August 9, 2024

Submitted via regulations.gov.

Re: Uses, Opportunities, and Risks of Artificial Intelligence in the Financial Services Sector

Upstart Network, Inc. ("Upstart") is the leading AI lending marketplace, connecting millions of consumers to more than 100 banks and credit unions that leverage Upstart's AI models and cloud applications to deliver superior credit products. With Upstart AI risk assessments, lenders can approve more borrowers at lower rates across races, ages, and genders, while delivering the exceptional digital-first experience customers demand. More than 80% of borrowers are approved instantly, with zero documentation to upload. Founded in 2012, Upstart's platform includes personal loans, automotive retail and refinance loans, home equity lines of credit, and simple <36% APR small-dollar loans. Upstart is based in San Mateo, California, and Columbus, Ohio.

Given our robust use of AI in financial services, Upstart appreciates the Department of the Treasury's Request for Information on Uses, Opportunities, and Risks of Artificial Intelligence in the Financial Services Sector. Below you will find Upstart's response to some of the questions posed in the request for information.

General Use of AI in Financial Services

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

The definition of AI in this RFI is generally appropriate for financial institutions, although it may be helpful to clarify what is meant by "virtual environment."

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

Upstart leverages the power of AI to more accurately quantify the true risk of making a loan. Our AI models have been continuously upgraded, trained and refined for over a decade. Our AI models use and analyze data from all of our lending partners.We have discrete AI models that target income fraud, acquisition targeting, loan stacking, prepayment prediction, and identity fraud and time-delimited default prediction. Our models incorporate more than 1,600 variables and benefit from a rapidly growing training dataset that currently contains more than 58 million repayment events, adding an average of 83,000 new repayments each business day.[1]

We have been able to demonstrate through several studies that AI-powered lending works. Each year, Upstart evaluates the ability of our personal loan model to underwrite applicants in comparison to a more "traditional" model–i.e. a benchmark model including criteria like credit score or debt-to-income ratio. We conduct the research looking back at the prior year. Our research showed that, in 2023, the Upstart model could approve more applicants, including Black and Hispanic applicants, at lower APRs than a more traditional underwriting model. In comparison to the traditional model, the Upstart model approved 101% more applicants and resulted in APRs that were 38% lower.[2] Additionally, for pools of securitized loans, our realized loss rates were only approximately half of those predicted by Kroll, a prominent credit rating agency; over that same time period, realized losses for the same pool of loans were on average only 5% different than our internal forecasts.[3]

Our AI models are provided to bank and credit union partners within a consumer-facing cloud application that streamlines the end-to-end process of originating and servicing a loan. We have built a configurable, multi-tenant cloud application designed to integrate seamlessly into a lender's existing technology systems. Our highly configurable platform allows each lending partner to define its own credit policy and determine the significant parameters of its lending program. As a result, these models are trained by every Upstart-powered loan, and each lending partner benefits from participating in a shared AI lending platform.

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

As an AI model developer that partners with state-chartered credit unions and community banks, Upstart has experienced first-hand situations where some smaller institutions feel unable to access the benefits of AI due to challenges such as uncertainties in regulatory compliance obligations and/or perceptions that they have inadequate technical expertise, or human and financial resources to fulfill those obligations.[4] Still, small banks and credit unions represent the majority of the over 100 institutions that use Upstart's technology, indicating that these challenges and perceptions are not uniform roadblocks across the system.

Upstart has found that certain best practices are key to addressing the challenges faced by small and medium-sized community institutions in deploying AI, including: (1) offering to participate transparently

[1] As of 12/31/2023, references to variables refer to all raw variables and certain combined variables considered in our AI models.
[2] https://www.upstart.com/lenders/regulatory-compliance/access-to-credit-report/
[3] In an internal study, Upstart compared the actual realized loss rates of Upstart loans securitized in five securitization transactions between June 2017 and September 2019 and the loss rate predictions for those loans obtained from KBRA Surveillance Reports published by Kroll Bond Rating Agency in December 2019. As compared to Kroll's loss predictions, actual realized losses were approximately 31% to 71% lower, with an average deviation across all five securitization transactions of -48%. As compared to our internal forecasts, actual realized losses ranged from approximately 35% higher (for the earliest securitization transaction) to approximately 17% lower (for the most recent securitization transaction), with an average absolute deviation across all of five securitization transactions of approximately 13%.
[4] Community institutions often express misgivings about whether they will be able to answer very technical model documentation questions from supervisors / regulators and express uncertainty as to whether they can rely on third party expertise to assist them with monitoring and oversight.

in discussions with their regulatory supervisors; (2) conducting regular independent statistical validations of the Upstart model that institutions can review; (3) providing access to loan data and reports that align with the regulatory examination schedule of the institution; and (4) articulating clear strategic goals early in the onboarding process for how the deployment of a third-party AI credit model will help the institution execute on its business strategy and better serve the community.

Finally, development of a well-designed and well-executed standard-setting process and voluntary certification of third-party models may provide an opportunity for regulators to be comfortable with third party risk management related to AI usage, and ease the path to adoption of sound third party models. This enables community banks to use these AI models in a safe and sound manner even when they do not have adequate resources to independently validate AI models developed by third-party technology providers themselves. Upstart believes that the 2020 FDIC proposal providing a framework that centralizes and standardizes certain model risk management and third-party relationship due diligence functions through a voluntary certification process, overseen by regulators, could over time, significantly reduce the barriers to adoption of certified models by individual community banks and smaller institutions, thereby increasing the speed of adoption of innovative technology via well-vetted partnerships.[5]

Actual and Potential Opportunities and Risks Related to Use of AI in Financial Services

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.

The AI Opportunity
AI can result in significant benefits for the financial services industry. First, it involves sophisticated decisioning for events that occur millions of times each day. Second, responsible use of the available data has the potential to be predictive and improve the accuracy of credit decisions. Third, given the costs and risks associated with lending, the economic wins from AI are dramatic for both banks and consumers. Fourth, AI facilitates credit access expansion by leveraging alternative data points beyond traditional attributes of an applicant (e.g. credit score) to reach those who have been traditionally overlooked for favorable credit. This means that the significant investment required to overcome the technical and regulatory hurdles is well worth the effort.

[5] In this regard, the FDIC issued a request for information on standard setting and voluntary certification for models and third-party services providers, overseen by regulators See https://www.federalregister.gov/documents/2020/07/24/2020-16058/request-for-information-on-standard-setting-and-voluntary-certification-for-models-and-third-party (July 24, 2020).

Our AI Lending Platform

Our AI models are central to our value proposition and unique position in the industry. Our models incorporate more than 1,600 variables, which are analogous to the columns in a spreadsheet. They have been trained by more than 58 million repayment events, adding an average of 83,000 new repayments each business day, analogous to rows of data in a spreadsheet. Interpreting these billions of cells of data are increasingly sophisticated machine learning algorithms that enable a more predictive model.

These elements of our model are co-dependent; the use of hundreds or thousands of variables is impractical without sophisticated machine learning algorithms to tease out the interactions between them. And sophisticated machine learning depends on large volumes of training data. Over time, we have been able to deploy and blend more sophisticated modeling techniques, leading to a more accurate system. This codependency presents a challenge to others who may aim to short-circuit the development of a competitive model. While incumbent lenders may have vast quantities of historical repayment data, their training data lacks the hundreds of columns, or variables, that power our model to better serve consumers. Despite their sophistication, our AI models are delivered to banks and credit unions in the form of a simple cloud application. Additionally, our platform allows lending partners to tailor lending applications based on their policies and business needs. Our lending partners can configure many aspects of their lending programs, including factors such as loan duration, loan amount, minimum credit score, maximum debt-to-income ratio and return target by risk grade. Within the construct of each lender's self-defined lending program, our platform enables the origination of conforming and compliant loans at a low per-loan cost.

Our platform benefits from powerful flywheel effects that drive continuous improvements as our business scales. Our platform benefits first from increasingly sophisticated models, variable expansion and rapid growth of training data. Upgrades to our platform allow lenders to offer higher approval rates and lower interest rates to consumers, which increases the number of borrowers on our platform. Upgrades to our platform also lead to better borrower selection, which lowers losses and lowers interest rates to borrowers. The flywheel effect created by self-reinforcing AI increases the economic opportunity that can be shared by borrowers and lenders over time.

Value Proposition to Consumers

- *Higher approval rates and lower interest rates*— Each year, Upstart evaluates the ability of our personal loan model to underwrite applicants in comparison to a more "traditional" model. We conduct the research looking back at the prior year. Our research showed that, in 2023, the Upstart model could approve more applicants, including Black and Hispanic applicants, at lower APRs than a more traditional underwriting model. In comparison to the traditional model, the Upstart model approved 101% more applicants and results in APRs that are 38% lower.[6]
- *Superior digital experience*—Whether consumers apply for a loan through www.Upstart.com or directly through a lending partner's website, the application experience is streamlined into a single application process and the loan offers provided are firm. Upstart enables banks and credit unions to respond instantly to a customer's

[6]https://www.upstart.com/lenders/regulatory-compliance/access-to-credit-report/

loan request 24/7 with no human intervention. More than 90% of borrowers are approved instantly, with zero documentation to upload.[7] Such automation improvements are due in large part to improvements to our AI models and the application of such models to different aspects of the loan process, including data verification and fraud detection.

Value Proposition to Bank Partners

- *Competitive digital lending experience*—We provide banks and credit unions with a cost effective way to compete with the technology budgets of their much larger competitors. The Net Promoter Scores, or NPS, for our bank partners' lending programs are approximately 83.5, well above the less than 30 NPS benchmarks for the top-tier banks.[8]
- *Expanded customer base*—We refer customers that apply for loans through www.Upstart.com to our lending partners, helping them grow both loan volumes and number of customers.
- *Lower loss rates*—An internal study comparing our model to that of several large U.S. banks found that our model could enable these banks to lower loss rates by almost 75% while keeping approval rates constant.[9]
- *New product offering*—Personal loans are one of the fastest-growing segments of credit in the U.S.[10] Our platform helps lending partners provide a product their customers want, rather than letting customers seek loans from competitors.
- *Institutional investor acceptance*—Analyses by credit rating agencies, loan and bond buying institutions, and credit underwriters help banks gain confidence that Upstart-powered loans are subject to significant and constant scrutiny from experts, the results of which are often publicly available.

Actual and Potential Risks and Risk Management

Oversight of AI—Explainability and Bias

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

[7] In Q1 2024. Percentage of Loans Fully Automated is defined as the total number of loans in a given period originated end-to-end (from initial rate request to final funding for personal loans and small dollar loans, and from initial rate request to signing of the loan agreement for auto loans) with no human involvement required divided by the Transaction Volume, Number of Loans in the same period.

[8] To determine Net Promoter Score (NPS) score, Upstart used a third-party service to administer surveys to personal loan applicants following an applicant's acceptance of a loan on Upstart's platform.While the NPS methodology used by Upstart's third-party service was designed to be consistent with the methodology used in the referenced benchmark study, any differences in the timing or method in which the surveys were administered could negatively impact the comparability of such NPSs. Source of bank NPS scores: Retently, "What is a Good Net Promoter Score? (2023 NPS Benchmark)," May 2023.

[9] In an internal study, Upstart replicated three bank models using their respective underwriting policies and evaluated their hypothetical loss rates and approval rates using Upstart's applicant base in late 2017. To compare the hypothetical loss rates between Upstart's model and each of the replicated bank models, Upstart held approval rates constant at the rate called for by each bank's respective underwriting policy. Such result represents the average rate of improvement exhibited by Upstart's platform against each of the three respective bank models.

[10] Beiseitov; see the section titled "Industry, Market and Other Data."

models and tools? What are the benefits and risks to financial institutions and consumers when the AI models and tools are interconnected among financial institutions?

Upstart develops our models in-house and our models are then used by our partner banks and credit unions. Well-regulated partnerships between financial institutions and third party technology companies, like Upstart, are critically important today for the financial health of consumers and the banking system. Because of Upstart's use of additional data and AI/ML techniques, lenders that work with Upstart are able to offer loans to more consumers who might not qualify using traditional underwriting methods. Upstart also provides technical integration support as part of the loan application processing and loan servicing services for our lending partners. Those functions are overseen based on the Interagency Guidance on Third-Party Relationships: Risk Management. Among other things, the use of comprehensive AI models enables our partners to increase the percentage of consumer loans that are made to low-and-moderate income communities.[11] It is also critically important that the use of AI leads to fair outcomes, and Upstart has worked proactively with the CFPB to demonstrate that using AI technology in lending can improve credit access and reduce interest rates for borrowers in all demographic groups, when compared to traditional underwriting approaches.[12]

Despite the proven benefits, financial institutions have been slow to embrace modern underwriting technology. All but the largest lenders face challenges in building advanced credit underwriting models in-house because the employees with the skill set required to build such models (*i.e.*, highly advanced computer science and mathematics) are not available to traditional financial institutions in all communities. Small and medium-sized lenders with more limited resources are especially disadvantaged given the resource-intensive nature of building and managing such technology. With physical bank branch networks continuing to shrink, effectively serving customers who need access to credit increasingly means offering convenient, fairly priced products *online*, *enabled for applications from mobile devices*. A study from Experian found that even before the pandemic, *50% of the US personal loan market was served online* by fintech platforms or lenders, up from 22% in 2015.[13] Offering technology that allows loan applications directly from smartphones can level the playing field; according to The Pew Charitable Trusts, more than eight in 10 Black and Hispanic Americans own smartphones today, nearly identical to the percentage of White Americans.[14]

Given these challenges (and opportunities), it makes more and more sense for banks and credit unions to access the benefits of technology through partnerships with third-party vendors like Upstart. That can only happen in a policy framework that supports banks' reliance on well-managed third party partnerships.

---

[11] As of April 18, 2024, 28.7% of Upstart-powered loans went to LMI communities. This is data based on loans originated on the Upstart platform from January 2017 to April 2024. LMI categorization is based on comparing median income in a customer's zipcode vs median income within the MSA of that zipcode.
[12] An update on credit access and the Bureau's first No-Action Letter. https://www.consumerfinance.gov/about-us/blog/update-credit-access-and-no-action-letter/
[13] https://www.experian.com/blogs/insights/2019/09/fintech-vs-traditional-fis-latest-trends-personal-loans/
[14] https://www.pewresearch.org/fact-tank/2021/07/16/home-broadband-adoption-computer-ownership-vary-by-race-ethnicity-in-the-u-s/

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.To what extent are financial institutions evaluating and addressing potential gaps in human capital to ensure that staff can effectively manage the development and validation practices of AI models and tools? What challenges exist for addressing risks related to AI explainability? What methodologies are being deployed to enhance explainability and protect against potential bias risk?

While AI offers a significant opportunity to improve the accuracy, fairness, and inclusiveness of the models used by financial institutions, Upstart believes it is also critical that AI model outputs meet a basic standard of being "explainable." The largest challenge for explainability of AI systems is the fact that they often make a large number of decisions before reaching their final output. Today there is a growing body of both established tools, and newer promising approaches, that can facilitate interpretation of complex AI and machine-learning models.[15] These include, for instance:

1. Shapley values (SHAP)[16]
2. Partial-dependence plots[17]
3. Relative importance[18]
4. Permuted feature importance[19]
5. Individual conditional expectation[20]
6. Local interpretable model-agnostic explanations (LIME)[21]

These techniques, either used separately or in combination, offer financial institutions ways to quantify the impact of particular data sets and even individual variables in a model. For example, one of the techniques Upstart uses is SHAP, which enables it to quantify the impact of certain variables on model outputs. The SHAP approach quantifies this impact by assessing the effect of a variable's removal from

[15] Papastefanopoulos, V., Kotsiantis, S. (2020). "Explainable AI: A Review of Machine Learning Interpretability Methods". Entropy, (2021), 23, 18. https://dx.doi.org/10.3390/e23010018
[16] Lundberg, S., Lee, S.I. "A Unified Approach to Interpreting Model Predictions". In: NIPS (2017). papers.nips.cc/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf
[17] Friedman, J. (2001). "Greedy Function Approximation: A Gradient Boosting Machine". The Annals of Statistics, 29:5, 1189-1232. See also, Goldstein, A., Kapelner, A., Bleich, J., Pitkin, E. (2015). "Peeking Inside the Black Box: Visualizing Statistical Learning with Plots of Individual Conditional Expectation". Journal of Computational and Graphical Statistics, 24:1, 44-65. https://doi.org/10.1080/10618600.2014.907095
[18] Breiman, L., Friedman, J., Olshen, R. and Stone, C. (1984). C. (1984). Classification and Regression Trees, Wadsworth, New York.
[19] Fisher, A., Rudin, C., Dominici, F. (2019). "All Models are Wrong, but Many are Useful: Learning a Variable's Importance by Studying an Entire Class of Prediction Models Simultaneously". Journal of Machine Learning Research, 20:177, 1-81. jmlr.org/papers/volume20/18-760-18-760.pdf
[20] Goldstein, A., Kapelner, A., Bleich J., Pitkin, E. (2015), "Peeking Inside the Black Box: Visualizing Statistical Learning With Plots of Individual Conditional Expectation". Journal of Computational and Graphical Statistics, 24:1, 44-65, DOI: 10.1080/10618600.2014.907095
[21] Ribeiro, M.T., Singh, S., Guestrin, C. (2016). "Why Should I Trust You?: Explaining the Predictions of Any Classifier". KDD '16: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 1135–1144. https://doi.org/10.1145/2939672.2939778

the model. This method allows each variable to be assigned some fraction of the model's output for which it is accountable (variables with a negative impact are assigned a negative proportion of this effect). In the context of Upstart's underwriting model, this method allows Upstart to ascertain which variables or variable groupings are most influential in producing credit decisions for a particular applicant. This method has numerous business applications including the determination of reasons that are included in adverse action notices ("AANs"). By observing the model's reliance not only on discrete variables but also of highly correlated variable groupings, more specific and nuanced adverse action reasons can be communicated to a consumer.

Upstart has found that to manage the risks associated with AI, financial institutions require rigorous testing and third party validation of AI model outputs in combination with the use of available tools and techniques to ensure accurate and relevant explanations of those final outputs. For example, during model development at Upstart, data scientists follow a rigorous statistical process of cross-validation to ensure that every added variable produces a robust improvement in model accuracy and the causal relationships are well understood. All testing is fully documented before any updated model is ready to be put into production.

Despite the growing body of explainability techniques outlined above, misplaced perceptions about insufficient "explainability" tools may delay the deployment of AI/ML systems in the banking system.[22] Institutional inertia or status quo bias may cause sluggish adoption of AI/ML applications, even if traditional models are less accurate and exhibit problems with interpretability.[23] These forces could impede the deployment of sound AI credit underwriting systems at financial institutions of all sizes, making these institutions less competitive.

Furthermore, in the case of credit underwriting models, these perceptions could become barriers to the availability of affordable consumer credit. Access to insured deposit funding typically means that banks are "...the most dependable, low cost, through-the-cycle source of credit for consumers, including LMI borrowers."[24] Fewer bank AI-powered choices will mean higher fees, higher interest rates, or, in certain cases, a lack of affordable access.

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

---

[22] For example, the Financial Stability Board released a report in 2017 assessing the costs and benefits of AI use in financial services. Lack of interpretability was identified as a key risk, with the claim that it could result in unpredictable and unforeseen actions with possible macroeconomic consequences. See: Financial Stability Board (FSB), Artificial intelligence and machine learning in financial services. Market developments and financial stability implications. Nov. 1, 2017, available at: https://www.fsb.org/2017/11/artificial-intelligence-and-machine-learning-in-financial-service/.

[23] Traditional credit models can suffer from problems with "interpretability." For example, an input variable in a simple multivariate regression model could be found empirically to influence default risk in a manner that is difficult to understand or rationalize.

[24] Bank Policy Institute and Covington, Artificial Intelligence Discussion Draft: The Future of Credit Underwriting: Artificial Intelligence and Its Role in Consumer Credit (2019) at p. 6.

Our models incorporate more than 1,600 variables and benefit from a rapidly growing training dataset that currently contains more than 58 million repayment events, adding an average of 83,000 new repayments each business day.[25] The network effects generated by our constantly improving AI models provide a significant competitive advantage to our bank and credit union partners—more training data leads to higher approval rates and lower interest rates at the same loss rate.

The AI models underpinning the Upstart platform are central to its efficacy and the high-quality experience we provide to borrowers. Our models have evolved rapidly since our founding, key aspects of our AI models include:

Variables
Variables in our AI models have increased from 23 in 2014 to more than 1,600. These include factors related to credit experience, employment, educational history, bank account transactions, cost of living and loan application interactions.

Training Data
Our models have been trained by more than 58 million repayment events, adding an average of 83,000 new repayments each business day, such as a successful repayment or a delinquency. Upstart's models learn from repayment data even while loan principal remains outstanding, allowing us to improve our models rapidly, while keeping the appropriate model governance and oversight guardrails in place. New model deployments go through a formal model approval process set out in the Upstart Underwriting Model Change Management Policy. As per this policy, we employ a variety of tests to ensure the new model will perform to expectations once deployed. We have automated tests that run on a large number of former applicants, and we compare the default predictions, approval rates, and APRs of the new model versus the prior model. These tests help us to minimize the chance of unexpected issues once new models and new variables are deployed.

Modeling Techniques
Growth in training data has enabled the development of increasingly sophisticated modeling techniques. For example, while earlier versions of our AI models were centered on logistic regression, our more recent models incorporate stochastic gradient boosting. We expect that our data science investments and continued growth of training data will unlock even more powerful techniques over time.

Model Applications
While our first model focused on predicting the likelihood of loan default, we have since applied models throughout the process of credit origination. These models quantify and reduce risk in various ways, while also increasing automation and funnel conversion.

The currently active AI models within the Upstart platform—shared by and available to all Upstart's bank and credit union partners—include:

- Income fraud—quantifies potential misrepresentation of borrower income;
- Acquisition targeting—identifies consumers likely to qualify for and have need for a loan;
- Loan stacking—identifies consumers likely to take out multiple loans in a short period of time;

[25] References to variables refer to all raw variables and certain combined variables considered in our AI models.

- Prepayment prediction—quantifies the likelihood that a consumer will make payments on a loan earlier than originally scheduled;
- Identity fraud—quantifies the risk that an applicant is misrepresenting their identity; and
- Time-delimited default prediction—quantifies the likelihood of default for each period of the loan term.

Upstart has demonstrated that helping financial institutions lend money safely to more consumers, or alternatively, reduce their credit losses, is an area uniquely suited to the use of AI and alternative data.[26] In 2017, the CFPB "estimate[d] that 26 million Americans are 'credit invisible,' meaning they have no credit history at all," and "another 19 million people have credit histories that are too limited or have been inactive for too long to generate a credit score" under traditional credit scoring models.[27] Furthermore, CFPB has found that Black and Hispanic Americans are more likely than white or Asian Americans to be credit invisible or to have un-scored records and typical approaches to building strong credit files – for example, "[t]he use of co-borrowers and authorized user account status – [are] notably less common in lower-income neighborhoods."[28] To promote fair access to credit for all individuals, including those in these circumstances, federal regulators have recognized that credit underwriting is an area where AI/ML and their use of alternative data can be particularly effective. The CFPB, for instance, has reported that:

> *"In addition to the use of alternative data, increased computing power and the expanded use of machine learning can potentially identify relationships not otherwise discoverable through methods that have been traditionally used in credit scoring. As a result of these innovations, some consumers who now cannot obtain favorably priced credit may see increased credit access or lower borrowing costs."*[29]

Upstart's AI/ML-powered credit underwriting model has demonstrated how AI/ML technology and alternative data use can significantly expand fair credit availability in a fair and responsible manner, consistent with fair lending laws and regulations, while also potentially improving bank safety and soundness.[30] Further, the automation of the loan applications, including the underwriting process, using AI/ML technology provides a more streamlined and efficient process that benefits both financial institutions and consumers; approximately 90% of loans originated through Upstart's platform by lending partners are fully automated.

There also are fair lending benefits to AI/ML credit underwriting models compared to traditional credit underwriting models. Use of AI/ML technology can help eliminate unconscious or conscious human bias in the credit underwriting process through the use of AI/ML credit underwriting models that require little,

---

[26] See Upstart.com Results To Date. https://www.upstart.com/about#results-to-date-3
[27] Schmidt & Stephens, supra note 8, at 141-142. 24 See id. 25 Cordray, supra note 1; see also Kenneth P. Brevoort, et al., CFPB Office of Research, "Data Point: Credit Invisibles," and https://files.consumerfinance.gov/f/201505_cfpb_data-point-credit-invisibles.pdf
[28] See Kenneth P. Brevoort & Michelle Kambara, CFPB Office of Research, "Data Point: Becoming Credit Visible," available at https://files.consumerfinance.gov/f/documents/BecomingCreditVisible_Data_Point_Final.pdf.
[29] "An Update on credit access and the Bureau's first No-Action Letter," CFPB Blog (Aug. 6, 2019), available at https://www.consumerfinance.gov/about-us/blog/update-credit-access-and-no-action-letter/.
[30] Regulators have consistently identified the importance of more accurate credit underwriting for safety and soundness of financial institutions: https://www.minneapolisfed.org/article/2014/underwriting-standards-lessons-from-the-past

if any, human intervention. Upstart has invested and continues to dedicate significant resources in the fair lending testing infrastructure to continuously monitor fairness of its models.

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?

Applying AI and richer data sets to lending has great potential to make lending more fair and more inclusive than the current traditional system.[31] In general, technological advances have benefitted consumers seeking credit by reducing the scope of human bias and enhancing the reach of human intelligence. While imperfect, automated credit scoring opened up credit to individuals who may have lacked the kind of personal history or relationships formerly needed to apply for credit at a local bank. In the same way, AI models open up credit to individuals who lack the kind of credit file formerly needed to secure an appropriately robust credit score.

Still, we must also be clear-eyed about the potential risks and take steps to address them.[32] Models that employ facially-neutral criteria and operate on large volumes of data could still end up doing little to improve on the legacy of discrimination, or even may exacerbate our credit system's deeply unequal status quo. Separately, the large number of data sources used by AI/ML algorithms could increase these risks if they are not selected with care and monitored. Upstart has demonstrated that these risks can be effectively mitigated in AI-based credit underwriting if variables are closely monitored for bias, in line with regulatory expectations.

Upstart has developed and routinely applied sophisticated fair lending tests, as well as access-to-credit tests, to all lending outcomes on its AI platform, covering nearly three million borrowers). To date, Upstart can report that financial institutions' use of Upstart's AI model has enabled a significant "expansion of credit access...across all tested race, ethnicity, and sex segments" and that the Upstart model does not introduce bias into financial institutions' credit decision process.[33]

---

[31] See, e.g., Richard Cordray, Director, CFPB, Alternative Data Field Hearing (Feb. 16, 2017), available at https://www.consumerfinance.gov/about-us/newsroom/prepared-remarks-cfpb-director-richard-cordray-alternative-data-field-hearing/.

[32] Klein, Aaron. "Reducing Bias in AI-Based Financial Services." Brookings Institute. 10 July 2020. https://www.brookings.edu/research/reducing-bias-in-ai-based-financial-services/. 28 June 2021.

[33] "An update on credit access and the Bureau's first No-Action Letter" https://www.consumerfinance.gov/about-us/blog/update-credit-access-and-no-action-letter/

Further, Upstart's fair lending reporting procedures ensure that its bank partners can validate that future versions of the model continue to be fair. To ensure sound, effective fairness testing, there are a number of techniques Upstart uses, and principles that Upstart adheres to, that could help guide best practices in fair lending evaluations of AI models. These techniques and principles may also help regulators evaluate whether additional guidance could be provided to market participants.

The current application of the Equal Credit Opportunity Act ("ECOA") and its implementing Regulation B ("Reg B") has produced a number of well-known approaches for financial institutions to analyze disparities in lending. Each approach has some strengths and weaknesses when applied to credit models of any type. However, a complete and optimal solution depends in part on the public policy objectives being pursued but likely *requires the use of several tests in concert*.

One of the simplest, yet effective, approaches to fair lending testing is to compare credit decisions on different groups of borrowers (statistical parity). For example, a test can assess whether different demographic groups have the same approval rate for a loan. The testing standard could use either a preset ratio threshold or a statistical significance test.[34] This test, however, suffers from an obvious drawback; because of large underlying socio-economic disparities between groups, different groups of borrowers in practice have different loan outcomes, i.e., they default at different rates when given loans. Forcing approval rate equality across groups would lead to either over-approving borrowers who would default, or under-approving borrowers who would repay. Both scenarios are bad for borrowers and conflict with financial institutions' core business objectives.[35]

A second essential testing methodology that can help address the shortcomings of statistical parity is a calibration test. Financial institutions will want to assess the parity of a model's output for different protected classes, conditional on the complete set of factors affecting creditworthiness, i.e., the actual outcomes that show which borrowers repay their loans. This can be implemented by testing the following question: When a model predicts a particular risk, does that translate to the same actual default rate across groups? If a lender finds that one demographic group actually defaulted at a higher rate than another group at the same predicted risk, that would justifiably cause concerns about fairness. Calibration is objective and generalizable. It doesn't incentivize distortionary behavior or incline financial institutions to approve overly risky borrowers. And it has practical relevance to borrowers; it evaluates both approval and APR decisions. This makes calibration a sound approach from a fair lending perspective, and calibration tests are widely used in practice. There are drawbacks here as well, however.[36] These are just

[34] This is known as statistical parity in the literature. See, e.g., https://arxiv.org/pdf/1703.09207: Berk, Richard, et al. "Fairness in criminal justice risk assessments: The state of the art." Sociological Methods & Research (2018): 0049124118782533.
[35] Given these realities, this standard could push financial institutions to avoid marketing to certain disadvantaged groups of consumers, even if many would repay, to avoid higher defaults.
[36] Details of a chi-squared test for calibration proposed by VantageScore for credit scoring models: https://www.vantagescore.com/wp-content/uploads/2022/02/VantageScore_Statistical_Bias_whitepaper-05-2015.pdf. The shortcomings of calibration include the fact that financial institutions don't observe outcomes at every level of predicted risk (they do not lend to borrowers whose predicted risk is higher than the risk tolerance). This is an effective test in the spectrum of risk scores where a lender lends, so the wider that spectrum is, the more complete this test is. Unfortunately, in some specific cases a model could be calibrated without being useful to consumers. Consider a hypothetical situation where we have a model based only on credit score, and two demographic groups. Suppose the credit score is very predictive for one demographic group, and therefore gives them a range of scores across the spectrum, but the score is not predictive for another group, for example because they lack credit history. A calibrated model could output a range of (cont.) predicted default probabilities for the first

two of the eight different types of fairness tests and supplemental approaches that Upstart and other industry participants can apply to evaluate the fairness of an AI model.[37]

Upstart has always encouraged the adoption of guidance on appropriate testing methods and basic principles that should be followed in fair lending testing. An effective fair lending testing regime must follow certain basic principles. First, testing should be objective rather than subjective, and the tests themselves should not leave room for human interpretation of whether a practice seems fair or reasonable. Furthermore, testing should be readily understandable and verifiable – a lender should not be able to manipulate or hide the true results of the test.

A second key principle is that fair lending tests should be universal and generalizable. This means that tests should be applicable to all competing approaches to both underwriting and pricing. Upstart's thesis is that an effective test should not *a priori* assume that any specific approach to lending is fair, even if, for example, it has been used historically. Furthermore, the test of fairness, at least conceptually, should be extensible to different technical approaches and innovations, including both traditional methods of underwriting consumers and any new or innovative technologies that might become available.

Third, the most effective and appropriate fair lending tests measure quantities that are relevant to the consumer with achievable target thresholds. Tests should be designed to measure impacts on the consumer, rather than operating as a purely theoretical or intellectual exercise. The more complete a test's coverage of the quantities relevant to the end consumer, the more relevant the test becomes. For example, tests should measure key issues like approval rates and interest rates/APRs. Fair lending testing standards should actually be achievable in practice and take into account real world constraints. For example, certain disparity thresholds in lending may be nearly impossible to achieve in practice by any individual lender because of the large systemic inequalities in socioeconomic conditions outside of the lender's control. Therefore, for any fair lending test to not be self-defeating, it must be possible for a lender to meet that test sustainably and without undermining its business because a lender that stops lending responsibly does not serve the credit needs of consumers.

Fourth, Upstart's view is that effective fair lending tests should not inadvertently limit progress towards important fairness objectives, such as equitable access and financial inclusion, because the testing focus is limited to only certain aspects of fairness or seeks elimination of only certain types of disparity. Given the

---

group, but would output the same (mean) score for everyone in the second group because they are indistinguishable from the perspective of the model. This would lead to a situation where the model approves nobody from the second group if their mean risk is above our risk tolerance. That would be unfair to applicants from the second group, some of whom at least would actually repay their loan. These concerns can be mitigated by specifically assessing their relevance to the lender in question, or by combining calibration with other fair lending tests.

[37] These eight tests are (i) demographic parity test, (ii) constant test, (iii) classification parity test, (iv) calibration test, as well as supplemental tests such as (v) equal accuracy tests, (vi) comparison test, (vii) debias test, and (viii) tradeoffs test. See: "Does Credit Scoring Produce a Disparate Impact?" Federal Reserve Finance and Economics Discussion Series. Divisions of Research & Statistics and Monetary Affairs Avery, Brevoort, Canner.. "Fairness in criminal justice risk assessments: The state of the art." Sociological Methods & Research (2018) Berk, Richard, et al. "Fairness through awareness." Dwork, Cynthia, et al. "On conditional parity as a notion of non-discrimination in machine learning." Ritov, Ya'acov, Yuekai Sun, and Ruofei Zhao. "The measure and mismeasure of fairness: A critical review of fair machine learning." (2018) Corbett-Davies, Sam, and Sharad Goel. "Inherent trade-offs in the fair determination of risk scores." Kleinberg, Jon, Sendhil Mullainathan, and Manish Raghavan. "Tracking and Improving Information in the Service of Fairness." Proceedings of the 2019 ACM Conference on Economics and Computation. (2019). Garg, Sumegha, Michael P. Kim, and Omer Reingold. "Certifying and removing disparate impact." Feldman, Michael, et al. ACM international conference on knowledge discovery and data mining (2015).

consensus that the status quo in credit scoring and access to credit is far from ideal, it is important for any testing regime not to lock-in the status quo. Among other things, this means a fairness test should not be structurally anti-change or incumbent-preferring.

Finally, fair lending standards should not encourage financial institutions to make loans to borrowers who will likely be unable to repay them. Extending consumer credit is not beneficial to all consumers in all circumstances. Any fair lending standards that compel a lender to extend credit to a consumer who is unlikely to be able to repay the loan may lead to default, bankruptcy or financial hardship. Fair lending tests should avoid inadvertently distorting lender incentives towards practices that would run counter to key policy goals, such as access; lending responsibly, i.e., to those with ability to repay; and fairness, including in approvals and pricing.

Third-Party Risks

Question 15: To the extent financial institutions are relying on third-parties to develop, deploy, or test the use of AI, and in particular, emerging AI technologies, how do financial institutions expect to manage third-party risks? How are financial institutions applying third-party risk management frameworks to the use of AI? What challenges exist to mitigating third-party risks related to AI, and in particular, emerging AI technologies, for financial institutions? How have these challenges varied or affected the use of AI across financial institutions of various sizes and complexity?

On July 25th, 2024, federal prudential regulators released a statement regarding middleware and deposit/payment partnerships and announced a more general request for information (RFI) covering third party financial technology partnerships. Upstart is engaged only in lending partnerships and plans to submit detailed comments on the relevant questions.[38] It is critical that regulators and bank supervisors acknowledge that partnerships with technology firms are likely the only way that the vast majority of banks and credit unions will be able to overcome the many barriers that stand in the way of a successful digital transformation of their traditional branch-based consumer lending programs. It is simply too much to expect that any but the largest banks in the United States will be able to organically develop the software, methods and the associated technical expertise, to manage a successful online consumer lending program that uses advanced AI/ML techniques. A successful program requires more than an "Apply Here" button on a bank's website. From online customer acquisition to fraud protection to underwriting and pricing, to meeting the demands of modern mobile and online experiences, to digital servicing and collections, in the current age, it is a complex enterprise.

Each of the prudential regulators has issued safety and soundness guidance for the institutions it supervises on managing risk in connection with the use of third-party vendors[39] and extensive oversight of third-party vendors is expected from financial institutions who use them. We continue to observe that

[38] https://www.fdic.gov/news/press-releases/2024/agencies-remind-banks-potential-risks-associated-third-party-deposit

[39] See Board of Governors of the Federal Reserve System, SR Letter 13-19, "Guidance on Managing Outsourcing Risk" (December 5, 2013); Federal Deposit Insurance Corp., FIL 44-2008, "Third-Party Risk: Guidance for Managing Third-Party Risk" (June 6, 2008); National Credit Union Administration, SL No. 07-01, "Evaluating Third-Party Relationships" (October 2007); Office of the Comptroller of the Currency, Bulletin 2013-29, "Third-Party Relationships" (October 30, 2013). See also Consumer Financial Protection Bureau, Bulletin 2012-03, "Service Providers" (2012); OCC, Bulletin 2020-10, "Third-Party Relationships: Frequently Asked Questions to Supplement OCC Bulletin 2013-29" (March 5, 2020); FDIC, FI-50-2016, "Request for Comment on Proposed Guidance for Third-Party Lending."

small and mid-sized institutions are reluctant to commit resources to vendor relationships without assurances from their regulators about how they can do so in a manner that is consistent with the regulators' expectations for prudent risk management. Furthermore, we have seen examples of examiners considering applying third party risk management expectations that a) inconsistent across similarly-situated market participants and b) which would be cost-prohibitive for small banks and credit unions. Current guidance directed more specifically to the management of model risk, which pre-dates the agencies' third-party risk guidance, does not resolve this uncertainty.[40]

First, there are often varying interpretations of the existing model risk management governance of these models and the exact oversight responsibilities banks have when they engage a vendor that employs an AI model. Uncertainty surrounding the appropriate method for applying the existing model risk management guidance to third party AI technology – and the supervisory application of the principles – can discourage banks from working with new vendors if they believe using a new third party risk management system will lead to much more scrutiny than remaining with a prior vendor. Note that banks can also be discouraged due to uncertainty about the appropriate fair lending testing regime that regulators expect will be applied.

Work done by the Bank Policy Institute ("BPI") suggests that although 2019 regulatory guidance technically "gives banks flexibility to modify the model risk management framework for validating vendor and other third party models," the reporting on the ground reveals that federal banking regulators "have not consistently afforded this flexibility to banks with regard to vendor-developed AI credit underwriting systems." According to BPI, regulators "have not applied a similar review or approval process to widely used conventional underwriting systems."[41] If this approach persists, it will create an unlevel playing field – one that fails to harness the benefits of new models or acknowledge that traditional models may be less accurate and more biased against protected groups.

Second, absent new or revised formal guidance, many financial institutions may not be able to participate in the growing adoption of AI and may not become aware of the growing recognition of responsible AI/ML model use by federal regulators focused on both prudential supervision and regulation and on consumer protection. Federal regulators, therefore, have the opportunity to significantly improve banks' ability to safely use AI technology in the near term by issuing examiner guidance that would clarify the application of existing model risk management and third-party relationship risk management principles to AI/ML credit underwriting models, including those sourced from third-party vendors.

We recognize that recalibrating the approach to model risk management on an interagency basis is a difficult undertaking that may take significant time and effort. There are steps the agencies can take sooner rather than later, however, that would address the current uncertainty and facilitate the use of responsible innovation while that process is ongoing. For example, issuing examiner guidance would be consistent with regulators' longer-term effort to modernize its digital activities regulations. The guidance would not displace the current Model Risk Management guidance or the Third-Party Service Provider

---

[40] Board of Governors of the Federal Reserve System and Office of the Comptroller of the Currency, "Supervisory Guidance on Model Risk Management" (2011); FDIC, FIL-22-2017, "Adoption of Supervisory Guidance on Model Risk Management" (June 7, 2017).
[41] Id.

guidance, which would remain applicable. It would complement those documents in a relevant and practical manner and could be updated as principles-based guidance or regulations evolve.

Key concepts for any updated examiner guidance could include:

- Federal regulators should affirm that vendors and technology providers harnessing AI and machine learning and providing similar services to supervised entities should be treated similarly.
- Federal regulators should confirm that banks' reliance on independently validated AI credit underwriting models managed by third parties is recognized/appropriate;
- The updated guidance should revise existing methods used and recommended by bank examiners for validating and fairness testing, so they are effective and relevant for evaluating a complex AI model, meaning that validation activities would be conducted largely at the model level, rather than the variable level.
- The guidance should make clear that while examiners should expect banks to develop a "detailed knowledge" of vendor-provided models, "detailed knowledge" does not require banks to have a detailed understanding of the model at the code level, just as financial institutions are not currently required to understand third-party proprietary credit scoring models at the code level.
- In an effective program, "detailed knowledge" means an understanding of the different categories of variables, the techniques used by the model, and the key metrics by which model outputs are measured – providing banks with the ability to confirm that the use of the model is consistent with its prudent operation, safety and soundness, and fair lending. This is best accomplished by using appropriate metrics and regular tests for model accuracy and fairness.
- Finally, regulators' approaches to supervision should reflect that the practical application of model risk management principles to vendor-provided AI models will be different than it is in cases of bank-developed models.

Further actions

Question 18: What actions are necessary to promote responsible innovation and competition with respect to the use of AI in financial services? What actions do you recommend Treasury take, and what actions do you recommend others take? What, if any, further actions are needed to protect impacted entities, including consumers, from potential risks and harms? Please provide specific feedback on legislative, regulatory, or supervisory enhancements related to the use of AI that would promote a financial system that delivers inclusive and equitable access to financial services that meet the needs of consumers and businesses, while maintaining stability and integrity, protecting critical financial sector infrastructure, and combating illicit finance and national security threats. What enhancements, if any, do you recommend be made to existing governance structures, oversight requirements, or risk management practices as they relate to the use of AI, and in particular, emerging AI technologies?

Upstart recommends that the regulators take actions to address the public policy points made by the Monitor in the Fair Lending Monitorship of Upstart Network's Lending Model.[42] More specifically, Upstart agrees that regulators should set the expectation that all lenders routinely assess their credit models for discrimination risks. Many companies, especially in financial services, have been testing

[42] Final Report of the Independent Monitor, Fair Lending Monitorship of Upstart Network's Lending Model, March 27, 2024. (https://www.relmanlaw.com/media/news/1512_Upstart%20Final%20Report.pdf)

models for discrimination for years, but there is variability in the effectiveness of these assessments. Effective systems should align with regulatory expectations and traditional anti-discrimination principles. These assessments should include both qualitative and quantitative reviews to ensure models are fair and do not have unlawful disparate impacts. Additionally, model testing should be part of a comprehensive fair lending compliance program addressing risks throughout the loan lifecycle, including with respect to less discriminatory alternative model searches. Regulatory agencies should ensure consistent fair lending model testing and provide guidance on appropriate methodologies. They should review model testing protocols and results in supervisory exams and share effective approaches publicly. Finally, regulators and legislators should enhance transparency in credit markets, encouraging companies to share data and modeling techniques to foster comprehensive understanding and improvements in lending practices.

Thank you for the opportunity to provide these comments. Please feel free to contact Gilberto Soria Mendoza, Government Relations Manager at gilberto.mendoza@upstart.com or Nat Hoopes, Head of Government & Regulatory Affairs at nat.hoopes@upstart.com if you have any questions or if you would like to discuss our comments further.

Sincerely,

Nat Hoopes
Head of Government & Regulatory Affairs, Upstart