and alternative. As a result, AI techniques can extract useful, subtle and predictable information from variables or factors which may not have been possible from traditional models.

Banks exercise extensive due diligence on data sources that may be used to develop ML algorithms, especially when using third-party or non-traditional data to manage explainability risks. Model developers take extreme caution when deciding which variables or factors to include in their models. Banks often tailor alternative and ML data sources per each business use to avoid accidental correlations between

selected variables. Many of the effective controls banks use today for traditional data sources in their own human decisioning can be used to minimize potential risks of using AI and ML.

One of the most important options to manage risk related to AI explainability is to exercise exhaustive due diligence on data sources that could potentially be used to develop algorithms, particularly if third party data or non-traditional data is also utilized. Secondly, model developers use extreme caution when deciding which variables or factor to include in their models. Variables which are suitable for the business use case are used to avoid accidental correlations between selected variables. For example, a model that ingests a lot of non-traditional data such as social media data is not appropriate for developing underwriting algorithms.

Banks are Well-Suited to Navigate Areas of Ambiguity

Due to the "black box" nature of some of the more complex AI or ML models, relationships between the variables and predicted outcomes may not be easily discernable. Therefore, model developers have to resort to either "post-hoc" techniques that are now beginning to be widely used such as Shapely values or LIME ("Local Interpretable Model-agnostic Explanations") or develop more intuitive challenger models using traditional techniques such as logistic regression or decision trees.

In terms of post-hoc methods, academic researchers, as well as practitioners in AI/ML, have begun utilizing some model explainability techniques which enable the end users to examine the decision making of the model both at a global and local level. On a global level, these techniques can help users understand which features contribute to the model outcome and the extent to which they influence the decision of the model. Additionally, on a local level, these techniques can help explain why the model produced certain decisions and provide reasons for the results, which may provide insight into the model's shortcomings. CBA notes the development and use of opaque models present material risks to consumers without these post-hoc tools or effective MRM framework.

For instance, Shapely value is one such method, a concept derived from game theory in which is used to determine contribution of an individual player in a coalition or a cooperative environment. Within the context of AI models, if one were to conceive each feature or variable as a player with the game being the prediction of the target variable and the score being the model output, the Shapley value shows the average marginal contribution of a feature across all possible combinations.

In addition, LIME generates an explanation by approximating the black-box AI model by using components from an interpretable model (for example, coefficients in a linear regression model) in the neighborhood or vicinity of the instance.

Lastly, many practitioners simultaneously develop different models using traditional algorithms such as Logistic Regression or Decision Trees, and then compare final model variables, their correlations and dependencies to predict outcomes obtained using these models to those obtained from the AI based model.

Banks rely on a combination of development processes and post-hoc evaluation and ongoing performance monitoring tools and controls to mitigate the risks of opaque models or decision making. These processes allow banks to better identify when an opaque model is not working as intended. Unless otherwise required by law (e.g. fair lending), opacity is not a bar to model use. Explainability tools are quickly evolving, so, if necessary, any guidance should remain flexible so as not to stifle innovation. CBA also

believes non-banks, especially non-bank consumer lenders, should be subject to similarly explainability risk management expectations as banks.

Data Quality and Processing

As heavily regulated and supervised entities, banks have a long history of safely validating models for consumer products and services. CBA members do not find data quality and processing for AI models significantly different from validating non-AI models. For example, an institution would continue to perform the same risk assessments, note potential issues and implement mitigating techniques to prevent consumer harm. The major difference with AI models is the need for computer resources since the data sizes are significantly larger. This is especially true when an AI model is built on transaction level data.

The Agencies should more closely study the use of unstructured data by nonbank financial companies. Unstructured data, e.g., texts, images, audio, may contain customer's personally identifiable information ("PII") which would need scrubbed before being processed. The absence of this step is a privacy concern to consumers.

CBA members view an AI/ML model's success as relying upon correctly annotated and labeled data, which needs to be consistently performed by human subject matter experts. There are negative impacts to the application of models if the quality and consistency of the data is not strictly followed.

Alternative Data Risks

A level playing field for all stakeholders, where non-bank lenders are held to the rigorous standards applied to banks, will ensure all consumers receive effective protections of their data. Consistent examination across lenders allows federal regulators to detect misuses of alternative data. Consistent regulation and robust compliance management provide effective guardrails to ensure diverse forms of alternative data are protected and improves financial inclusion and access to responsible credit.[2]

Alternative data has helped refine AI/ML models for more efficient banking. The RFI defines alternative data to mean, "information not typically found in the consumer's credit files of the nationwide consumer reporting agencies or customarily provided by consumers as part of applications for credit." With that being said, there is a broad range of alternative financial data available.

All lenders using alternative data in a credit models should be subject to the appropriate risk management oversight, similar to the current standards.[3] The CFPB should join this guidance and examine non-bank lenders for adherence. The risks of alternative data are magnified when lenders rely on vendor models that offer little insight into the type and quality of the data being used. Models sourced from vendors

[2] Statement of Michael J. Hsu, Acting Comptroller of the Currency before the Committee on Financial Services United States House of Representatives (May 19, 2021) ("My primary concern is that the regulatory community is taking a fragmented agency-by-agency approach to these trends, just as it did in the 1990s and 2000s. To the extent there is interagency coordination, it tends to be tactical, to deal with a pressing issue, such as Facebook's Diem. The key strategic question which the regulatory community must answer collectively is: Where should we set the regulatory perimeter? To my knowledge, there is not a shared understanding of the answer to that question and no overarching strategy to achieve it").

[3]

should be subject to consistent third-party risk management ("TPRM") vendor oversight and data governance expectations.

The Agencies should also implement certain consumer rights for alternative data for non-bank financial companies to ensure effective consumer protection. Customers who ascertain credit from non-bank financial companies should have protections similar to the Fair Credit Reporting Act ("FCRA") safeguards, including the right to see the alternative data used in credit decisioning and dispute rights to ensure its accuracy. Like traditional data subject to the FCRA, there should be an age limit on negative alternative data. CBA urges consistent oversight across lenders, bank and non-bank, to better detect misuses of alternative data as consistent and comprehensive oversight will ensure the data and models using the data are appropriate, related to creditworthiness, and fair to all consumers.

Use of alternative credit data should be held to the same rigorous regulatory and examination standards as traditional credit data. This data should have the same ability to dispute, error resolution, accuracy, and transparency standards as required by the various regulations and Agency guidelines. The Agencies should continue engaging with stakeholders to ensure rules, guidance and examination standards best serve consumer protection initiatives.

Overfitting and Dynamic Updating

Banks validation and model governance frameworks have developed effective ways to mitigate AI/ML technology risks. The MRM Guidance ensure models are managed appropriately throughout their lifecycle, regardless of model methodology. Overfitting model risks are more prone to occur with AI/ML based models as compared to models developed using traditional methods. This is primarily because AI/ML based models are extremely efficient at learning and extracting signals from all presented information even from factors that may be sparsely populated or could even be noise, which in turn causes the model to become too attuned to or specific to the data on which the model was developed. As a result, the model loses general applicability and is unable to perform as well on other data sets leading to inaccurate results.

To better understand and to mitigate the risks of overfitting, FIs will typically partition the model development data into training and holdout validation samples to evaluate model performance and stability on the holdout sample. Additionally, banks may also use cross-validation in which the training data set is iteratively partitioned into different model development and validation samples, and then model performance is averaged on each of the validation partitions in order to measure overall model performance. This is done primarily to reduce sample bias, since the original validation partition of the data set was by chance and not representative of the overall population. If not, the resulting model could falsely appear to exhibit a good fit and the model may not be generally applicable to future data sets. Moreover, in order to validate that AI/ML models exhibit stable performance over time, banks will also generally score out-of-time samples to evaluate model performance metrics to ensure stability in model's predictive power.

AI/ML practitioners also use other techniques to control for overfitting such as Feature Removal, Early Stopping, Regularization and Ensembling. In Feature Removal, the model developer can manually remove seemingly irrelevant features to improve model generalization while in Early Stopping the model developer can stop the model training process before the AI/ML algorithm passes the point when additional iterations yield marginally less improvements to the model. With Regularization, model

developers force the AI/ML model to be simpler by adding penalty parameters or cost functions in the model fit process. Finally, AI/ML model developers will also often use Ensembling in which predictions from several separate AI models are combined to yield more "smoothened" predictions which also tend to reduce overfitting.

Model explainability also can be useful in illustrating variables with unintuitive relationships to the target variable and which should be considered for exclusion. Certain machine learning algorithms also have special regularization techniques as well, which can be applied during model training, tree boosting algorithms like XGBoost and neural networks are two key cases.

Explainability assessments can also help detect illogical drivers of model behaviors due to spurious correlation in data artifacts. Banks also use adversarial testing through perturbing data inputs to gauge the extent to which it degrades the AI/ML model's performance and reduces its effectiveness.

CBA believes effective independent model validation and model governance help ensure AI/ML will adapt to new and different populations. To better protect consumers, we urge the agencies to extend these model risk management requirements to nonbank financial companies.

Cyber Security Risk and Third-Party Oversight

Cybersecurity is Already Deeply Embedded in Bank Risk Management Programs

Banks already implement and maintain information security programs, consistent with Interagency Guidelines mandated by GLBA, to identify and control potential increased adversarial machine learning threats. Applications aimed at detecting malicious behavior such as fraud are susceptible to adversarial countermeasures. This is not unique to such applications that rely on machine learning to aid in defense. In some situations, machine learning makes these systems less susceptible than traditional rule-based approaches. At the same time, there are new and unique cyber-attack vectors specific to the application of machine learning. Monitoring for these emerging threats falls within model risk management and where appropriate banks coordinate with law enforcement.

Third Party Oversight is Subject to Strict Due Diligence Standards.

Effective third-party risk management ("TPRM") processes can control for increased risk created by using third-party products and services. Banks comply with existing TPRM guidance from the regulators to develop and manage third-party relationships. Effective due diligence and compliance with the Agencies' third-party risk management guidance helps banks better manage third-party model risk.

However, there may still exist some challenges to utilizing third-party AI and ML models. The most significant challenge with these engagements is in fully understanding the AI functionality and services provided since many AI-based solutions tend to allow for limited visibility into their inner workings and processes. Additionally, due to the proprietary and confidential nature of their solutions, most third-parties are reluctant to share the details behind their solutions. However, the onus often falls solely on banks to continue their due diligence obligations by trying to understand these solutions to whatever extent possible. CBA encourages the Agencies to provide additional clarity on regulatory expectations with respect to due diligence on models provided by third parties.

Fair Lending

Techniques to Determine the Fair Lending Compliance of AI-Based Models

Many lending products offered by banks are subject to demographic data collection restrictions under either the Equal Credit Opportunity Act ("ECOA") or the Home mortgage Disclosure Act ("HMDA") and are limited in conducting direct testing or evaluation of compliance of AI-based determination approaches with fair lending laws. However, in order to comply with their obligations to conduct the fair lending analysis, many FIs resort to either proxy assignments or the Bayesian Improved Surname Geocoding (BISG) proxy method. Under the proxy assignment methodology, a consumers name along with his or her geographic location is used to arrive at a "best guess" of the consumer's demographics (e.g., sex, race, and ethnicity). Due to the heuristic nature of this methodology, it is subject to relatively high degree of inaccuracy, which could lead to misleading results. As an alternate, many FIs utilize the same methodology which the CFPB uses in their analyses, and which combines geography- and surname-based information into a single proxy probability for race and ethnicity. While this method also has its limitations, the CFPB acknowledged that "this approach produces proxies that correlate highly with self-reported race and national origin and is more accurate than relying only on demographic information associated with a borrower's last name or place of residence alone."

Many analyses can potentially be conducted in order to evaluate the fair lending impact of AI-based credit models. These primarily include examination of two types of impacts:

1. Direct impact in which protected attributes explicitly result in non-favorable outcomes toward consumers in protected classes, or

2. Indirect impacts in which consumers may appear to be evaluated based on seemingly neutral and/or non-protected attributes but results in protected classes appearing to be treated negatively.

When conducting the fair lending analyses, for example in the context of underwriting, FIs can examine disparate impact on various metrics including requiring equal approval rates within each demographic segment or requiring equal predicted default rates of approved applicants within each demographic segment, or requiring equal model precision within each demographic segment.

When models are evaluated or controlled for equal approval rates, the requirement would be that each demographic class have the same percent of applications be approved; however, since different demographic groups perform differently from credit risk perspective, this would result in different default rates within each demographic segment. On the other hand, when the model is controlled for equal predicted default rates of approved applicants, the requirement would be to approve different proportions within each demographic segment so that each group then exhibits the same credit risk default rate. As an example, this could be achieved by having lower approval rates for demographic segments that perform worse and exhibit higher default rates. Finally, when the model is controlled for precision, as measured by the ratio of correct outcomes (i.e., actual defaults) divided by the number of all predicted outcomes (i.e., predicted defaults), results are equalized across each demographic segment. Or in other words, the model should accurately be able to predict the same percent of credit risk defaults within each demographic segment.

However, it is very difficult for models whether they be traditional or AI-based to be able to control for all these metrics simultaneously. Hence, FIs have to evaluate the tradeoffs of each of these approaches

against each other in order to determine which exposes the FI to the least fair lending or disparate impact risk.

Fair Lending Risk Management Techniques with AI Model Testing

Any model whether it be traditional or AI-based can potentially be biased and/or result in discrimination on prohibited bases. However, with AI based models the analysis becomes a little more challenging due to the "black box" nature and lack of transparency into the inner workings of these models. However, in order to attempt to reduce potential risks of discrimination, FIs adopt various methods to try to mitigate biases in the data used for training, impose bias mitigation techniques during model training as well as other approaches post model training.

First FIs will suppress factors or attributes that may be overtly correlated to prohibited factors under ECOA. Next FIs model builders will attempt to correct for class imbalances in the model development data set by, for example, sampling under-represented groups or down sampling over-represented segments. During the model training process, FIs will also attempt to correct models for fairness by eliminating or dropping variables that may seem to be highly correlated to the outcome being predicted but have no obvious relationship to the predicted outcome. Post model training, FIs will ensure there is strong business justification for the factors that are include in the final model and validate the models against holdout and out of time samples to determine model stability and predictive accuracy. Finally, by the use of proxy assignments, e.g. BISG, FIs will evaluate whether the model results in non-favorable outcomes toward consumers in protected classes.

Flexibility to Innovate with ECOA

CBA believes there is enough regulatory flexibility within the Equal Credit Opportunity Act and Regulation B to engage in the use of machine learning-based credit decisioning through risk-based business determinations. There should be consistent standards for using artificial intelligence and machine learning between banks and non-banks.

Conclusion

CBA encourages the Agencies to recognize that while AI/ML models currently have the potential to present risk, effective MRM practices ensure AI/ML models are no more risky than traditional models. MRM guidance requires critical analysis through the development, implementation and use of complex algorithms like AI, and sets supervisory expectations for independent review of models to confirm they are functioning as intended. Most importantly, consistent regulation and robust compliance management provide effective guardrails for AI/ML models including, but not limited to, the following: privacy; improved financial inclusion; and responsible access to credit.

If you have any additional questions or concerns, please do not hesitate to contact André Cotten at 202-552-6360 or at Acotten@consumerbankers.com.

Sincerely,

CONSUMER BANKERS ASSOCIATION

Andre' B. Cotten, Esq.
Assistant Vice President, Regulatory Counsel
Consumer Bankers Association