July 1, 2021

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

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

James P. Sheesley
Assistant Executive Secretary
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

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

To Whom It May Concern:

The American Financial Services Association (AFSA)[1] appreciates the opportunity to write to the Office of the Comptroller of the Currency, Board of Governors of the Federal Reserve System, Bureau of Consumer Financial Protection, and the National Credit Union Administration (collectively, the "Agencies") in response to their request for information (RFI) regarding the use of artificial intelligence (AI), including machine learning (ML), by financial institutions (FIs). We support the Agencies' efforts to gather feedback from all stakeholders related to appropriate governance, risk management, and controls over AI.

AFSA believes that AI and ML can help expand access to responsible consumer credit while at the same time satisfying the principles of anti-discrimination, accuracy, and education underlying Equal Credit Opportunity Act (ECOA) and its implementing regulation (Regulation B). Below, we respond to the questions posed by the Agencies in the RFI.

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?

[1] Founded in 1916, AFSA is the national trade association for the consumer credit industry, protecting access to credit and consumer choice. AFSA members provide consumers with many kinds of credit, including traditional installment loans, mortgages, direct and indirect vehicle financing, payment cards, and retail sales finance.

One way that FIs can identify and manage risks relating to AI explainability is by starting with a development dataset with which the institution is already familiar. In this respect, it is important to distinguish AI/ML models developed in a closed, non-production environment from AI/ML models that learn dynamically.[2] This helps to manage risk by containing the model's "learning" to parameters that are similar to the outcomes produced by a regression model.

As with any model, whether AI/ML or regression, FIs adhere to statistically sound development standards and practices to ensure predictability and avoid overfitting. Therefore, if FIs were to use dynamic learning where AI explainability challenges could be more pronounced, FIs could establish parameters to contain and control the capabilities of the AI/ML as discussed further in Answer #8.

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 (whether to explain an AI approach's overall operation or to explain a specific prediction or categorization)? If so, please provide details on such limitations.

FIs monitor the performance of models (*i.e.* post-hoc) against predicted outcomes observed in model development. Consistent with the official interpretation of 12 C.F.R. § 1002.2(p), "Periodic revalidation," post-hoc monitoring provides assurance that variables in the models have values that are distributed as expected (*i.e.* as developed). For example, each variable in a model may have its own one-time distribution created when the model was developed (reference vintage) and a newly created distribution (current vintage) specifically for a given period being reviewed. Institutions can examine performance differences between those two distributions. If a comparison is consistently different from the reference value, then action may or may not be taken. This process can be done with AI/ML models, just as it can with regression models. In addition to these post-hoc processes, FIs commonly use other conceptual tools to ensure statistical soundness, such as concordance and accuracy measures between development datasets and production datasets.

One limitation for AI/ML models specifically is that variables often have interrelated effects on, and from, other variables in the model. While the validation comparison can still be performed, it may be more challenging to tweak the models—*i.e.* forced directionality, reduced weight. Where needed, institutions may neutralize variables to return models to maintain predictive ability.

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?

Uses of AI that do not have a baseline standard for a "good" or "bad" outcome present more of a challenge than those that do. For example, AI/ML that operates in consumer-facing environments

[2] Throughout this response, AFSA's discussion of AI and ML in model use is primarily focused on "credit process models," or those used in the extension, servicing, and/or collection process of consumer credit. It is important to distinguish other types of AI/ML use cases from credit process models because in the latter cases, development is centered on "good" vs. "bad" outcomes (*e.g.* payment). This concept is discussed more fully in Answer #3. However, AFSA is not suggesting that emerging uses of AI/ML (ex. behavioral models) are less significant or valuable for purposes of this RFI.

may use a range of possible paths to determine the "next best action" to suggest or resolve a question or concern. To illustrate, a chatbot may use AI/ML to make recommendations to a consumer on resolving an issue or learning more about a product. On the other hand, where a standard exists for good or bad outcomes—such as in repayment for an extension of credit—an AI/ML model is more intuitively explained. While this area is not entirely without "explainability" challenges, such as in identifying the reasons for a particular decision in credit underwriting, modeling techniques and Regulation B provide flexibility to produce these "explanations" in ways that are both accurate and consistent with law. As the CFPB has previously noted—ECOA and Regulation B offer flexibility in adhering to requirements such as producing the list of reasons for adverse action.[3]

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.

Data quality and processing present similar risks under both a regression and an AI/ML model scenario when used on traditional data.[4] Therefore, when AI/ML techniques are used on traditional data, similar risk management processes used for regression models would also be used. These processes may include comprehensive review and testing of the statistical soundness of a model (*i.e.* concordance, accuracy) and its component variables. Additionally, each variable can be reviewed for any potential fair lending risks, as appropriate. Models developed with AI/ML techniques can produce variables that are non-monotonic, making them somewhat different from traditional regression variables. However, any variable may be removed if its effect cannot be reasonably explained.

Managing risks for data quality and processing with alternative data[5] presents a different challenge because an FI may have to take additional steps to gain confidence that the data satisfies any applicable regulatory expectations. While AFSA believes this obligation primarily falls on consumer reporting agencies under the Fair Credit Reporting Act, FIs need to have confidence in the data for legal and business purposes. These challenges are more pronounced when the data comprises a vended score that the FI is either not technically or contractually able to audit, examine, and review. This can because the vendor wishes to protect valuable trade secret information in the way that its score or data is produced. FIs, however, take their responsibility seriously to ensure

---

[3] Ficklen et al.,"Innovation spotlight: Providing adverse action notices when using AI/ML models," *Consumer Financial Protection Bureau* (blog), July 7, 2020, https://www.consumerfinance.gov/about-us/blog/innovation-spotlight-providing-adverse-action-notices-when-using-ai-ml-models/

[4] "Traditional data" refers to data assembled and managed in the core credit files of the nationwide consumer reporting agencies, which includes tradeline information (including certain loan or credit limit information, debt repayment history, and account status), and credit inquiries, as well as information from public records relating to civil judgments, tax liens, and bankruptcies. It also refers to data customarily provided by consumers as part of applications for credit, such as income or length of time in residence.

[5] We refer to "alternative data" here as defined by the Request for Information Regarding Use of Alternative Data and Modeling Techniques in the Credit Process.

that their service providers, including vendors providing such models, comply with applicable law and avoid consumer harms.
AFSA notes that the CFPB is currently drafting a rule under § 1033 of the Dodd-Frank Act which allows consumers to access to data that financial institutions have accumulate about consumers. Other federal agencies should coordinate with the CFPB to ensure consistent rules across agencies on the same subject.

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

While still in its early phases, alternative data may have potential for providing access to the underbanked and credit invisible consumers. This could offer benefits to both consumers and FIs, but AFSA does not believe the benefits are limited to AI/ML models, as the data could also be used in regression models.

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)?

As discussed in Answer #2, as part of both their regulatory and risk management processes, FIs routinely conduct analysis of models to ensure that the results from development are consistent with those observed in production.

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?

The Gramm-Leach-Bliley Act (GLBA) requires FIs to explain how they share and protect their customers' private information. FIs already implement and maintain information security programs, consistent with the Interagency Guidelines mandated by GLBA, to identify and control 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 ML to aid in defense. In some situations, ML makes these systems less susceptible than traditional rule-based approaches. At the same time, there are new and unique attack vectors specific to the application of ML. Monitoring for these emerging threats falls within MRM and, where appropriate, FIs coordinate with law enforcement.

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?

Dynamic updating is an aspect of AI/ML modeling that may not be necessary to potentially expand upon traditional models. In general, managing risks in the model development process can always be mitigated through human review and oversight to ensure that outputs are logical and explainable. FIs may choose to manage model complexity and over-fitting by setting minimum and maximum parameters on the AI/ML "learning" that takes place. In addition, "randomness" can be introduced in the training data to guard against overfitting and ensure that the modeling techniques produce robust, valid results. As with any predictive modeling, consistency is critical not only to an FI's compliance program, but also with its own interest in maintaining a prudent risk portfolio. FIs benefit from extending credit to consumers who can afford to repay, and therefore the incentives of using AI/ML in developing models are aligned with the goals of expanding responsible access to credit to consumers who are creditworthy.

Question 9: Do community institutions face particular challenges in developing, adopting, and using AI? If so, please provide detail about such challenges. What practices are employed to address those impediments or challenges?

No additional comments.

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?

The primary challenge that FIs face in using AI solutions developed or provided by third parties is gaining sufficient confidence in the performance of the solution and its compliance with regulatory expectations.[6] For example, if a third party refuses to share its inputs or methods that helped build the AI solution for fear of disclosing trade secrets, it may be impossible for the FI to confirm the solution was appropriately validated and complies with the FI's compliance program. Currently, FIs that aim to expand access to credit do not have a transparent way to confidently determine fair lending compliance of third-party solutions.

Question 11: What techniques are available to facilitate or evaluate the compliance of AI-based credit determination approaches with fair lending laws or mitigate risks of noncompliance? Please explain these techniques and their objectives, limitations of those techniques, and how those techniques relate to fair lending legal requirements.

FIs could use various techniques to evaluate the compliance of AI-based credit determination approaches; however, today there is little guidance on the way that any particular approach would be treated. Because ECOA prohibits the collection of race/ethnicity information, FIs may perform fair lending testing (*e.g.,* using the Bayesian Improved Surname Geocoding (BISG) proxy method) to evaluate the compliance of AI-based credit determination approaches. However, not all FIs believe such testing is required or appropriate under ECOA and Regulation B. AFSA believes that FIs should be given flexible guidance or tools to safely evaluate fair lending compliance, and not mandated to use a specific form of analysis.

[6] As discussed in Response #4, this challenge can be exacerbated by the use of alternative data.

Question 12: What are the risks that AI can be biased and/or result in discrimination on prohibited bases? Are there effective ways to reduce risk of discrimination, whether during development, validation, revision, and/ or use? What are some of the barriers to or limitations of those methods?

See answers to Questions 4, 10, and 11.

Question 13: To what extent do model risk management principles and practices aid or inhibit evaluations of AI-based credit determination approaches for compliance with fair lending laws?

See answers to Questions 4, 10, and 11.

Question 14: As part of their compliance management systems, financial institutions may conduct fair lending risk assessments by using models designed to evaluate fair lending risks (''fair lending risk assessment models''). What challenges, if any, do financial institutions face when applying internal model risk management principles and practices to the development, validation, or use of fair lending risk assessment models based on AI?

While an FI may use proxy methodologies (*e.g.,* BISG) to test for fair lending risk, it is important to note that these methodologies are not 100% accurate. Furthermore, in considering whether any potential fair lending concerns exist within a certain set of alternative data variables, the analysis can quickly become difficult when trying to identify whether a novel variable could potentially be a proxy (however attenuated it may be) for a protected class under ECOA. This type of analysis is also very time consuming and labor intensive—which can hinder the adoption of discrete proposals for using AI/ML.

Question 15: The Equal Credit Opportunity Act (ECOA), which is implemented by Regulation B, requires creditors to notify an applicant of the principal reasons for taking adverse action for credit or to provide an applicant a disclosure of the right to request those reasons. What approaches can be used to identify the reasons for taking adverse action on a credit application, when AI is employed? Does Regulation B provide sufficient clarity for the statement of reasons for adverse action when AI is used? If not, please describe in detail any opportunities for clarity.

As a preliminary comment, although AI presents a new way of evaluating credit eligibility, the fair lending risks presented by the use of AI are not necessarily greater than other fair lending risks that FIs adeptly navigate. The guiding principles outlined in current regulatory guidance on fair lending is equally applicable to the use of AI in credit-decisioning. Current regulatory guidance, including the CFPB ECOA Baseline Review, expects regulated entities to address the use of models in the fair lending context, including the use of third-party models. Any additional guidance must continue to be flexible yet clear to allow regulated entities the ability to effectively manage the differing, but not necessarily increased risk, presented by the use of AI in credit-decisioning.

As such, AFSA members believe there is sufficient clarity on the approaches that can be used to identify reasons for taking adverse action when AI is employed. Specifically, sufficient regulatory flexibility exists within the ECOA and Regulation B to engage in the use of ML-based credit decisioning through risk-based business determinations without running afoul of the Act or its implementing regulations. AFSA appreciates the "CFPB Innovation Spotlight on Providing Adverse Action Notices When Using AI/ML Modes," and the approach articulated therein of encouraging the use of the Trial Disclosure Program Policy and the Compliance Assistance Sandbox Policy to give the CFPB a better understanding of the marketplace for future policy making and/or rulemaking. Encouraging the use of programs must continue in these early stages of AI use in credit-decisioning so that any future policy or rulemaking is properly informed by sufficient experience and data.

Question 16: To the extent not already discussed, please identify any additional uses of AI by financial institutions and any risk management challenges or other factors that may impede adoption and use of AI.

No additional comments.

Question 17: To the extent not already discussed, please identify any benefits or risks to financial institutions' customers or prospective customers from the use of AI by those financial institutions. Please provide any suggestions on how to maximize benefits or address any identified risks.

No additional comments.

* * *

AFSA appreciates the outreach the Agencies' have undertaken in issuing this RFI. Seeking input from stakeholders on this important topic is critical. We look forward to continuing to work with the Agencies on this issue. Please contact me by phone, 202-776-7300, or email, cwinslow@afsamail.org, with any questions.

Sincerely,

Celia Winslow

Celia Winslow
Senior Vice President
American Financial Services Association

919 18th Street NW, Washington, DC, 20006 ▲ 202.296.5544 ▲ www.afsaonline.org ▲ @AFSA_DC