Consumer
Federation
of America

August 12th, 2024

Deputy Assistant Secretary for Financial Institutions Policy Jeanette Quick
Director Moses Kim
Senior Policy Advisor Liang Jensen
United States Department of the Treasury
1500 Pennsylvania Avenue, NW
Washington, DC 20220

RE: Treas-DO-2024-0011-0001 Uses, Opportunities, and Risks of Artificial Intelligence in the Financial Services Sector

Dear Deputy Assistant Secretary Quick, Director Kim, and Senior Policy Advisor Jensen:

Thank you for the opportunity to comment on how the Department of the Treasury can consider opportunities for artificial intelligence (AI) to inform and shape the national financial inclusion strategy.

The Treasury Department must recognize the risks posed by artificial intelligence to consumers of financial services. Equitable access to financial services is a critical component of a fair economy and essential to ensuring that everyone can buy a home or finance a business. Unfortunately, there is still work to realize an economy where all people have full access to fair and transparent credit. For these reasons, Congress has given regulators the responsibility to implement protections from unfair, discriminatory, and dangerous practices. Because these laws and regulations predate the emergence of AI, regulators must adapt them to ensure the same legal structures remain viable and effective.

It is imperative that policymakers prioritize human rights above the gains of private industry, insist on the right to test products before they go to market, and assert the authority of existing laws and regulations.

Financial service providers may claim that consumer protections threaten their business interests. They will construct narratives of fundamental tradeoffs between accuracy and fairness or accuracy and explainability. Policymakers must ensure that when there are conflicts, the rights of consumers take priority over the interests of algorithmic efficiency and profit.

In the European Union, legislators crafted a regime whose approach to identifying risk was defined by AI's impacts on people. In the US, we must consider the risks of AI to humans above other factors. Risks to innovation, profits, or national security may be important to consider, but they should still be secondary to implications for risk to humans.

As a foundational principle, policymakers must apply the principle of precaution to the introduction of new AI-driven technologies. In other areas where a product's failure can lead to substantial consumer harm—pharmaceuticals are one example—policymakers always take steps to prevent risks before a product is approved for the marketplace. Similarly, when the decisions made by an algorithmic system have effects on important aspects of life, they must undergo review and testing.

Unfortunately, the introduction of generative AI technologies has, to date, been of the opposite course. OpenAI released ChatGPT and subsequent updates to its protocols without any testing. When considered more fully, the fact that a technology of this significance was introduced without regulatory review is shocking. We are concerned that policymakers may succumb to industry pressure from Big Tech to permit the deployment of untested algorithms. The threat of such a possibility must be given serious consideration. We cannot permit the temptation to support innovation if it means compromising the safety of consumers.

To this end, policymakers should determine when public institutions can build infrastructure to support critical AI testing and assert their views on how to test AI. In financial services, regulators should provide guidance on testing for discrimination and correct models with disparate impacts. The availability of data sets for model training could serve markets. Policymakers must be proactive in asserting their authority over evolving markets. If they are too passive, regulators create vulnerable openings for legacy institutions to use their market power to create de facto standards. Proactive rule-setting, guidance issuing, and other steps are preferable to permitting private firms to deploy untested models in regulatory "sandboxes."

One tension point centers on applying existing authorities to these new technologies. To date, policymakers in the US have proceeded with the view that risk can be mitigated without significant changes to laws and regulations. This contrasts with the European Union's (EU) response, which has been to pass legislation that recasts the protection of consumers to fit with the dynamics of new technologies.

Speed is important. From a pragmatic point of view, developing new laws that create new agencies and legal protections may be slow. Given this constraint, we believe policymakers should use their existing regulatory authorities.

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)? How has AI been used in financial services to improve fair lending and consumer protection, including substantiating information? To what extent does AI improve the ability of financial institutions to comply with fair lending or other consumer protection laws and regulations? Please be as specific as possible, including details about cost savings, increased customer reach, expanded access to financial services, time horizon of savings, or other benefits after deploying AI.

i. *Mature AI technologies already exist for fraud prevention, but the ecosystem is undermined by the lack of fraud information sharing.*

Artificial intelligence has already proven useful in preventing fraud in financial services. As the Treasury Department notes in the introduction of its 2024 report on cybersecurity risk, financial institutions have deployed AI to defend their systems from hackers and prevent scammers from perpetrating fraud inside their networks.[1]

[1] US Department of the Treasury. "Managing Artificial Intelligence-Specific Cybersecurity Risks in the Financial Services Sector." AI Report, March 2024. https://home.treasury.gov/system/files/136/Managing-Artificial-Intelligence-Specific-Cybersecurity-Risks-In-The-Financial-Services-Sector.pdf.

However, financial institutions could do better. One problem is that too few of them are working together. Their reticence to share information undermines fraud prevention. A 2024 survey of 158 financial institutions revealed that more than three in five believed consortium data was ‘critically’ or ‘very’ important, but only 16 percent said their firms participated in such an effort.[2]

In April 2024, the Treasury Department published a report criticizing banks for their reticence to share information to prevent fraud, even as it acknowledged that banks do share information to thwart cyberattacks and money-laundering networks. The report questioned why a clearing house to allow rapid sharing of fraud data does not already exist. It noted that some institutions build models for internal use, but that advanced AI-driven modeling was practiced at larger banks.[3]

Working to ensure that financial institutions share fraud information and make use of the information they receive should be a goal of financial regulators independent of the role played by the use of AI.

ii. *Fraud information sharing is a widespread practice in other countries.*

Internationally, regulators from advanced and developing countries require financial institutions to participate in fraud-information-sharing regimes. For example:

- The EU’s Payment Systems Directive 2 (PSD2) requires all financial institutions to report fraudulent financial activity.
- In the Netherlands, five Dutch Banks partner through the Transaction Monitoring Netherlands initiative to jointly scan payment traffic for unusual patterns. By working together, they prevent scams that banks would not be able to thwart if acting alone.[4]
- The Society for Worldwide Interbank Financial Telecommunication’s (Swift) Payment Controls Service (PCS), a cooperative that facilitates common payment messaging formats for international funds transfers, uses AI to detect transactions with a high risk of fraud using historical patterns on its payment messaging networks.[5] PCS can identify unusual payments to create alerts or permit transaction blocks in real-time.
- As a condition of participating in the Central Bank of Brazil’s (BCB) Pix payment system, banks and non-banks must submit reports of fraudulent financial activity to the Transaction Accounts Identifier Directory (the “DICT”). In turn, the BCB shares this information with all other Pix participants. In turn, those participants are required to use it.

As Pix is a real-time payment system, efforts to stop fraud require automated systems to effectively surveil against potential fraud. One interesting observation made by the BCB was that most scams

[2] Pape, Carter. “Call Centers and Bank Branches Are Major Fraud Liabilities.” *American Banker*, July 16, 2024. https://www.americanbanker.com/news/call-centers-and-bank-branches-are-major-fraud-liabilities.

[3] US Department of the Treasury. “Managing Artificial Intelligence-Specific Cybersecurity Risks in the Financial Services Sector.” AI Report, March 2024. https://home.treasury.gov/system/files/136/Managing-Artificial-Intelligence-Specific-Cybersecurity-Risks-In-The-Financial-Services-Sector.pdf.

[4] Hopman, Paul, Aline Kiers, and Rik van Weenen. “PSD3 and PSR: Sharing Data on Fraudulent Payment Transactions.” *DLA Piper* (blog). Accessed August 5, 2024. https://www.dlapiper.com/en/insights/publications/2023/08/psd3-and-psr-sharing-data-on-fraudulent-payment-transactions.

[5] Zschach, Tom. “Harnessing AI in the Fight against Payments Fraud.” SWIFT, May 30, 2024. https://www.swift.com/news-events/news/harnessing-ai-fight-against-payments-fraud.

occurred using accounts open for at least a year. Had the BCB relied only on AML screens at the moment of account opening, they would have had little or no ability to identify scams[6] on aged accounts.[7]

iii. AI technologies can protect consumers from payment scams. The ability to review account activity, as opposed to relying solely on identity verification during the account opening process, is an important step to improving fraud prevention.

Financial institutions and their regulators must better understand the ultimate destination of funds to identify scammers.

AI-enhanced fraud detection systems can identify high-risk payments using pattern analysis. These capabilities strengthen anti-fraud efforts that would otherwise consist of verifying identities when approving new accounts. These systems are useful for identifying account takeover, spoofing, or sender-authorized fraud.

Regulators should insist financial institutions implement multi-factor authentication when accounts are accessed through devices. At the moment, regulatory guidance does not mandate multi-factor authentication. It requires institutions to use "stronger" authentication when transactions are high-risk, but it stops short of making MFA a standard. Instead, MFA is only a best practice.[8] Financial regulators should determine why some institutions are not implementing MFA and consider how they could encourage and support those institutions to introduce MFA.

iv. Treasury and other regulators should ensure that AI used to thwart fraud does not prevent consumers from having fair access to financial services.

Today, many consumers are denied access to bank accounts because their applications trigger know-your-customer and anti-money-laundering flags. The presence of more data sharpens decision-making, and conversely, a paucity of documentation increases uncertainty. Moreover, as systems become more automated, there are fewer opportunities for a human to intervene to reject a faulty model. One of the main reasons consumers do not have a bank account is that they do not have enough documentation to fulfill KYC requirements. Outcomes such as this are a problem. When a well-intentioned consumer's application is rejected because they cannot muster data required by a machine, the system has failed. Financial services must meet the needs of all consumers. Unfortunately, this is the status quo. With greater reliance on automated systems, the problem could worsen.

AI has an opportunity to rectify the false-positive problem. For example, it is possible that models may be able to incorporate PII information from government agencies in foreign countries. That is an optimistic outcome. It is also a vision that is outside the providence of regulators. Regulators should focus on the downsides. This is the essence of a policy approach that prioritizes risks to people above the motives for industry to innovate.

[6] As terms of art in the financial services industry, "frauds" refer to funds lost by unauthorized transactions and "scams" characters those resulting from authorized ones.

[7] Carlos Eduardo Brandt. "Brazilian Payments Innovation." Presented at the Brookings Institution Center on Regulation and Markets, The Brookings Institution, September 7, 2023.

[8] "Authentication in Internet Banking: A Lesson in Risk Management – Winter 2007 Vol. 4, Issue 2, Updated July 10, 2023." Supervisory Insights. Federal Deposit Insurance Corporation, July 10, 2023. https://www.fdic.gov/regulations/examinations/supervisory/insights/siwin07/siwinter2007-article05.html.

To this end, regulators must ensure that when financial institutions use AI to review applications for transaction accounts (and credit), they have a person "in the loop" to review denials for account applications.

Additionally, we are concerned that some financial institutions will delay or object to changing account closure practices that are unfair and discriminatory on the grounds of fraud prevention. There is no basis to say that protected class members are more likely to engage in fraudulent activities.

Actual and Potential Risks and Risk Management Oversight of AI—Explainability and Bias
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. 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. What challenges exist for addressing risks related to AI explainability? What methodologies are being deployed to enhance explainability and protect against potential bias risk?

A. *Lenders must understand how to interpret their models and explain the decisions made by those models to consumers.*

i. *As a separate condition of meeting the goal of explainability in modeling, supervisors should hold lenders accountable for demonstrating their ability to interpret their models.*

Supervisors should work to prevent harm by verifying that banks understand the content of outsourced models.

Many financial institutions rely on third-party non-bank service providers to build and test algorithmic models. In some cases, the choice of their third-party modeler may be made by a banking-as-a-service provider that is also a non-bank and separate from the insured depository. The interagency third-party risk management guidance established by prudential regulators clearly states that banks are responsible for any risks arising from these vendors. Nonetheless, this expectation requires persistent supervision as algorithmic models and bank staffing change.

ii. *Adverse-action notices must be redesigned to fit the complexities of algorithmic underwriting.*

The ECOA requires creditors to explain the reasons for an adverse credit decision. This expectation exists if an application is denied or under terms less desirable than originally sought. The CFPB has authority over the ECOA. The CFPB can use its existing authority to hold financial institutions accountable to explain adverse decisions resulting from algorithmic underwriting. No new regulatory authority is necessary.

The CFPB has been clear that "black boxes" cannot become "brick walls."[9] For example, the CFPB has signaled that it expects lenders to provide clear, accurate, and specific reasons for adverse credit decisions. It published a circular (2022-03) on adverse action requirements for explaining adverse credit

[9] Chopra, Rohit. "Director Chopra's Prepared Remarks on the Interagency Enforcement Policy Statement on 'Artificial Intelligence.'" *Consumer Financial Protection Bureau* (blog), April 25, 2023. https://www.consumerfinance.gov/about-us/newsroom/director-chopra-prepared-remarks-on-interagency-enforcement-policy-statement-artificial-intelligence/.

decisions derived from algorithmic underwriting in May 2022[10] and another (2023-03) requiring creditors to provide accurate and specific reasons for denials in September 2023.[11]

The use of AI has been an element of the CFPB's work in other areas as well. The CFPB published an interpretive rule on the use of AI by digital marketers in August 2022,[12] and a spotlight on the use of AI chatbots in June 2023.[13] In an April 2023 guidance, the CFPB provided an analytical framework for how it would address abusive conduct and noted that its prohibitions could cover service providers who use an algorithm in ways that are abusive, including if the algorithm is used in marketing.[14] In June 2023, the CFPB proposed a rule for the use of automated valuation models to address the use of algorithms in real estate valuation assessments (but not appraisals).[15] Lastly, the CFPB participated in an interagency statement outlining a commitment to enforce its authority to prevent discrimination and enforce existing consumer financial protection laws when creditors use automated systems.[16]

Steps must be taken to modernize adverse action notices to realize Congress's intent of explainability when it passed the ECOA.

The complexity of some AI models may exceed the capabilities embedded inside current adverse-action forms. These forms have a small list of potential reason codes and were designed when the most advanced decision-making techniques incorporated linear regression. As structured, they cannot account for the volume of data points and the dynamic nature of modeling systems.

In addition to making decisions explainable to consumers, the CFPB must verify through its supervisory reviews that financial institutions can interpret algorithmic models. Unless they can interpret models, financial institutions cannot fulfill requirements for managing risk when contracting with third-party service providers.

[10] "Consumer Financial Protection Circular 2022-03: Adverse Action Notification Requirements in Connection with Credit Decisions Based on Complex Algorithms." CFPB, May 26, 2022. https://www.consumerfinance.gov/compliance/circulars/circular-2022-03-adverse-action-notification-requirements-in-connection-with-credit-decisions-based-on-complex-algorithms/.

[11] "Consumer Financial Protection Circular 2023-03: Adverse Action Notification Requirements and the Proper Use of the CFPB's Sample Forms Provided in Regulation B." CFPB, September 19, 2023. https://www.consumerfinance.gov/compliance/circulars/circular-2023-03-adverse-action-notification-requirements-and-the-proper-use-of-the-cfpbs-sample-forms-provided-in-regulation-b/.

[12] Interpretative Rule: Limited Applicability of Consumer Financial Protection Act's "Time or Space" Exception with Respect to Digital Marketing Providers. CFPB, August 10, 2022.https://files.consumerfinance.gov/f/documents/cfpb_time-or-space_interpretive-rule_signed_2022-08.pdf.

[13] Issue Spotlight: Chatbots in consumer finance. CFPB, June 6, 2023. https://www.consumerfinance.gov/data-research/research-reports/chatbots-in-consumer-finance/chatbots-in-consumer-finance/.

[14] "CFPB Issues Guidance to Address Abusive Conduct in Consumer Financial Markets." CFPB, April 3, 2023. https://www.consumerfinance.gov/about-us/newsroom/cfpb-issues-guidance-to-address-abusive-conduct-in-consumer-financial-markets/.

[15] "Proposed Rule: Quality Control Standards for Automated Valuation Models." CFPB, June 1, 2023. https://www.consumerfinance.gov/rules-policy/rules-under-development/quality-control-standards-for-automated-valuation-models/.

[16] "CFPB and Federal Partners Confirm Automated Systems and Advanced Technology Not an Excuse for Lawbreaking Behavior." CFPB, April 25, 2023. https://www.consumerfinance.gov/about-us/newsroom/cfpb-federal-partners-confirm-automated-systems-advanced-technology-not-an-excuse-for-lawbreaking-behavior/.

Financial regulators must not accept an argument from lenders that obligations for explainability and interpretability are at loggerheads with accuracy. Research shows that explainability does not compromise the efficacy of underwriting.[17]

B. *To exercise caution, it is essential that lenders use unbiased data sets. Efforts to test for bias are undermined by a lack of demographic data. Lenders can only ask for demographic data when offering credit for mortgages and small business financing. In other lending segments, regulations prohibit lenders from requesting demographic data about applicants.*

In exercising its authority as the regulator responsible for the Equal Credit Opportunity Act (ECOA) and Regulation B, the CFPB should clarify how lenders can solicit sensitive demographic data for testing model fairness.

i. *Regulations surrounding the solicitation of demographic data are inconsistent. While there are valid concerns about discouragement, data collection has demonstrated benefits for enforcing fair lending laws.*

Regulators have prevented lenders from soliciting demographic information from credit applicants outside mortgage lending. Pending legal resolution, small business lenders will be required to ask for demographic data, as well.

When it finalized Regulation B to implement rules for ECOA,[18] the Federal Reserve Board (FRB) prohibited lenders from asking or documenting protected class characteristics of non-mortgage credit applicants. It noted that solicitations could discourage applicants. In 1999, the FRB considered an amendment to remove the prohibition, but after consideration and comment, it left it intact. It also raised the possibility that the data could be used for new discriminatory practices and that there would be variation in how information was solicited.[19]

There is a strong record showing the benefits of data transparency to support fair lending efforts. Since the passage of the Home Mortgage Disclosure Act, mortgage lenders have been required to collect and report demographic data. Many stakeholders have used HMDA data. It has been an important tool for fair lending enforcement.[20]

Over time, a variety of permissible alternatives have developed. For example, a now out-of-date interagency fair lending manual permitted lenders to assume that an applicant is Hispanic based on the last name, or in the case of female applicants, on the first name, or to consider an applicant to be African-

[17] Rudin, Cynthia and Radin, Joanna. "Why Are We Using Black Box Models in AI When We Don't Need To? A Lesson from An Explainable AI Competition." *Harvard Data Science Review* 1, no. 2 (2019). https://doi.org/10.1162/99608f92.5a8a3a3d.

[18] In the Dodd-Frank Act, Congress transferred authority for ECOA to the CFPB from the Federal Reserve.

[19] Williams, Orice M. "Testimony Before the Subcommittee on Oversight and Investigations, Committee on Financial Services. House of Representatives: Race and Gender Data Are Limited for Nonmortgage Lending." Government Accountability Office, July 17, 2008. https://www.gao.gov/assets/gao-08-1023t.pdf.

[20] Ficklin, Patrice, Tim Lambert, and Abby Hogan. "Fair Lending at the CFPB and the Role of the Home Mortgage Disclosure Act (HMDA) in Protecting America's Consumers." Presented at the Consumer Financial Protection Bureau Office of Fair Lending Consumer Protection Week, July 2020. https://files.consumerfinance.gov/f/documents/cfpb_hmda-data-browser_cfpw-presentation_2020-07.pdf.

American based on their census tract.[21] Obviously, these approaches are unreliable and do not provide a sound foundation on which to undertake ongoing, robust testing for fairness.

ii. *Bayesian Improved Surname Geocoding (BISG) is an imperfect statistical methodology.*

Currently, algorithmic model testing for credit outside of mortgage lending relies on BISG methodologies to assess the demographic composition of applicant pools. Some fintech service providers offer customized versions of BISG, but even these "BISG+" tools are constrained.

When model reviews incorporate BISG techniques into fairness assessments, they introduce challenges that make testing unnecessarily difficult. While they are too blunt and simplistic to be useful in inferring the demographic data of individuals, BISG techniques can be applied to make inferences about the demographic makeup of groups.

The CFPB conducted a study on BISG in 2014, acknowledging that information on consumer race and ethnicity was required for fair lending testing, and non-mortgage credit products could not solicit the information.[22] Since then, the CFPB has completed its rulemaking to require lenders to report demographic data for small business lending.

CFPB research on using BISG with mortgage lending data yielded three important findings. First, BISG techniques improved predictive power compared to systems that relied only on zip codes or surnames on their own. Second, estimates of the demographic makeup of groups still differed from the reported classifications revealed in HMDA, and those errors had a concerning dynamic: BISG tended to overestimate the share of protected class members and undercount the number of white non-Hispanic mortgage applicants. The CFPB also opined the possibility of using BISG for individuals if BISG probabilities were high enough to cross an 80 percent estimate threshold. The Bureau noted that it would produce some false positives for estimates above the threshold but could artificially bias estimates of the number of protected class members downward for estimates below the threshold.[23] The last two findings show that all BISG methodologies force testers to accept statistical compromises.

iii. If lenders are going to implement testing of disparate impact, they will require access to demographic data.

Lenders must have data to test their models to ensure algorithmic lending does not disparately impact credit access for protected class members. As mentioned earlier, they can use BISG, but it is not optimal.

As a condition, the CFPB must convey acceptable mechanisms for soliciting data. Compliance staff may insist that requests for data be made with legally approved forms. Testers would benefit from clarity on acceptable sample sizes would be valuable.

[21] "Interagency Fair Lending Examination Procedures." Office of the Comptroller of the Currency, Federal Deposit Insurance Corporation, Federal Reserve Board, Office of Thrift Supervision, and National Credit Union Administration, August 2009. https://www.ffiec.gov/pdf/fairlend.pdf.

[22] Consumer Financial Protection Bureau. "Using Publicly Available Information to Proxy for Unidentified Race and Ethnicity: A Methodology and Assessment." Washington, DC, September 2014. https://files.consumerfinance.gov/f/201409_cfpb_report_proxy-methodology.pdf.

[23] "Using Publicly Available Information to Proxy for Unidentified Race and Ethnicity: A Methodology and Assessment." CFPB, September 2014. https://files.consumerfinance.gov/f/201409_cfpb_report_proxy-methodology.pdf.

The CFPB should determine how data can be solicited for testing without causing discouragement. Additionally, financial regulators should insist that lenders should are prevented from sharing solicited information with affiliates and unaffiliated third parties or using the data for marketing. Specific rules will provide support for internal champions at financial institutions to press forward against "fairness through unawareness" perspectives.

We urge the CFPB to establish clearer rules on collecting and using demographic data for testing for disparate impact. These instructions will advance the infrastructure needed to test models for disparate impact and will support efforts to prevent harm. While financial institutions will need such data to test for disparate impact effectively, strong safeguards should be implemented to ensure that such data is collected and used appropriately.

iv. Prudential regulators or the CFPB should consider creating a model data set to train algorithmic models for providing access to financial services.

If an algorithmic model has been trained on a data set whose contents are compromised with bias, it will learn to replicate those errors.

Modelers are free to use data sets for testing, but some may choose to begin testing after model deployment. Such an approach will inevitably lead to consumer harm. Addressing harm with enforcement is important, but it is suboptimal to preventing the harm in the first place.

To support precautions, financial regulators could improve the market by creating data sets for training. By doing so, they would prevent harm from untested models. Moreover, they would have control over the possibility that models will be trained from biased information.

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

Question 10: How are financial institutions addressing any increase in fair lending and other consumer-related risks, including identifying and addressing possible discrimination related to the use of AI, particularly emerging AI technologies? What governance approaches throughout the development, validation, implementation, and deployment phases do financial institutions expect to establish to ensure compliance with fair lending and other consumer-related laws for AI models and tools prior to deployment and application? In what ways could existing fair lending requirements be strengthened or expanded to include fair access to other financial services outside of lending, such as access to bank accounts, given the rapid development of emerging AI technologies? How are consumer protection requirements outside of fair lending, such as prohibitions on unfair, deceptive, and abusive acts and practices, considered during the development and use of AI? How are related risks expected to be mitigated by financial institutions using AI?

Some financial institutions already use AI for underwriting, and many are interested in using it in the future. While these advancements have the potential to enhance efficiency and advance financial inclusion, there is growing evidence that they can also perpetuate and exacerbate existing and historical biases, leading to discriminatory outcomes that adversely affect marginalized and underserved communities.

The remarks of Acting Comptroller Michael Hsu summarize the quandary faced by policymakers on the impacts of AI to the fairness of our financial system:

> "From a fairness perspective, AI holds both promise and peril. AI has the potential to

> reduce bias and enable fair access to credit and banking services in ways that humans have been challenged to do. We should welcome this and create spaces to safely explore and develop this potential. At the same time, however, AI has the potential to perpetuate and exacerbate the biases, discrimination, and unfairness that are deeply embedded in the data feeding AI systems. The challenge is that there can be strong echoes of race, gender, and other characteristics in large data sets, even when race data are removed. This means that even if a data set were "color blind," the redundant encodings in the other variables would likely reflect and reveal race as a factor."[24]

The tension underscores the need for action – before harm occurs. It is entirely possible that absent action lenders whose understanding of their models is limited will use unfair models. They may do some without being aware of the impact – the "fairness through unawareness problem – and neither test for disparate impacts nor even conduct searches to verify that better alternatives exist. The legacy stances of many lenders have been to assert that their use of FICO is implicitly fair, even if evidence shows it is not, and reify that approach. They may do so even if fintech service providers disagree. Lenders and non-bank third-party service providers will inevitably push back on any accommodation for fairness if they believe it comes at the expense of accuracy.

As if to illustrate this point, a monitorship between an algorithmic modeling company and a fair lending firm ended due to a dispute on when to search for a lesser discriminatory alternative. The conflict centered on reducing a model's accuracy to support a fairer version. The modeler insisted it would undermine its competitive position in the market if it made any concession on accuracy.[25] This tension underscores the need for regulator intervention to ensure that anti-discrimination consumer protection laws are followed in this new era of underwriting.

i. There is a need for regulatory clarity on how lenders can search for less discriminatory algorithms (LDAs).

When financial institutions fail to search for and implement LDAs, they undermine consumer trust and contravene principles of fairness and equity that are foundational to our financial system and broader society. In some instances, these failings may rise to the level of violating the law.

The CFPB has authority for the Equal Credit Opportunity Act (ECOA) and Regulation B. To clarify that algorithmic underwriting is held to the same anti-discrimination standards as traditional methods, the CFPB has repeatedly expressed its conviction that fair lending laws and regulations apply regardless of how credit is offered.

The CFPB should respond by providing clear guidance on how lenders should search for and implement LDAs when using algorithms for credit underwriting and pricing, which is currently lacking. Clear guidance is needed to complement supervisory and enforcement action, which by themselves are insufficient to provide necessary clarity for the market.

[24] Hsu, Michael. "Remarks at the National Community Reinvestment Coalition Just Economy Conference 2024 'Elevating Fairness, 2024.'" Washington, D.C., April 4, 2024. https://www.occ.gov/news-issuances/speeches/2024/pub-speech-2024-38.pdf.

[25] Relman Colfax. "Fair Lending Monitorship of Upstart Network's Lending Model." Monitorship, March 27, 2024. https://www.relmanlaw.com/cases-406.

While enforcement has a positive effect on safeguarding the market, it occurs after harm has taken place. Regulatory clarity, combined with guidance on supervisory expectations, will make both supervision and enforcement more effective and can prevent future harm.

Providing carefully crafted guidance will also counter claims of innocence from financial institutions that do not conduct robust searches for LDAs while aiding those who have good intentions but are unaware of proper testing techniques. Guidance and examples of how to effectively search for and develop LDAs and demonstrating where institutions have already fulfilled the expectations of the CFPB will address both scenarios.

ii. Clarity on how to apply fair lending anti-discrimination laws to algorithmic underwriting will support competition.

Clear, explicit guidance will also encourage more financial institutions to use these technologies in ways that lead to more competition in the market. Some lenders are introducing ML models, either of their own making or through a relationship with a vendor, but many remain on the sidelines due to a lack of regulatory clarity. We are concerned that too many "disruptors" will forge ahead with dangerous models, while conversely, many small and medium-sized financial institutions will fall behind, missing the opportunity to leverage ML models to reach underserved consumers.

The risk of harm to consumers due to prolonged resistance to LDA searches calls for moving beyond the status quo. If financial institutions do not see irrefutable evidence of expectations to perform statistically effective reviews, many will fail to do so. Nonetheless, given the current political and legislative environment, we understand the need to proceed with prudence, particularly with rapidly evolving technology. We suggest a set of practical options that can be implemented within the current regulatory framework to help make significant progress in addressing these issues.

Question 12: How are financial institutions, technology companies, or third-party service providers addressing and mitigating potential fraud risks caused by AI technologies? AND Question 13: How do financial institutions, technology companies, or third-party service providers expect to use AI to address and mitigate illicit finance risks? What challenges do organizations face in adopting AI to counter illicit finance risks? How do financial institutions use AI to comply with applicable AML/CFT requirements? What risks may such uses create? AND *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 they 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?*

Third-Party Risks

The financial services sector has been one of the early adopters of AI technology. As mentioned earlier, it has been used as a tool for preventing fraud and money laundering for some time. Increasingly, though, non-bank fintechs are building lending models.

Financial institutions using AI are largely limited to select groups: (1) large financial institutions with resources to build models and conduct testing on their own, (2) bank and non-bank disruptors who will likely proceed forward without due regard for the potential of harm to consumers, (3) banks who conduct partnerships with fintechs, and (4) tech firms that contract with depositories to provide proprietary models. At the same time, many institutions are making little use of AI. Those not participating include many of the smaller community-focused banks and credit unions that often focus on reaching underserved consumers, as well as other depository institutions that feel hesitant to invest resources in technologies where compliance risk remains uncertain.

i. Important players in third-party relationships are not supervised by banking regulators.

Recently, a crisis exposed shortcomings with independent banking-as-a-service (BaaS) platforms. The event follows a multi-year period where prudential regulators have issued enforcement actions against the most leading BaaS partner banks. The event could result in tens of thousands of depositors being unable to recoup their insured deposits.

The event made it painfully clear how there is a dangerous gap in supervisory authority. The independent middleware provider may have served depositories supervised by banking regulators but was itself not subject to any supervision. The failure of the independent company did not cause its banking partners to fail, and as a result, depositors were not protected by deposit insurance.

It also highlighted a gap in third-party risk management guidance and actual supervisory implementation. For example, one of the reasons regulators cannot verify deposit holdings is because a vendor has not been paid for services. Pending payment, the vendor may or may not continue to grant access to its database. The prudential regulators contemplated this possible outcome when they crafted the latest third-party guidance. In the document, they called for banks to ensure funds are escrowed to pay vendors of the third-party partner. Nonetheless, an escrow fund did not exist to satisfy the debt.

Third-party relationships that involve the provision of consumer-facing products must be supervised by prudential regulators and the CFPB.

i. Competition among third-party BaaS providers should not prioritize partner banks' costs above consumer protections.

Using third-party vendors permits banks to seek cost efficiencies by outsourcing activities outside their core competencies. In isolation, such an arrangement leads to better use of capital. However, in practice, it may undermine accountability.

Unless banks understand how AI works, they may not be prepared to select vendors to perform AI-driven services. The problem is an upstream version of the classic issue of black box underwriting models whose decision-making cannot be explained to credit applicants. In this case, the shortcoming is a lack of understanding among bankers of an algorithm's predictive power and compliance with consumer protection regulations. Absent the understanding that could only be attained by someone with expertise, procurement decisions will be made based on price or salesmanship.

The weight given to vendor compliance with consumer protection laws would grow if there was some understanding of which firms and methods comply with Regulation B and other anti-discrimination regulations. Most likely, any third-party firm that could not prove its compliance with these regulations would be challenged to attract bank partners. In the last two years, prudential regulators have clamped down on BaaS by issuing a series of enforcement actions. These efforts are improving the market, underscoring how regulation prevents a race-to-the-bottom and corrects markets toward pro-consumer competition.

ii. If financial regulators are slow to provide regulatory clarity will lead to re-entrenchment by legacy credit modeling firms.

If regulators cannot clarify how AI can be deployed, there is a risk that dominant credit analytics firms will, because of their market power, become standard setters for the use of AI. At least two factors contribute to this likelihood.

First, given their scarce financial resources, many smaller depositories will refrain from using unproven AI tools without assurance that they will not take on compliance risk. Their relative lack of financial resources is only one driver behind this scenario. An equally determinative factor could be understanding that banks are accountable for the mistakes of their vendors.

Smaller depositories are unlikely to have the staff to judge the efficacy of AI vendors. They know that third-party risk-management guidance will hold them accountable for their vendors—and rightfully so. Unlike the case with more standard procurement decisions, choosing an algorithmic modeler and tester requires an in-house understanding of how algorithms work. It is easy to observe diversity, equity, and inclusion goals when contracting with vendors for event management, advertising, graphic design, or other straightforward business lines. It is far more difficult to do so when assessing how an AI vendor constructs its neural networks.

Second, demonstrated patterns and practices exist to show how standardization facilitates highly interdependent networks in financial services. FICO scores have become an important input for the securitization of mortgage-backed securities, and the grades of investment-grade bonds made by rating agencies have become embedded in standards for investor protection. Even as recently as 2022, the GSEs renewed their emphasis on the primacy of FICO when they gave it the status as default credit scoring system for mortgages.[26] Already, calls are being made by industry[27] and advocates[28] for policymakers at the Federal Housing Finance Administration to determine standards for government-sponsored entities that purchase mortgages underwritten with cash-flow underwriting.

In that vacuum, it is entirely likely that one or both leading credit analytics firms will purchase the small fintech modelers who currently supply most of the AI-driven algorithms. This would supplant the reliance on unknown and nascent algorithmic modeling fintech with new and potentially greater problems. Lenders might seek to apply political pressure on financial regulators to provide regulatory accommodation. Existing market power in non-algorithmic lending would lead to dominance in algorithmic modeling. The anti-competitive practices already observed would be extended.

These risks underscore the need for regulators to act expeditiously to ensure consumers are protected from harm associated with using algorithms in underwriting.

iii. While prudential regulators need supervisory authority over middleware providers, the answer is not to create a "fintech charter."

[26] Fannie Mae, and Freddie Mac. "Joint Enterprise Credit Score Solicitation," June 2022. https://singlefamily.fanniemae.com/media/22061/display.

[27] National Association of Realtors. Comment Letter. "Letter to FHFA Director Sandra Thompson on Fintech Request for Information." Comment Letter, October 31, 2022. https://www.fhfa.gov/sites/default/files/discussion_topics/Attachments/2045/2022.10.30%20NAR%20Letter%20to%20FHFA%20Re%20Fintech%20RFI_Final.pdf.

[28] National Consumer Law Center. "Comments/RIN 2590-AA98, Validation and Approval of Credit Score Models," March 21, 2019. https://www.fhfa.gov/sites/default/files/rulemaking_comments/Attachments/15401/Consumer%20group%20comment%20to%20FHFA%20re%20credit%20scoring%20models%20with%20sign%20ons.pdf.

Regulators should not revisit the question of creating a fintech charter. The Office of the Comptroller of the Currency issued a proposal to create a special charter for non-depositories. The possibility of such a structure raised many questions. It is important that the problems arising from the recent middleware BaaS event do not become grounds to support the return of such an idea.

Fintech charters may provide a basis for supervision, but they extend many benefits on non-depositories. For example, once granted a fintech charter, a non-depository could gain direct access to Federal Reserve payment services. It would open a loophole that gave entities the benefit of a charter without a meaningful community reinvestment obligation.

iv. Sandboxes are also inappropriate and will put consumers at risk of harm from untested models.

The use of regulatory sandboxes as a tool for cultivating advancements in financial services is a mistaken approach. Some proponents of sandboxes have used the term “unleash” to describe the benefits of sandboxes. Oddly, this term also describes its fundamental risks. AI is fundamentally a black box. Regulators must insist that lenders and third-party service providers test models before algorithms are unleashed from black boxes.

Conclusion

Thank you for the opportunity to provide input on your consideration of how the Treasury Department could advance the use of artificial intelligence.

All modelers must be held accountable for testing their algorithms before releasing them to the public. Biases in automated systems can result in incorrect or inaccurate decisions, leading to real harm to consumers due to financial exclusion or unfair pricing. The potential for automated systems to produce biased outcomes and “automate discrimination” is significant. Some financial institutions may go ahead with models before they fully understand their implications, leading to huge loopholes that undermine the public. The industry's tendency to claim “fairness through unawareness” is too real. Similarly, financial regulators must prioritize consumer protection standards over private profit motives. Already, conflicts in how fair lending rules are observed have revealed fissures along these lines.

Due to the pace of innovation, regulators should rely on existing authorities whenever possible. The rate of technological advance will outpace the response from legislators, and as a result, the responsibility to protect consumers from the risks of AI falls squarely on our agencies.

Sincerely,

Adam M. Rust

Adam Rust
Director of Financial Services
Consumer Federation of America
arust@consumerfed.org