EQUITABLE ALGORITHMS: HOW HUMAN-CENTERED AI CAN ADDRESS SYSTEMIC RACISM AND RACIAL JUSTICE IN HOUSING AND FINANCIAL SERVICES

VIRTUAL HEARING

BEFORE THE

TASK FORCE ON ARTIFICIAL INTELLIGENCE

OF THE

COMMITTEE ON FINANCIAL SERVICES

U.S. HOUSE OF REPRESENTATIVES

ONE HUNDRED SEVENTEENTH CONGRESS

FIRST SESSION

MAY 7, 2021

Printed for the use of the Committee on Financial Services

Serial No. 117–23

U.S. GOVERNMENT PUBLISHING OFFICE
44–838 PDF WASHINGTON : 2021

HOUSE COMMITTEE ON FINANCIAL SERVICES

MAXINE WATERS, California, *Chairwoman*

CAROLYN B. MALONEY, New York
NYDIA M. VELÁZQUEZ, New York
BRAD SHERMAN, California
GREGORY W. MEEKS, New York
DAVID SCOTT, Georgia
AL GREEN, Texas
EMANUEL CLEAVER, Missouri
ED PERLMUTTER, Colorado
JIM A. HIMES, Connecticut
BILL FOSTER, Illinois
JOYCE BEATTY, Ohio
JUAN VARGAS, California
JOSH GOTTHEIMER, New Jersey
VICENTE GONZALEZ, Texas
AL LAWSON, Florida
MICHAEL SAN NICOLAS, Guam
CINDY AXNE, Iowa
SEAN CASTEN, Illinois
AYANNA PRESSLEY, Massachusetts
RITCHIE TORRES, New York
STEPHEN F. LYNCH, Massachusetts
ALMA ADAMS, North Carolina
RASHIDA TLAIB, Michigan
MADELEINE DEAN, Pennsylvania
ALEXANDRIA OCASIO-CORTEZ, New York
JESÚS "CHUY" GARCIA, Illinois
SYLVIA GARCIA, Texas
NIKEMA WILLIAMS, Georgia
JAKE AUCHINCLOSS, Massachusetts

PATRICK McHENRY, North Carolina, *Ranking Member*
FRANK D. LUCAS, Oklahoma
BILL POSEY, Florida
BLAINE LUETKEMEYER, Missouri
BILL HUIZENGA, Michigan
STEVE STIVERS, Ohio
ANN WAGNER, Missouri
ANDY BARR, Kentucky
ROGER WILLIAMS, Texas
FRENCH HILL, Arkansas
TOM EMMER, Minnesota
LEE M. ZELDIN, New York
BARRY LOUDERMILK, Georgia
ALEXANDER X. MOONEY, West Virginia
WARREN DAVIDSON, Ohio
TED BUDD, North Carolina
DAVID KUSTOFF, Tennessee
TREY HOLLINGSWORTH, Indiana
ANTHONY GONZALEZ, Ohio
JOHN ROSE, Tennessee
BRYAN STEIL, Wisconsin
LANCE GOODEN, Texas
WILLIAM TIMMONS, South Carolina
VAN TAYLOR, Texas

CHARLA OUERTATANI, *Staff Director*

TASK FORCE ON ARTIFICIAL INTELLIGENCE

BILL FOSTER, Illinois, *Chairman*

BRAD SHERMAN, California
SEAN CASTEN, Illinois
AYANNA PRESSLEY, Massachusetts
ALMA ADAMS, North Carolina
SYLVIA GARCIA, Texas
JAKE AUCHINCLOSS, Massachusetts

ANTHONY GONZALEZ, Ohio, *Ranking Member*
BARRY LOUDERMILK, Georgia
TED BUDD, North Carolina
TREY HOLLINGSWORTH, Indiana
VAN TAYLOR, Texas

CONTENTS

EQUITABLE ALGORITHMS: HOW HUMAN-CENTERED AI CAN ADDRESS SYSTEMIC RACISM AND RACIAL JUSTICE IN HOUSING AND FINANCIAL SERVICES

Friday, May 7, 2021

U.S. HOUSE OF REPRESENTATIVES,
TASK FORCE ON ARTIFICIAL INTELLIGENCE,
COMMITTEE ON FINANCIAL SERVICES,
Washington, D.C.

The task force met, pursuant to notice, at 12 p.m., via Webex, Hon. Bill Foster [chairman of the task force] presiding.

Members present: Representatives Foster, Sherman, Casten, Pressley, Adams, Garcia of Texas, Auchincloss; Gonzalez of Ohio, Loudermilk, Budd, Hollingsworth, and Taylor.

Ex officio present: Representative Waters.

Chairman FOSTER. The Task Force on Artificial Intelligence will come to order. Without objection, the Chair is authorized to declare a recess of the task force at any time.

Also, without objection, members of the full Financial Services Committee who are not members of this task force are authorized to participate in today's hearing.

As a reminder, I ask all Members to keep themselves muted when they are not being recognized by the Chair. The staff has been instructed not to mute Members, except when a Member is not being recognized by the Chair and there is inadvertent background noise. Members are reminded that they may only participate in one remote proceeding at a time. If you are participating today, please keep your camera on, and if you choose to attend a different remote proceeding, please turn your camera off.

Today's hearing is entitled, "Equitable Algorithms: How Human-Centered AI Can Address Systemic Racism and Racial Justice in Housing and Financial Services."

I now recognize myself for 4 minutes to give an opening statement.

Thank you, everyone, for joining us today for what should be a very interesting discussion. We have a great panel of witnesses that I know will provide some stimulating and thought-provoking points of view. Today, we are here to explore how artificial intelligence (AI) can be used to increase racial equity in housing and financial services. There has been extensive discussion around this topic, mostly focusing on the real problems that can occur when we use AI that can inherently or unknowingly be biased. I think that a lot of these issues can be more complicated and nuanced than how they are portrayed in the media, but it is clear that the use of AI is hitting a nerve with a lot of folks, and that concern is for a good cause. No one should be denied the opportunity to own a home, a pillar of the American Dream, because of a non-human, automated, and, often, unlawfully discriminatory decision. Regulators and policymakers have a big responsibility here, too.

We must actively engage in these sorts of discussions to determine what the best practices are and to enact laws that reflect and encourage those practices, while also fostering innovation and improvements. Ideally, we should get to a space where AI is not only compliant with and meeting the standards that we have set for fairness, but exceeding those standards. It should be a tool that augments and automates fairness, not something that we have to babysit to make sure that it is still meeting our standards. The real promise of AI in this space is that it may eventually produce greater fairness and equity in ways that we may not have contemplated ourselves. So, we want to make sure that the biases of the analog world are not repeated in the AI and machine-learning world.

I am excited to have this conversation to see how we can make AI the best version of itself, and how to design algorithmic models that best capture the ideals of fairness and transparency that are reflected in our fair lending laws. Thank you all again for being part of this important discussion, and the Chair will now recognize the ranking member of the task force, Mr. Gonzalez of Ohio, for 5 minutes for an opening statement.

Mr. GONZALEZ OF OHIO. Thank you, Chairman Foster. First of all, I want to say how pleased I am to work with you as I take on the role of ranking member of this important task force. You have always shown a great willingness to be a thoughtful, bipartisan partner, and I look forward to continuing our work together. I also want to thank Ranking Member McHenry, ranking member of the full Financial Services Committee, for putting his trust in me to lead on this task force. He has been a tremendous mentor to me, and a thoughtful leader on policies that promote and expand the use of innovative technologies.

Financial services is an industry that continues to be on the cutting edge of technology, as is evident through the use of AI and other emerging technologies. I believe that this committee, and particularly this task force, should embrace this innovation and continue to consider ways that Congress can provide helpful clarity to industry without stifling innovation. Technology can help to not only propel forward our advancements in the financial services industry, but can also foster further inclusion and opportunities to our unbanked and underbanked communities.

Advanced credit decision models can use AI to improve the confidence of lenders in extending credit, reducing defaults, and finding data that is not readily available for traditional assessments of creditworthiness.

Additionally, it is my belief that AI technologies can provide Federal regulators with additional oversight tools to reduce and prevent financial crimes. We should be encouraging Federal agencies to be working more with the industry in a way that fosters adoption and can assist on money laundering efforts. On top of using AI to catch bad actors, Federal entities can take steps to work with industry to further adopt the use of artificial intelligence through the use of RegTech, in order to help automate and streamline regulatory compliance.

Today's hearing is an important one. We are having an important discussion about some of the challenges the industry faces by employing this technology, specifically on bias in algorithms. I believe these discussions are important to have. We must acknowledge and recognize that these technologies, at times, are not perfect due to the inherent nature of a technology created by humans. It is vital, though, that we do not take steps backwards by overregulating this industry, which may have a chilling effect on the deployment of these technologies. Instead, my hope is that we will continue to work with the experts in industry in order to move forward in a bipartisan way that both celebrates the technological advancements and ensures that there is transparency and fairness through the use of artificial intelligence.

I look forward to hearing from our witnesses today about the importance of this technology in the financial services sector and how Congress can act to encourage innovation and promote fairness. And with that, I yield back.

Chairman FOSTER. Thank you. The Chair will now recognize the Chair of the full Financial Services Committee, the gentlewoman from California, Chairwoman Waters, for 1 minute.

Chairwoman WATERS. Thank you so very much, Chairman Foster. I am so delighted and excited about artificial intelligence, and I am very pleased that you chose to provide the leadership for this task force that will help us to understand how we can get rid of bias in lending, and other efforts that should be made throughout our society in dealing with, simply, fairness and justice. I am very pleased, and I think that our committee will provide the leadership in the Congress of the United States for dealing with this issue.

As a matter of fact, we created a Subcommittee on Diversity and Inclusion, and your Task Force on Artificial Intelligence works very well with that subcommittee, because actually, you are going down the same paths, looking at the same issues, and dealing with what we can do to get rid of injustice and unfairness. Thank you so very much, and, please, go forward, and you are the one to do it. Thank you very much. I yield back.

Chairman FOSTER. Thank you, Madam Chairwoman. Today, we welcome the testimony of our distinguished witnesses: Stephen Hayes, a partner at Relman Colfax PLLC; Melissa Koide, the founder and CEO of FinRegLab; Lisa Rice, the president and CEO of the National Fair Housing Alliance; Kareem Saleh, the founder of FairPlay AI; and Dave Girouard, the founder and CEO of Upstart.

Witnesses are reminded that their oral testimony will be limited to 5 minutes. You should be able to see a timer on your screen that will indicate how much time you have left, and a chime will go off at the end of your time. I would ask you to be mindful of the timer and quickly wrap up your testimony if you hear the chime so we can be respectful of both the witnesses' and the task force members' time.

And without objection, your full written statements will be made a part of the record.

Mr. Hayes, you are now recognized for 5 minutes to give an oral presentation of your testimony.

STATEMENT OF STEPHEN F. HAYES, PARTNER, RELMAN COLFAX PLLC

Mr. HAYES. Chairwoman Waters, Chairman Foster, Ranking Member Gonzalez, and members of the task force, thank you for giving me the opportunity to testify. My name is Stephen Hayes, and I am a partner at Relman Colfax, a civil rights law firm. We have a litigation practice focused on combating discrimination in housing and lending. We also provide legal counsel to entities, including counsel on testing algorithms for discrimination risks. I previously worked at the Consumer Financial Protection Bureau (CFPB).

Credit markets reflect our nation's history of discrimination. There are stark gaps in credit access and disparities in credit scoring and in populations with thin or no credit histories. There is evidence that some alternative data and AI-based machine-learning models (ML models) can help lenders make credit decisions for these groups, and so have the potential to expand access. Whether that is true in practice and whether any increases will improve or exacerbate disparities is a context-specific question. Use of alternative data and alternative models can also raise serious risks related to explainability, validity, and, of course, discrimination.

The Equal Credit Opportunity Act (ECOA) and the Fair Housing Act prohibit lending and housing discrimination. They prohibit intentional discrimination, sometimes called disparate treatment, as well as an unintentional type of discrimination called disparate impact. Disparate impact focuses on fair outcomes. Unlawful disparate impact occurs when: one, a policy disproportionately harms members of a protected class; two, either the policy does not advance an interest; or three, there is a less discriminatory way to serve that interest. And what that means in practice is that entities should not adopt policies, like models, that unnecessarily cause disparities.

These frameworks, in particular, disparate impacts, translate well to lending models, including to ML models. Some banks have been testing models for discrimination for years, and, of course, disparities remain in credit markets, and model fairness alone is not going to solve that problem. But these programs demonstrate that discrimination testing is possible, and it can be effective.

As a general matter, the best programs align with legal principles, so first disparate treatment. The programs ensure that models don't include protected classes or proxies as variables, and that the models are accurate across groups, which is important, but it is insufficient to eliminate discrimination. The programs include a disparate impact assessment using the three-step framework that I mentioned before.

The final step in that framework, minimizing the disparities caused by models, is key to this process. In the case of traditional models, this involves substituting variables in the models with the goal of identifying variations of models that maintain performance, but that have less disparate impact, and newer methods exist now that can improve upon that process for ML models.

Disparate impact testing can benefit businesses and consumers. It can create more representative training samples and increase access to credit over time. It can also counteract the legacies of historic and of existing discrimination. These tests are also paired with more holistic measures, like fair lending training for modelers, ensuring that teams have diverse backgrounds, reviewing policies within which models operate, and monitoring areas of discussion.

Finally, banks are expected to comply with agency model risk guidance, which is meant to help mitigate safety and soundness risks. And these principles are not focused on discrimination, but they can help facilitate discrimination testing because they create an audit trail for models, and they help establish monitoring systems for models.

In my experience, many companies understand that models can perpetuate discrimination, and they don't want to use discriminatory models. But at the same time, discrimination testing is very uneven, and oftentimes nonexistent, which is the result of legal and structural background characteristics that incentivize testing in some areas, but not in others.

Policymakers can take steps to ensure more uniform and effective testing. First, agencies like the CFPB can routinely test models for discrimination, including assessing whether less discriminatory models exist.

Second, agencies should announce the methodologies that they use to test models, and they should encourage adoption of discrimination-specific model risk principles.

And third, agencies should clarify that discrimination, including unnecessary disparate impact, is illegal across markets outside of traditional areas like credit and housing.

Thank you for considering my testimony today.

[The prepared statement of Mr. Hayes can be found on page 34 of the appendix.]

Chairman FOSTER. Thank you. Ms. Koide, you are now recognized for 5 minutes.

STATEMENT OF MELISSA KOIDE, FOUNDER AND CEO, FINREGLAB

Ms. KOIDE. Thank you so much, Chairman Foster. Good afternoon. And thank you, Chairwoman Waters, Ranking Member McHenry, Ranking Member Gonzalez, and the entire AI Task Force. My name is Melissa Koide, and I am the founder and CEO of FinRegLab. FinRegLab is a nonprofit research organization evaluating the use of new technologies and data in financial services to drive greater financial inclusion.

FinRegLab has focused on the use of alternative financial data and machine learning algorithms in credit underwriting because credit not only helps bridge short-term gaps, but it is critical for enabling longer-term investments for families and homes, education and small business.

The credit system, as we all realize, reflects and influences the ability of families and small businesses to participate in the broader economy, yet I think we also realize that about 20 percent of adults in the U.S. lack a sufficient credit history to be scored under the most widely-used models. Another 30 percent have struggled to access affordable credit because their scores were non-prime. Communities of color and low-income populations are substantially more likely to be affected. Nearly 30 percent of African Americans and Hispanics cannot be scored under traditional means compared to 16 percent of Whites and Asians.

Our work at FinRegLab directly intersects with the task force's inquiry into ways to safely harness the power of AI and data to increase opportunity, equity, and inclusiveness. FinRegLab's first empirical research evaluated cash flow data as a means to risk-assess underserved people in small businesses for credit. We found cash flow data has substantial potential to increase credit inclusion.

Our latest project, launched last month, focuses on machine learning algorithms and their use in credit underwriting. We are empirically evaluating the capability and performance of diagnostic tools that seek to explain machine learning underwriting models with respect to reliability, fairness, and transparency.

Financial services providers have begun using machine learning models in a variety of contexts because of the potential to increase the prediction accuracy. There are many ways AI and machine learning may be beneficial for consumers and small businesses, but the technology could also be transformational where information gaps and other obstacles currently heighten the costs and risks of serving particular populations. Yet, we all realize that the complexity of AI and machine learning models can make it harder to understand and manage, and they raise important concerns around exacerbating historical disparities as well as flaws in the underlying data.

Publicly-available research is limited, but what there is supports the general predictiveness benefits of machine learning. Yet, it also suggests the effects of fairness and inclusion may vary depending upon—and this is important—the underlying data used. Some sources suggest it can increase inclusion when used to analyze traditional credit bureau data, while other studies find mixed or even negative effects when additional supplemental data source is used. For this reason, we believe more research is needed to better understand the effect of machine learning alone and in conjunction with promising types of financial data.

So, what is happening in the market today? Some banks and non-banks are beginning to use machine learning algorithms directly in their underwriting models in order to evaluate applications for credit cards, and personal auto and small business loans. They are doing so to improve the credit risk accuracy, to leverage the speed and efficiency of the technology, and to keep up with competitors. Yet, while interest in machine learning is increasing, there are fundamental questions about the ability to diagnose and manage these model, and might both have general concerns about reliability, transparency, fairness, and specific Federal regulatory requirements that Steve just discussed.

FinRegLab is, therefore, partnering with researchers from the Stanford Graduate School of Business to evaluate the performance and the capabilities of explainability tools designed to help lenders develop and manage machine learning algorithms in credit underwriting. We will use the Federal requirements concerning risk model governance, fair lending, and adverse action disclosures as a starting point, but expect that our research may be useful to address broader questions about machine learning reliability and the use of diagnostic tools for managing algorithmic decisions in a range of contexts.

In addition to focusing on the machine learning explainability, we intend to continue to study the role of alternative financial data, both alone and in conjunction with AI and machine learning, to foster greater financial inclusion. Thank you very much.

[The prepared statement of Ms. Koide can be found on page 40 of the appendix.]

Chairman FOSTER. Thank you, Ms. Koide. Ms. Rice, you are now recognized for 5 minutes to give an oral presentation of your testimony.

STATEMENT OF LISA RICE, PRESIDENT AND CEO, NATIONAL FAIR HOUSING ALLIANCE

Ms. RICE. Chairman Foster, Ranking Member Gonzalez, and members of the task force, thank you so much for inviting me to testify at today's hearing. The National Fair Housing Alliance is the country's only national civil rights agency dedicated solely to eliminating all forms of housing and lending discrimination, and this includes eliminating bias- and algorithmic-based systems used in housing and financial services through our recently-launched Tech Equity Initiative.

How AI systems are designed, the data used to build them, the subjective renderings applied by the scientist creating the models, and other issues, can cause discrimination, create or further entrench structural inequality, and deny people critical opportunities. On the other hand, innovations in the area of artificial intelligence have the potential to reduce discriminatory outcomes and help millions of people. Much as scientists used the coronavirus to develop lifesaving vaccines, we can use AI to detect, diagnose, and cure harmful technologies that are extremely detrimental to people in communities.

We have biased AI systems because the data used to build the models is deeply flawed. Technicians developing the systems are not educated about how technology can render discriminatory outcomes, and regulators are not equipped to sufficiently handle the myriad manifestations of bias generated by the technologies we use in financial services and housing. Let's start with the data.

The building blocks for algorithmic tools are tainted data that is embedded with bias generated from centuries of discrimination. Not only are we building systems with biased data, but oftentimes datasets are underinclusive and not representative of underserved groups. As a result, for example, traditional credit scoring systems, as you just heard Melissa say, oftentimes cannot see the behavior of consumers that are not represented in the data. This is why communities of color are disproportionately credit invisible or inaccurately scored. For example, in Detroit, Michigan, almost 40 percent of Black adults are credit invisible. This pattern is common throughout our nation.

So, how do these consumers access quality credit opportunities, rent apartments, obtain affordable insurance, or access other important opportunities necessary for people to lead productive lives?

Technology does not have to be biased. There are mechanisms for producing fair systems, and I will mention just a few. One method of de-biasing tech is to integrate the review of racial and other forms of bias into every phase of the algorithm's life cycle, including data selection, development, deployment, and monitoring. The European Union's newly-proposed regulation for AI offers one way of addressing this issue. It creates a risk-based framework that considers technologies, like credit scoring, as a high-risk category because of the grave impact it has on people's lives. The proposal holds high-risk models to a higher standard and incorporates a review for discrimination risk in all aspects of the algorithm life cycle.

To help de-bias tech, all AI stakeholders, including regulators, scientists, engineers, and more, should be trained on fair housing and fair lending issues. Trained professionals are better able to identify red flags and design solutions for de-biasing tech. In fact, recent innovations in building fair tech have come from AI experts trained on issues of fairness. Increasing diversity will also lead to better outcomes for consumers. Research shows that diverse teams are more innovative and productive. Moreover, in several instances, it has been people of color working in the field who are able to identify potentially discriminatory AI systems.

I will close by calling out the need for the creation of a publicly-available dataset to be used for research and educational purposes. Congress should encourage the release of more loan-level data from the National Mortgage Survey and the national mortgage databases so researchers, advocacy groups, and the public can study bias in housing and finance markets and, in particular, as it may relate to AI systems.

Thank you so much for the opportunity to testify today.

[The prepared statement of Ms. Rice can be found on page 55 of the appendix.]

Chairman FOSTER. Thank you, Ms. Rice. Mr. Saleh, you are now recognized for 5 minutes.

STATEMENT OF KAREEM SALEH, FOUNDER AND CEO, FAIRPLAY

Mr. SALEH. Thank you, Chairwoman Waters, Chairman Foster, Ranking Member Gonzalez, and members of the task force, for the opportunity to testify today. My name is Kareem Saleh, and I am the founder and CEO of FairPlay, the world's first fairness-as-a-service company. I have witnessed firsthand the extraordinary potential of AI algorithms to increase access to credit and opportunity, but I have also seen the risks these algorithms pose to many Americans. If we are to fully harness the benefits of AI, we must commit to building infrastructure that embeds fairness in every step of the algorithm decisioning process.

Despite the passage of the fair lending laws almost 50 years ago, people of color and other historically-underprivileged groups are still denied loans at an alarming rate. The result is a persistent wealth gap and fewer opportunities for minority families and communities to create a prosperous future.

Why are we still so deeply unfair? The truth is that the current methods of bias detection in lending are completely unsuited to the

AI era. Even though lending has become AI-powered and automated, fair lending compliance is stuck in the analog past.

So how can we bring fair lending compliance into the 21st Century? We must give lenders the tools and guidance they need to increase fairness without putting their businesses at risk. Today, lenders are required to measure and remediate bias in their credit decisioning systems. If, say, Black applicants are approved at materially lower rates than White applicants, lenders must evaluate whether this disparity is justified by a business necessity or determine whether the lender's objectives could be met by a less discriminatory alternative. It is at this stage, the search for alternatives and the invocation of business justifications, where our current fair lending system has the greatest potential to evolve.

The way most lenders search for less discriminatory models involves taking credit scores out of an algorithm, re-running it, and evaluating the differences in outcomes for protected groups. This method almost always results in a fairer model, but also a less profitable one. This puts lenders in a catch-22. They would like to be fair, but they would also like to stay in business, plus there is no guidance on what constitutes an appropriate tradeoff between profitability and fairness, creating uncertainty for lenders about how to meet regulatory requirements. Worse still, lenders fear that the very act of trying to find a fairer, better means of underwriting or pricing loans could be used against them as evidence they knew their algorithms were biased to begin with.

Faced with this problem, most lenders opt for safety, writing explanations for the use of unfair models instead of searching for alternatives that may yield fairer results. The upshot is that fair lending compliance has become an exercise in justifying unfairness rather than an opportunity to increase inclusion.

Today, a better, fairer option exists, using AI fairness tools to debias algorithms without sacrificing profitability. Several AI techniques allow lenders to take a variable, like credit score, and disentangle its predictive power from its disparity-driving effects. In many instances, these AI fairness tools have increased approval rates for protected groups anywhere from 10 to 30 percent without increasing risk.

Of course, industry will need support in order to fully embrace the benefits of AI fairness. Here, Congress and regulators can play an important role by ensuring that fairness testing is being done by more lenders more often, applied to their underwriting, pricing, marketing, and collections models, and includes a robust search for less discriminatory alternatives.

In addition, policymakers should ease the fear of liability for lenders who commit to thoroughly searching for disparities and less discriminatory alternatives, to reward rather than punish those who proactively look for fairer systems. Regulators can provide guidance on how lenders should view the tradeoffs between profitability and fairness, and set expectations for what lenders should do if disparities are identified.

To bring fairness to AI decisions, we must build the fairness infrastructure of the future, not justify the discrimination of the past. Using AI de-biasing tools, we can embed fairness into the algorithmic decisions to promote opportunity for all Americans while allowing financial institutions to reap the rewards of a safe and inclusive approach. If we prioritize fairness, the machines we build will follow.

Thank you. I am happy to answer your questions.

[The prepared statement of Mr. Saleh can be found on page 69 of the appendix.]

Chairman FOSTER. Thank you, Mr. Saleh. Mr. Girouard, you are now recognized for 5 minutes to give us an oral presentation of your testimony.

STATEMENT OF DAVE GIROUARD, CEO AND CO-FOUNDER, UPSTART

Mr. GIROUARD. Chairwoman Waters, Chairman Foster, Ranking Member Gonzalez, and members of the Task Force on Artificial Intelligence, thank you for the opportunity to participate in today's conversation. My name is Dave Girouard, and I am co-founder and CEO of Upstart, a leading artificial intelligence lending platform headquartered in San Mateo, California, and Columbus, Ohio.

I founded Upstart more than 9 years ago in order to improve access to affordable credit through application of modern technology and data science. In the last 7 years, our bank and credit union partners have originated more than $9 billion in high-quality consumer loans using our technology, about half of which were made to low- and moderate-income borrowers. Our AI-based system combines billions of cells of training data with machine learning algorithms to more accurately determine an applicant's creditworthiness.

As a company entirely focused on improving access to affordable credit for the American consumer, fairness and inclusiveness are issues we care about deeply. The opportunity for AI-based lending to improve access to credit for the American consumer is dramatic, but equally dramatic is the opportunity to reduce disparities and inequities that exist in the traditional credit scoring system.

In the early days at Upstart, we conducted a retroactive study of a large credit bureau, and we uncovered a jarring pair of statistics: just 45 percent of Americans have access to bank quality credit, yet 83 percent of Americans have never actually defaulted on a loan. That is not what we would call fair lending. The FICO score was introduced in 1989 and has since become the default way banks judge a loan applicant, but, in reality, FICO is extremely limited in its ability to predict credit performance because it is narrow in scope and inherently backward-looking. And as consumer protection groups, such as the National Consumer Law Center, have highlighted, for the past 2 decades, study after study has found that African-American and Latino communities have lower credit scores as a group than White borrowers.

At Upstart, we use modern technology and data science to find more ways to prove that consumers are indeed creditworthy, to bridge that 45 percent versus 83 percent gap. We believe that consumers are more than their credit scores, and going beyond the FICO score and including a wide variety of other information, such as a consumer's employment history and educational background, results in significantly more accurate and inclusive credit modeling. While most people believe a more accurate credit model means saying, “no” to more applicants, the truth is just the opposite. Accurately identifying the small fraction of borrowers who are unlikely to be able to repay a loan is a better outcome for everyone. It leads to significantly higher approval rates and lower interest rates than a traditional model, especially for underserved demographic groups, such as Black and Hispanic applicants.

Since our early days, skeptics have asked whether AI models will hold up in a down economy. The tragedy of the COVID pandemic, where unemployment rose from 4 percent to more than 14 percent in just a few weeks, required that we prove our mettle, and, in fact, we did just that. Despite the elevated level of unemployment, the pandemic had no material impact on the performance of Upstart-powered loans held by our bank holders. With the support of a more accurate credit model powered by AI, our bank and credit union partners can have the confidence to lend regardless of the state of the economy. Imagine banks lending consistently and responsibly just when credit is needed most. That is an outcome for which we can all cheer.

The concern that AI in credit decisioning could replicate or even amplify human bias is well-founded. We have understood since our inception that strong consumer protection laws, including the Equal Credit Opportunity Act, help ensure that good intentions are actually matched by good outcomes. This is especially true when it comes to algorithmic lending. For these reasons and more, we proactively met with the appropriate regulator, the Consumer Financial Protection Bureau, well before launching our company. Quite simply, we decided to put independent oversight into the equation. After significant good-faith efforts, starting in 2015, between Upstart and the CFPB to determine the proper way to measure bias in AI models, we demonstrated that our AI-driven model doesn’t result in an unlawful disparate impact against protected classes of consumers.

Because AI models change and improve over time, we developed automated tests with the regulator’s input to test every single applicant on our platform for bias, and we provide the results of these tests to the CFPB on a quarterly basis.

In September 2017, we received the first no-action letter from the CFPB recognizing that Upstart’s platform improves access to affordable credit without introducing unlawful bias. Thus far, we have been able to report to the CFPB that our AI-based system significantly improved access to credit. Specifically, the Upstart model approves 32 percent more consumers and lowers interest rates by almost 3½ percentage points compared to a traditional model. For near prime consumers, our model approves 86 percent more consumers and reduces their interest rates by more than 5 percentage points compared to a traditional model.

Upstart’s model also provides approval rates and lower interest rates for every traditionally-underserved demographic. For example, over the last 3 years, the Upstart model helped banks that use Upstart approve 34 percent more Black borrowers than a traditional model would have, with 4-percentage-point lower interest rates. That is the type of consumer benefit we should all get excited about.

I apologize that I am running long, so I will be happy to just cut it here if that is what the committee would prefer.

[The prepared statement of Mr. Girouard can be found on page 30 of the appendix.]

Chairman FOSTER. Thank you, Mr. Girouard, for your testimony.

The Chair will now recognize himself for 5 minutes for some questions.

One big prerequisite to racial and gender equity is socioeconomic integration. Minorities and traditionally-disenfranchised individuals should have the same access to communities with quality schools, banks, grocery stores, and other community staples, all of which stem from where they are able to work and live. Additionally, socioeconomically-integrated communities foster a greater sense of understanding and tolerance across people from different walks of lives and experiences. So to that end, I am interested in exploring how AI, as well as optimally-designed subsidies, can help improve socioeconomic integration.

There are many possibilities on how to proceed. For example, one might decide to subsidize investments in communities that have historically suffered from redlining, but if those communities have subsequently gentrified, then blanket subsidies in those areas might not be justified, so a broader set of data would be needed.

Or perhaps we should just acknowledge that there are many situations where there is an essential tradeoff between fairness and profitability, so we should explicitly subsidize lenders to adopt a more fair model while retaining the power of AI to identify the most promising loans to subsidize. For example, there is a program in Ottawa, Canada, that has been using AI to identify areas undergoing gentrification or disinvestment by analyzing home improvements that are visible by Google Earth and satellite images. This sort of technology might be showing where we are gaining or losing socioeconomic integration and where subsidies might be appropriate.

My question is for, I guess, all of the witnesses here. If our goals are not only to eliminate unfairness going forward, but also to correct for past unfairness, what sort of changes to the objective functions or explicit subsidies would we want to optimize an AI program to measure and reward socioeconomic integration and other things that we are interested in promoting? You can take it in any order you want.

Ms. RICE. I can kick it off. One of the things that we have been championing, Chairman Foster, is the building and development of a really robust publicly-available dataset for research purposes and to help fashion technology that is more fair. What we are finding is that a lot of discrimination and biases that we are seeing in AIs that we use are not just in financial services and housing, but in every area—criminal justice, education, employment, et cetera. One of the challenges is that the datasets upon which the models are used are extremely flawed and insufficient. They are underrepresentative.

So, if we can build more robust datasets, we can even use synthetic data so we don't have to use completely pure original data that may raise privacy concerns. But if we had more robust datasets, not only could we ensure that we are building better models that are less discriminatory and that provide more socio-economic benefits for everyone in our society, but it would also give us better tools for a better foundation for diagnosing different forms of discrimination and building more accurate tools for rooting out discrimination in algorithmic-based systems.

Chairman FOSTER. Thank you. Does anyone else want to take on the sort of optimal subsidy part of the question?

Mr. SALEH. Congressman, I will say that our experience working in emerging markets is that if you can provide some sort of credit enhancement for lenders to incentivize them to lend into these sub-populations that are not well-represented in the data, you can both give people a bridge to being scorable in the future, and also incentivize the creation of a more robust corpus of data that is truly representative of the ability and willingness of some of these historically-underprivileged communities to pay back loans. So, I endorse very much the comments Lisa made, and I think that we should look at credit enhancement programs for lenders to incentivize exactly the kind of lending development you are talking about.

Ms. RICE. Yes. And Kareem's statement just reminded me that Canada has a program that does that. They actually subsidize, on the insurance base, consumers who get declined from the voluntary market, and so there is a subsidy program to provide insurance for those consumers. And it has actually helped build a more robust dataset, and we can provide more information about that later.

Chairman FOSTER. Yes, thank you. I think this is a very important area to pursue, to really use AI to promote what we want instead of just looking at it to prevent it from acting badly.

I now recognize the ranking member of the task force, Mr. Gonzalez of Ohio, for 5 minutes.

Mr. GONZALEZ OF OHIO. Thank you, Chairman Foster. Mr. Girouard, I want to start with you. I find your testimony and your entire business model, frankly, to be inspiring and interesting in so many ways. But I am curious as to how scalable the process was with the CFPB from the very beginning, because I think one concern I have is that the CFPB, or any other entity, might not be able to handle, say, 100 companies, Mr. Girouard, sort of what you guys did.

So I guess my first question would be, from a structure standpoint, how did you go about approaching the CFPB from the beginning, because you sort of embedded compliance in the very beginning, which makes perfect sense. But I am curious how that all played out, how that evolved, and whether or not you think whatever program you used could handle, let's say, 100 Upstarts if we ever got to that point. So, I will just kind of turn it over to you to comment on that.

Mr. GIROUARD. Sure. Thank you, Congressman. First of all, I will say one thing, which is that the Equal Credit Opportunity Act actually is quite useful. You might think of it like old legislation from decades ago being irrelevant today or just not keeping up with the times, but it actually does, to a large extent. It works and it can be implemented. But, of course, there is some ambiguity when you get into sort of algorithmic lending and such.

So, we introduced ourselves to the Consumer Financial Protection Bureau (CFPB) before we ever launched as a company because we were naive. People told us, you shouldn't go talk to the regulators, just sort of hide out, but we didn't believe that was the right path, so we introduced ourselves, and told them what we were hoping to achieve. And after years of good work, we got what is termed a no-action letter, which basically means trying to provide some clarity where there is ambiguity in the regulation. That, of course, is not a scalable path for anybody.

And we also necessarily took on a bit of risk in our early days because we didn't know what the outcomes of our models would be, but we were a startup, so we had the capacity to take on that risk. The reality is, if there is going to be a path forward where these tools are broadly used, and used in a responsible manner where they do not introduce bias, they do improve credit outcomes, it is going to require some form of legislation or rulemaking to standardize how testing is done. We have sort of done that one-off, but it is really not scalable to the larger industry, which is, I think, what is necessary.

Mr. GONZALEZ OF OHIO. Yes, I couldn't agree more, and I would love to follow up with you—I only have 3½ minutes left—to get your ideas on what that might look like because I think it is really important.

Ms. Koide, I want to move to you. We know that bank regulators are increasingly open to new kinds of underwriting as a driver for more inclusive lending and even for sounder lending. The agencies put out a joint statement on this. The CFPB provided the no-action letter with Upstart, as we all know. What are the obstacles to industry adoption of these new models? Is it mostly regulatory risk, or technological or cultural, or something else, and what else could be done to sort of clear the obstacles?

Ms. KOIDE. Yes, thank you for the question. We have been quite focused in providing some of the empirical analysis on alternative financial data cash flow information. And to clarify here, it is transaction data that you can see in a bank account and, importantly, even a prepaid card transaction product which we have greater coverage, especially among underserved communities and populations in terms of bank and prepaid access as compared to credit records and histories. And that research, I think, helped to inform the regulators' awareness. They had been thinking about alternative data for a while as well, but, nevertheless, providing that kind of research and empirical insight, I think, helped to inform the steps that the regulators took jointly to issue that statement.

There are, nevertheless, important questions around using new types of data in underwriting, and more generally as well. They extend from, how are we ensuring consumer permission information is able to flow—we have Section 1033 under the Dodd-Frank Act, for which we do not have rules written that would articulate that process and the data that would be then flowing under that authority—to how adverse action notices are ultimately sufficiently responded to? If you are going to be extending credit to somebody that is different from what they expected to receive or under different terms than they expected, you have to explain it. And I think articulating those explanations to consumers are areas where the industry has continued to think about, how do they provide those kinds of explanations in a way that is comfortable for consumers and responsive to [inaudible].

Mr. GONZALEZ OF OHIO. Great. Thank you so much, and I yield back.

Chairman FOSTER. Thank you, and I will now recognize the Chair of the Full Committee, Chairwoman Waters, for 5 minutes.

Chairwoman WATERS. Thank you so very much. This will be directed to Ms. Rice and Mr. Hayes. The Equal Credit Opportunity Act and the Fair Housing Act prohibit discrimination for protected classes in the extension of credit in housing. Earlier this year, the Federal Reserve, the FDIC, the OCC, the NCUA, and the Consumer Financial Protection Bureau sent out a request to financial institutions and other stakeholders on how AI and ML are being used in the financial services space, and how these activities conform with these laws. Additionally, the Federal Trade Commission issued a separate guidance that racial or gender bias in AI can prompt law enforcement action.

Ms. Rice and Mr. Hayes, are these Federal agencies doing enough to ensure that existing loans prevent bias and discrimination or providing sufficient accountability for disparate impacts that can result from the use of AI models? What should they be doing? Ms. Rice?

Ms. RICE. Chairwoman Waters, thank you so much for the question. The National Fair Housing Alliance is currently working with all of those institutions and all of those Federal agencies that you have just named on the issue of AI fairness. And one of the challenges that we face is that the institutions themselves don't necessarily have sufficient staff and resources in order to effectively diagnose AI systems, detect discrimination, and generate mechanisms and solutions for overcoming bias.

As an example, financial services institutions have been using credit scoring systems, automated underwriting systems, risk-based pricing systems for decades, right? And we are now finding out, in part by using AI tools, that these systems have been generating bias for decades and decades, but for all of these years, the financial regulators were really not able to detect the deep level of bias ingrained in these systems. So, we really have to support the Federal regulatory agencies, make sure they are educated, make sure they are well-equipped so that they can do an efficient job, not only working with financial services institutions, but also to make their systems more fair.

Chairwoman WATERS. Let me interrupt you here for a minute, Ms. Rice and Mr. Hayes. We would like this information brought to us because when we talk about the longstanding biases, we should be on top of fighting for resources and insisting that the agencies have what they need to deal with it. And because they are embedded now, it is because we have not done everything we could do to make sure that they are equipped to do what they needed to do to avoid and to get rid of these biases. So, we want the information. We want you guys to bring the information to us so that we can now legislate and we can go after the funds that are needed. I thank you for continuing to work on these issues, but I want you to bring that information to us so we can do some legislation.

Mr. Hayes, do you have anything else to add to this?

Mr. HAYES. I completely agree with Lisa. I am hearing what you are saying. I think that is a great idea. I say the agencies have been in learning mode for a few years, and now it is actually time to provide more guidance on how you should test AI models. I think industry is ready for that. We are ready for that. We would like to help inform that process, but I do think now is the time for some more generally applicable guidance and action in this space.

Chairwoman WATERS. I think that Mr. Foster would welcome additional information, as would other Members of Congress, including me, the Chair of this Financial Services Committee, because we cannot just wait, wait, wait, and tell the agencies to do better. We have to force them to do better. And enforcing them to do better means that we understand where the biases are, and we actually legislate and we tell the agencies what they have to do.

So, I am so pleased about this hearing today. And I am so pleased about the leadership of Mr. Foster. But this is a moment in history for us to deal with getting rid of discrimination and biases in lending and housing and all of this, and so help us. Help us out. Don't just go to them. Come to us and tell us what we need to do. Is that okay?

Thank you very much. I yield back the balance of my time.

Chairman FOSTER. Thank you, Madam Chairwoman. And I just wanted to say that if any of the Members or the witnesses are interested in sort of hanging around informally after the close of the hearing—it is something that we often do with in-person hearings, and we are happy to try to duplicate that in the online era here.

And the Chair will now recognize the gentleman from Georgia, Mr. Loudermilk, for 5 minutes.

Mr. LOUDERMILK. Thank you, Mr. Chairman. I appreciate having another very intriguing hearing on a very important matter here, especially as we adopt newer technologies in the financial services sector.

Last year, the FDIC issued a request for information regarding standard setting and voluntary certification for technology providers. The idea was to have a voluntary certification program to streamline the process for banks and credit unions to partner with third-party FinTech and AI providers. The proposal is intriguing to me because when I met with both financial institutions and technology providers, one of their biggest concerns with the current regulatory requirements is that it takes an enormous amount of time and due diligence every time they want to form a partnership. I believe streamlining the onboarding process is an important step toward encouraging these type of partnerships.

Mr. Girouard: what are your thoughts on this issue?

Mr. GIROUARD. Yes, this is a really important issue. We tend to serve community banks, smaller banks which are often struggling to compete with the larger banks that have a lot more technical resources and people they put against the diligence they are required to do to use any type of third-party technology in their business. And if you are Wells Fargo, or Chase, or PNC, you can spend all day and millions of dollars evaluating technology solutions. But if you are a community bank, that is not possible.

Mr. LOUDERMILK. Right.

Mr. GIROUARD. I think if you want to even the playing field, if you want to keep the smaller banks alive, valid in the communities they serve, you need to make it easier for them to adopt technology. And that doesn't mean sort of foregoing the evaluations or the prudence that you need to responsibly adopt it. It just means allowing them to essentially put their efforts together on some sort of standard that would allow small banks across the country to keep up with all the investment going on in the top handful of banks out there.

Mr. LOUDERMILK. So if we were able to streamline the ability to form these partnerships, would that benefit consumers by expanding the FinTech and AI products?

Mr. GIROUARD. Oh, for sure. Every month or so, we turn on another community bank who suddenly offers attractively-priced products with higher approval rates, lower interest rates, in their communities, and it is happening regularly. But, honestly, it is just the tip of the iceberg. The opportunity is so much larger, and most banks, frankly, just don't have those kinds of resources. This is a process that can take 6 months. You can go through hundreds of hours of meetings and discussions. You have your regulator come in that you talk to, whether it is the FDIC, the OCC, et cetera. There is this incredible process that most banks just don't have the time and resources to take on, so it just gets sidelined.

Mr. LOUDERMILK. Another topic that I have brought up in these hearings before is dealing with the issue of bias. We need to recognize the difference between what types of bias we want to have in AI versus those that need to be rooted out. Obviously, you have to have a level of bias to discriminate against those who can and cannot pay a loan back. Not all types of biases are bad. If you think about it, the whole purpose of using AI in loan underwriting is to be biased against those who are unable to repay a loan, or at least identify those who have the dataset that would say these folks are unlikely to pay a loan, or even just to set an interest rate. At the same time, algorithms obviously should not contain bias that is based on factors that are irrelevant to the actual creditworthiness of the borrower, like race, or gender, or any other factor.

Mr. Girouard, do you agree that we need to be careful not to eliminate all bias in AI, but, rather, we should be working to eliminate the types of bias that really don't belong there?

Mr. GIROUARD. Congressman, perhaps it is a bit of semantics, but we believe that bias is always wrong. Accuracy in a credit model is what we seek. And giving a loan to somebody who is going to fail to pay it back is not doing any good for them, so, of course, wanting to lend to people who have the capacity to pay it back is always our goal. But we don't view an accurate credit model or making offers of credit as good as possible for people who are likely to pay it back in any sense biased against everybody else. It is really just accuracy in predicting and understanding who has the capacity to repay.

Mr. LOUDERMILK. And maybe it is semantics, but what we are looking at is for AI to look at data, just hard data, regardless of any other demographic factor, just looking at the creditability of the borrower. And I see that as a technical term as a level of bias just to be able to determine, is this person able to pay back the loan in the amount that they are borrowing or are they not? Set all that other stuff aside. That is really what we want AI to be able to do, not look at race, or gender, or any of those factors. Just, are they of the income level, do they have the credit history, do they have a history of paying back loans, et cetera? That is really what we are trying to get to, correct?

Mr. GIROUARD. It is true that we are trying to have an accurate model that will lend to people who can pay it back, and we constantly strive to make our model more accurate because when we do that, it tends to approve more people at lower rates, and it actually disproportionately improves more underserved people—Black Americans, the Hispanic community—so that is all good. But having said that, my thorough belief is that you need a supervisory system, a separate system that watches and makes sure that we are not introducing bias.

Mr. LOUDERMILK. I agree, and I appreciate your answer. And I yield back.

Chairman FOSTER. Thank you. The Chair now recognizes the gentlewoman from Massachusetts, Ms. Pressley, for 5 minutes.

Ms. PRESSLEY. Thank you, Mr. Chairman, for convening this task force hearing, and to each of our witnesses for their testimony. Last year, I had the opportunity to ask the former CFPB Director about a practice that remains a serious concern to me: the use of information about people's education, including where they went to college, when making decisions about access to credit and the cost of credit. An investigation by consumer advocates shows that the artificial intelligence lending company, Upstart, was charging customers who went to Historically Black Colleges and Universities more money for student loans than customers who went to other schools, holding all else equal. Now, I know Upstart has vigorously denied these allegations, but I have here the first report prepared by Mr. Hayes and his colleagues as a part of a settlement the company reached with the NAACP Legal Defense Fund and the Student Borrower Protection Center.

On page 23, it appears to say that Upstart made significant changes to its business model after coming under fire for its lending practices. I will certainly be watching closely see if Mr. Hayes' firm can independently verify that these changes actually address the disturbing effects of Upstart's approach to lending. It is hard to imagine a practice that better illustrates the deep and lasting legacy of systemic racism in American higher education than educational redlining. That is why I was so troubled to see that yet another FinTech lender that uses AI, a company called Stride Funding, was engaged in what sounds like the very same discriminatory practices as Upstart. Mr. Hayes, should we be worried that these practices are driving racial inequality and leading to disparate outcomes for former students?

Mr. HAYES. Thank you, Representative. I will say as a general matter, every time you use data in a model, part of the reason for using that data is to replicate some patterns in that data, and we also know that there are disparities in our education system. As you pointed out, they are with respect to race, national origin, and sex. Those could be replicated if you use that data model that is risk. It is not inevitable. There are lots of ways to use data to design models so that you don't do that.

Our role in the Upstart and Student Borrower Protection Center matters was as an independent monitor, so I don't have views at this point about whether that has happened, whether those reports are accurate or not. That is part of our charge as an independent monitor. I think it is a risk. It is one that should be guarded against, and I think any company that uses this type of data should be very careful with it and test its intuition.

Ms. PRESSLEY. Okay. So, Mr. Hayes, how can Congress and financial regulators ensure that complex algorithms and machine learning [inaudible] have skewered the disparate and illegal impact of these lending practices? What can we do?

Mr. HAYES. That is a great question. I will say as an initial matter, there is a [inaudible] in AI and ML models, and some of them are quite difficult to explain, or may be impossible to explain. Others are not. Others are explainable. And as an initial matter, if an institution cannot explain its model, why it is reaching certain conclusions, it should be very hesitant or maybe not use it at all for important decisions. I think that is pretty key.

This goes also back to the point that Chairwoman Waters had made. I think it is a great opportunity for the CFPB to come in and start actively testing some of these models, to test some of these intuitions, to test if these risks are real. That is a role it can play. As an outside advocate, there is only so much you can do with the model. It takes an agency with supervisory authority to really help institutions understand how their models work and make sure they are not going to violate the law.

Ms. PRESSLEY. Okay. Thank you. These patterns are certainly very disturbing, and it seems that people have not learned from Upstart's errors. The discrimination against students who have gone to HBCUs and minority-serving institutions exacerbates the disproportionate burden of student loans on Black Americans and perpetuates economic discrimination. If the use of AI in lending is to continue and expand in the financial services sector, Congress and Federal regulators must be positioned to provide proper oversight. And, as I mentioned, I will be watching closely. Thank you. I yield back.

Chairman FOSTER. Thank you. The Chair now recognizes the gentleman from Texas, Mr. Taylor, for 5 minutes.

Mr. TAYLOR. Thank you, Mr. Chairman. It is great to be on the task force, and I appreciate the opportunity for this hearing. Ms. Pressley, I certainly hope you won't discriminate against me for having gone to college and business school in your district. Since Upstart has been named here, I would love to give the CEO an opportunity to respond to that question set.

Mr. GIROUARD. Sure. Thank you. And, Congresswoman, I certainly appreciate your concern, but I will say, first and foremost, I have dedicated my career to improving access to credit, and I stand proud with what we have accomplished and how we have done it. The use of education data, without question, improves access to credit for Black Americans, for Hispanic Americans, for almost any demographic that you can speak to. Our models aren't perfect, but they certainly are not discriminatory.

We had a disagreement with the Student Borrower Protection Center, and their conclusions, in our view, were inaccurate. Having said that, we very willingly began to work with them and to engage with them to figure out, are there ways we can make even more improvements to our testing and to our methodology, and we continue to do that, as well as with the NAACP Legal Defense Fund. So, I think Upstart has demonstrated good faith in trying to improve credit access for all and to do it in a fair way that is working proactively with regulators, is here working with lawmakers, and we will work with consumer advocates if they want to. We have nothing to hide, and frankly, we are proud of the effort we are making to improve access to credit for Americans.

Mr. TAYLOR. Ms. Pressley, do you want to ask a follow up? I would be happy to yield the floor to you to ask a follow up to Mr. Girouard, or I can continue on with my questioning.

[No response.]

Mr. TAYLOR. Okay. So, Mr. Girouard, I really appreciate what you are doing. I think you have an impressive model, and it is amazing to see the application of AI in the way you have done it. How do you source your loans? Are you doing those directly or are you doing those through traditional banking platforms?

Mr. GIROUARD. Borrowers come either to Upstart through our brand and recognizing our marketing efforts to say, come here and you can get a better loan than you can get elsewhere. They can also come directly through our bank partners. There are more than 15 banks on our platform which also can, using our technology, offer loans to their own customers. So, they can find us in many different ways.

Mr. TAYLOR. How big are your 15 banking partners? Are those kind of regional banks? Are those G-SIBs? Are those community banks?

Mr. GIROUARD. They vary from community banks to credit unions, and credit unions are, on our platforms, growing quite quickly.

Mr. TAYLOR. What is your average loan size?

Mr. GIROUARD. In the range of $10,000 to $12,000.

Mr. TAYLOR. Okay. I just want to put this card on the table—I was on a bank board for 12 years, and I sat on the loan committee, and so, I was part of approving every loan for 12 years. I can honestly say that never once was credit score determinative of a loan. To be very honest, in the director discussions, I would say that credit score didn't come up in [inaudible] percent of our loan decisions. So, the statement that you made about it being a primary means of making decisions at least was antithetical to my own limited experience. We were one of the 5,000 banks in the United States, in terms of how we thought about credit. And I will say that—

Mr. GIROUARD. I have yet to meet a bank that doesn't have a minimum credit score requirement for a loan, typically 680 or something of that nature. So if they are out there, I haven't met them yet.

Mr. TAYLOR. Okay. I see where you are coming from. I think I understand what you are saying. Thank you for that. That just kind of clarifies where you are coming from in that particular assessment. But again, I would just say that underwriting credit is very important, and the other thing is you want to have costs be lower. The final thing I would say is, if I add a whole bunch of regulations on UI commerce, doesn't that make it more expensive for you to do business and then, in turn, force you to raise your rates?

Mr. GIROUARD. It depends what that regulation is. A lot of times regulation can be clarity that actually helps adoption of the technology—

Mr. TAYLOR. If I make it more expensive for you to operate, doesn't that increase the cost of operating?

Mr. GIROUARD. Oh, by definition, it for sure does, Congressman.

Mr. TAYLOR. Okay. Thank you. I just would encourage my colleagues as we think about this, to make sure that we don't increase the cost of operating, and then, in turn, lower access to capital, which I think is our mutual objective. I yield back.

Chairman FOSTER. Thank you. The Chair will now recognize the gentlewoman from North Carolina, Ms. Adams, for 5 minutes.

Ms. ADAMS. Thank you, Mr. Chairman. Thank you for calling this hearing, and Chairwoman Waters, we appreciate your support as well. And to the witnesses, thank you for offering your expertise and your insights.

I am grateful to Representative Pressley for diving into educational redlining and its harmful impacts on HBCU students and graduates. Over the past year, we have seen examples of how using such data and algorithms by lenders could result in borrowers facing thousands of dollars in additional charges if they attended a minority-serving institution, like an Historically Black College or University (HBCU). I am a proud product of an HBCU, a 2-time graduate of North Carolina A&T, and a 40-year professor at Bennett College, also an HBCU. And I do know how invaluable these schools have been to my success, and their outsized role in the economic and social mobility of millions of Black people in this country. They play a critical role in diversifying the workforce, particularly the tech sector.

Ms. Rice, and Mr. Saleh, we know that AI bias is real. Can you speak to the importance and value of increasing the diversity among AI researchers, scientists, and developers to improve quality of algorithm development and datasets, and how can we ensure that HBCUs play a greater role in diversifying the AI pipeline?

Ms. RICE. Congresswoman Adams, thank you so much for that question. It is critically important. I mentioned earlier that the National Fair Housing Alliance has launched the Tech Equity Initiative. One of the major goals of the Tech Equity Initiative is to increase diversity in the tech field, and one of the ways of doing that, of course, as you just mentioned, is partnering with Black, Indigenous, and People of Color (BIPOC)-serving financial institutions and HBCUs. I hinted in my statement that the National Fair Housing Alliance has been working on tech bias issues since our inception almost 40 years ago. So, these issues—tech bias, AI algorithmic bias—are not new. They are just gaining more media attention.

But we have found that as we work with financial services institutions on the issue of tech bias, and we have been doing this, again, for almost 40 years, the more these financial services institutions—lenders, insurance companies, et cetera—as they diversify their employee base, they yield better policies that are more inclusive and fair, they also themselves design better systems that are not only more accurate, but have less discriminatory outcomes. And oftentimes, it is because those people of color who are working inside those institutions can see signs of discrimination. They can pick up on variables that are being used in the algorithm and, from their own personal experience, can detect and sort of understand how those variables can generate a discriminatory outcome.

I mentioned that a lot of the innovations that we are seeing in the AI field, a lot of the tech bias that has been documented has come from scientists like Joy Buolamwini, who is one of the most noted data scientists in the world. How did she detect that facial recognition systems were discriminatory? Because she was working on a project and facial recognition technology did not work for her Black face.

Ms. ADAMS. Right. Okay.

Ms. RICE. If she had not been Black, she wouldn't have noticed that. So, I yield to my colleague, Mr. Saleh.

Ms. ADAMS. Mr. Saleh?

Mr. SALEH. I don't have much to add to Lisa's excellent comments. Congresswoman, you are absolutely right. We must do more to diversify the population of people who are building AI systems, governing AI systems, and monitoring AI systems. The technology industry has not been sufficiently good in that regard.

Ms. ADAMS. We know that tenant-screening algorithms have been increasingly employed by landlords, but there is evidence that algorithms adversely affect Black and Latino renters. For example, when a Navy veteran named Marco Fernandez returned from deployment, and was trying to rent a house, the tenant-screen algorithm [inaudible]. I am going to have to yield back, Mr. Chairman. Thank you so very much, and thank you to our guests for your responses.

Chairman FOSTER. Thank you. The Chair now recognizes the gentleman from Indiana, Mr. Hollingsworth, for 5 minutes.

Mr. HOLLINGSWORTH. I appreciate the Chair, and I certainly appreciate the ranking member for having this great hearing today, talking about these very important topics. I certainly welcome and hope for more diversity in the technology field writ large, and to find more opportunities for more people to contribute their great talents to this country. I think that is what has made us a leader around the world in technology, and I hope it is what will continue to make us a leader of technology around the world.

Mr. Girouard, I wanted to talk a little bit about this for a second. I certainly know that you are a fan of making sure that your workforces and other workforces are very diverse. But I also want to recognize the desire that you have for ensuring that your platform isn't biased in some way, that you make money by making loans, and if you can find more creditworthy individuals, no matter what walk of life they come from, no matter what color their skin, no matter what background they may have than other potential technologies, then you are better off because of that. Wouldn't you agree that you are incentivized to make sure that you find as many opportunities to make creditworthy loans as possible?

Mr. GIROUARD. Yes, absolutely. The way my company grows is the AI models get smarter at identifying who will and won't pay a loan, and that might seem odd. You might think that could make you shrink, not grow, but, in reality, millions and millions of people who are actually creditworthy, in reality are not recognized as such by a credit score.

Mr. HOLLINGSWORTH. Right.

Mr. GIROUARD. And that little oddness there means the better our models get unbalanced, the more people get approved, and the lower the interest rates are. So, it is a sort of win for everybody as long as the technology keeps improving, and, thus far, it has worked well for us.

Mr. HOLLINGSWORTH. And I definitely want to get back to, how do we keep improving the technology, but I just want to hit this point once again because I think, frequently, it goes unsaid, that the wind is at your back. The goal is to increase the number of loans and, frankly, to find opportunities to make loans where others might not be able to make those loans or may not find that same opportunity. So it is not as if we are struggling to hold back a problem, but, instead, the problem resolution and the market incentive here are working in the same direction. And I think that is really important for us to remember because in many other places, they work in opposite directions.

Second, I want to come back to exactly what you said, which is, how do we improve this technology over time? How do we expand the breadth of this technology over time? And I wondered whether there are stories or narratives or specific points as to how we might do that, how we as policymakers might empower you, your cohorts, your colleagues, your counterparts, and, frankly, the next generation of "you's" to develop this technology and be able to make it mainstream so that we can empower more Americans, no matter the color of their skin, no matter their background, to be able to get access to financial capital.

Mr. GIROUARD. Yes. First, thank you for the question, Congressman. I think, first of all, one of the most important things that could happen, just to provide clarity, we are all for testing, as you can see. We believe we are leading the charge on how rigorous testing for bias can be and should be. And as much as it is probably to our benefit that no one else figured out how to do it and deploy this technology, it is to the country's benefit that there is as much of this used responsibly as possible.

The problem, of course, is that banks are regulated not by one agency, but by at least four, if not more than that, and you have State-level regulators as well. So, it is really difficult for technology like this to get a hold when, even within one regulator, there is not a consistent opinion. A supervisor of this bank might say one thing, and a supervisor of another bank says another thing, so the adoption ends up being very slow.

There is one other important matter I want to raise, which is that banks have to worry about consumer protection, et cetera. But on the other side, they have the bank solvency, the people who care about whether the bank is going to go out of business, and these are sometimes at odds because they are prevented from making loans to what the regulator would perceive as risky borrowers. So, you have this sort of governance of banks that is oftentimes in conflict with moving toward a more equitable, more inclusive lending program. And that is difficult—

Mr. HOLLINGSWORTH. Mr. Girouard, I think that is a great point and something we really need to hit home. What you are saying is, we care about the solvency of our financial markets, the safety, but we also care about the efficiency, and making sure we don't push one too far in favor of the other is a really important dynamic going forward. And I think Van Taylor hit on this, but regulation can both help efficiency, but it can also hurt efficiency greatly, and making sure we monitor that is very important. I yield back to the Chair.

Chairman FOSTER. Thank you. The Chair now recognizes the gentleman from Massachusetts, Mr. Auchincloss, for 5 minutes.

Mr. AUCHINCLOSS. Thanks, Mr. Chairman, for organizing this hearing, and to our witnesses for their terrific testimony and Q&A. Massachusetts has been really on the cutting edge of artificial intelligence and its use in computational biology, in insurance, in the provision of legal services, in investing in real estate, and also in thinking about the regulatory dimensions.

The Massachusetts State House has formed a Facial Recognition Commission, led by State Senator, Cindy Creem, in my district, because of concerns over facial recognition application. A study from MIT in 2018 found that while accuracy rates for White men were north of 99 percent with facial recognition technology, for Black women, it was significantly less. And, Ms. Rice, this is why I was very happy to hear you raise this issue.

I was wondering if I could really bring up two questions with you. The first is concerns you may have on proposed regulations for the introduction of facial recognition technology into the setting of housing. We are seeing already that smart home technology, like Latch, or smart keypads and Nests are really becoming standard fare, and I don't think it is very far behind to have cameras that are linked up for recognition as well. Has this been an area that you have looked at in regards to housing, and are there safeguards in place?

Ms. RICE. Yes, Congressman. Thank you for the question, and one other area that we have particularly been focusing on is the use of facial recognition technology in the area of financial services. So, for example, more transactions have been happening in the virtual space, and there is certainly the opportunity to use facial recognition technology as a fraud detection mechanism, for example. So, yes, this is an area of deep and grave concern. It is one of the reasons why we have been calling for the building and development of more inclusive, robust datasets in many different areas. One of the ways that Joy Buolamwini and other data scientists were able to work with IBM, and Google, and Facebook, et cetera, to help them improve or lessen the discrimination on their systems was by building better training datasets.

Mr. AUCHINCLOSS. That was actually the second point I wanted to raise. You have been ahead of me this whole hearing. You had mentioned earlier in your comments the idea of synthetic data as a way to buttress training sets. My understanding for how the original facial recognition training sets were composed is that the faces were really scraped off of a lot of media sites and elsewhere, and they were pulling, it seems like, disproportionately White faces. Has there been work done, and maybe just describe more how those training sets have been fixed because, as you say, really the raw data is the core of undoing bias in the actual outcomes?

Ms. RICE. Yes, and I should have been more specific. I was sort of myopically focused on financial and housing services in terms of my reference to a synthetic dataset, publicly-available dataset, for research and education only. I don't think we should be building real systems and models using a lot of synthetic data, so I am sorry I didn't get a chance to make that distinction.

Mr. AUCHINCLOSS. Absolutely. Ms. Koide, maybe you could weigh in here as well about any oversight that you think is necessary for facial recognition technology.

Ms. KOIDE. Thank you for the question. We have been much more focused on tabular data, data that is being contemplated or used in credit underwriting. We have not been evaluating visual recognition data, but it is a great question.

Mr. AUCHINCLOSS. Understood. Yes, it is an area that we have been leaning into in Massachusetts and, I think, increasingly nationally just because, in some ways, the technology is both really good and really bad. Really good in the sense that it has been incredibly effective and has created some kind of compelling results in its accuracy, but very bad in the sense that these kinds of biases have snuck through in a way that, as Ms. Rice pointed out, were not identified for too long. So, it has been an area of concern for me both at the State and the Federal level, and I will yield back the balance of my time, Mr. Chairman.

Chairman FOSTER. Thank you, and I would like to thank all of our witnesses for their testimony today.

The Chair notes that some Members may have additional questions for this panel, which they may wish to submit in writing. Without objection, the hearing record will remain open for 5 legislative days for Members to submit written questions to these witnesses and to place their responses in the record. Also, without objection, Members will have 5 legislative days to submit extraneous materials to the Chair for inclusion in the record.

This hearing is now adjourned.

[Whereupon, at 1:24 p.m., the hearing was adjourned.]

A P P E N D I X

May 7, 2021

FSC Artificial Intelligence Task Force Hearing: "Equitable Algorithms: How Human-Centered AI can Address Systemic Racism and Racial Justice in Housing and Financial Services"

Statement for the Record from Representative Sylvia Garcia

Mr. Chairman, thank you for having us here to talk about this important issue. I regret not being able to attend in person, but I would like to convey my appreciation for you bringing this important issue the attention that it deserves. I hope that following this hearing, we can pursue informed legislative options to address this issue and facilitate equal access to credit.

I have said before that access to credit is the first step in building generational wealth. Access to consumer credit paves a path for borrowers to demonstrate creditworthiness and prepare themselves for one of the biggest financial transactions of their lives – purchasing a home. Thus, it is important that credit can be earned equally across all lines of business and demographics, and it's important that all consumers – regardless of race, ethnicity, gender identity, or sexual orientation – can build credit at the same rate, without paying higher prices despite equal risk.

Access to the middle class is through homeownership. Building equity allows wealth to accrue; when one sells their home, they are able to do things such as move to a more expensive home, send their children to college, invest in business ventures, or to simply utilize the freedom to diversify their investment portfolio and grow their wealth.

The use of technology in lending has democratized finance, and it has made banking products easier to use and more accessible. Credit underwriting from behind a screen has made it easier to avoid the face-to-face discrimination of the past, wherein borrowers were treated differently directly to their face. But discrimination has not simply disappeared, rather it has become enshrouded in data, with decisions automated by algorithms. Data transparency is crucial – if we don't understand how consumers are being evaluated for credit, we cannot ensure that it is fair.

While credit underwriting using data has many merits, the discriminatory barriers faced by consumers in face-to-face interaction are the same, with disparate pricing hidden behind obscure algorithms and unknown proxy variables. Algorithmic pricing has a significant impact on people of color in particular. A 2018 study by researchers at UC- Berkeley found that algorithmic underwriting resulted in Black and Latino borrowers paying an average 5.3 basis points higher interest rates in mortgage loans obtained online.[1]

It is critical that we work to ensure equitable credit access and pricing, and that we facilitate an economy that is fair to all. A report by the Urban Institute found that between 2020-2040, there will be 4.8 million more Hispanic homeowners. For market growth of this size, we as leaders must work to ensure that consumers are being treated fairly, with equitable pricing and transparency.

[1] Consumer-Lending Discrimination in the Era of FinTech. https://faculty.haas.berkeley.edu/morse/research/papers/discrim.pdf

Testimony of
Dave Girouard
CEO and Co-Founder, Upstart Network, Inc.

Before the
Task Force on Artificial Intelligence
United States House Committee on Financial Services

Hearing on Equitable Algorithms Revisited: How Human-Centered AI Can Address Systemic Racism and Racial Justice in Housing and Financial Services;

May 7, 2021

Chairman Foster, Ranking Member Gonzalez, and Members of the Task Force on Artificial Intelligence, thank you for the opportunity to participate in today's conversation.

My name is Dave Girouard, co-founder and CEO of Upstart, which is a leading artificial intelligence ("AI") lending platform headquartered in San Mateo, California and Columbus, Ohio. I founded Upstart more than 9 years ago in order to improve access to affordable credit through the application of modern technology and data science.

In the last seven years, our bank and credit union partners have originated more than $9 billion in high quality consumer loans using our technology, approximately half of which were made to low and moderate income borrowers. Our AI-based system combines billions of cells of training data with machine learning algorithms to more accurately determine an applicant's creditworthiness.

As a company focused entirely on improving access to affordable credit for the American consumer, fairness and inclusiveness are issues we care about deeply. The opportunity for AI-based lending to improve access to credit for the American consumer is dramatic. But equally dramatic is the opportunity to reduce disparities and inequities that exist in the traditional credit-scoring system.

In our early days at Upstart, we conducted a retroactive study with a large credit bureau and uncovered a jarring pair of statistics: just 45% of Americans have access to bank-quality credit, yet 83% of Americans have never actually defaulted on a loan. That's not what we would call fair lending.

The FICO score was introduced in 1989 and has since become the default way banks judge a loan applicant. But in reality, FICO is extremely limited in its ability to predict credit performance because it's narrow in scope and inherently backward looking. And as consumer protection groups such as the National Consumer Law Center have highlighted, for the past two decades, study after study has found that African American and Latino communities have lower credit scores as a group than White borrowers.[1]

At Upstart, we use modern technology and data science to find more ways to prove that consumers are indeed creditworthy - to bridge that "45% versus 83%" gap. We believe that consumers are more than their credit scores. And by going beyond the FICO score, and including a wide variety of other information such as a consumer's employment history and educational background, we've built a significantly more accurate and inclusive credit model.

While most people believe a more accurate credit model means saying "no" to more applicants, the truth is just the opposite. Accurately identifying the small fraction of borrowers who are unlikely to be able repay a loan is a better outcome for everyone. It leads to significantly higher approval rates and lower interest rates than a traditional model, especially for underserved demographic groups, such as Black and Hispanic applicants.

Since our early days, skeptics have asked whether our AI models would hold up in a down economy. The tragedy of the COVID pandemic, where unemployment rose from 4% to more than 14% in just a few weeks, required that we prove our mettle. And in fact we did just that; despite the elevated level of unemployment, the pandemic had no material impact on the performance of Upstart-powered loans held by our bank partners.

With the support of a more accurate credit model - powered by artificial intelligence - our bank and credit union partners can have the confidence to lend, regardless of the state of the economy. Just imagine - banks lending consistently and responsibly, just when credit is needed most. That's an outcome for which we can all cheer.

The concern that the use of AI in credit decisioning could replicate or even amplify human bias is well-founded. We have understood, since our inception, that strong consumer protection laws - including the Equal Credit Opportunity Act, help ensure that good intentions are actually matched by good outcomes. This is especially true when it comes to algorithmic lending. For these reasons and more, we proactively met with the appropriate regulator - the Consumer Financial Protection Bureau (CFPB) well before launching our company.

Quite simply, we decided to put independent human oversight into the equation. After significant good faith effort starting in late 2015 between Upstart and the CFPB to determine the proper way to measure bias in AI models, we demonstrated that our AI-driven model doesn't result in unlawful "disparate impact" against protected classes of consumers. Because AI models change

[1] https://www.nclc.org/images/pdf/credit_discrimination/Past_Imperfect050616.pdf

and improve over time, we developed automated tests with the regulator's input to test every single application on our platform for bias. And we provide the results of these tests to the Bureau on a quarterly basis.

In September 2017, we received the first "No-Action" letter from the CFPB, recognizing that Upstart's platform improves access to affordable credit without introducing unlawful bias. We have been reporting this information to the CFPB for the last three years, and we will continue to do so for the next three years under our latest "No-Action" letter.

Thus far we have been able to report to CFPB that our AI-based system has significantly improved access to affordable credit.

Specifically looking over the past three years of our access to credit results:

- **Overall, the Upstart model approves 32% more consumers and lowers interest rates (APRs) by almost 3.5 percentage points, compared to a traditional lending model**
- **For near-prime consumers (620-660 FICO), our model approves 86% more consumers and reduces their interest rates by more than 5 percentage points compared to a traditional model.**
- **Upstart's model also provides higher approval rates and lower interest rates for every traditionally underserved demographic. For example over the last three years, the Upstart model helped the banks that use Upstart approve 34% more Black borrowers than a traditional model would have, with 4 percentage point lower interest rates.**

That's the type of consumer benefit we should all get excited about.

We're proud of the proactive and transparent approach we've taken to working with regulators, because while we might be the first, we certainly won't be the last platform to leverage AI to improve credit outcomes. We have worked with the CFPB through what is now three different administrations. We've also realized that working with other stakeholders who care about fairness and consumer protection is important as well.

Late last year, we embarked on an effort to do fairness testing and analysis with the NAACP Legal Defense Fund and the Student Borrower Protection Center to ensure that we are identifying every possible opportunity to minimize the credit impacts of underlying inequality in society, while maintaining model accuracy and improving credit outcomes for all.

What has all of this yielded in terms of policy implications? In Upstart's experience, the fair lending laws enacted in the 1970s and the substance of fair lending regulation enforcement—that is, monitoring and testing the impact on actual consumers who apply for loans—translates very well to the AI-driven world of today.

But in reality, the path we have walked at Upstart is likely insufficient to create a robust and competitive market that will maximize the use of AI to promote financial inclusion and credit access. In our early days at Upstart, we couldn't know for certain whether our model would be biased. It wasn't until loans were originated that we were able to demonstrate that our platform was fair. As an early-stage startup, this was a risk worth taking, but it's not a risk that many large banks would have considered, and that's a problem.

If broader and deeper financial inclusion among American consumers is important to this Committee, it's worth considering rule-making or legislation that would provide a wide-scale, supervised opportunity for model development and testing, and a robust data collection effort in all areas of consumer lending.

By combining regulatory support for model innovation with rigorous and standardized testing, we can ensure that we don't forego the clear and obvious benefits that AI-enabled lending can offer to the American consumer.

Thank you.

Equitable Algorithms: How Human-Centered AI Can Address Systemic Racism and Racial Justice in Housing and Financial Services

Testimony of Stephen F. Hayes
Partner, Relman Colfax PLLC

Before the Task Force on Artificial Intelligence
United States House Committee on Financial Services

May 7, 2021

Chairman Foster, Ranking Member Gonzalez, and distinguished members of the Task Force on Artificial Intelligence, thank you for hosting this important hearing and for giving me the opportunity to submit this testimony.

My name is Stephen Hayes and I am a Partner at Relman Colfax PLLC. Relman Colfax is a national civil rights law firm. We have a litigation practice focused on combating discrimination in areas such as housing and lending, and we regularly represent individuals, non-profits, and municipalities bringing redlining, reverse redlining, and other civil rights claims. We also provide legal counsel to entities such as financial institutions, Internet-based companies, and nonprofits on civil rights compliance. My work focuses largely on providing fair lending and fair housing advice, including legal counsel on testing models for discrimination risks. I previously worked in the Legal Division of the Consumer Financial Protection Bureau ("CFPB"). I hope that my testimony today furthers the Committee's understanding of algorithms and discrimination, and the potential for using Artificial Intelligence and Machine Learning ("AI" and "ML") to increase opportunity, equity, and inclusiveness in financial markets.

A. Credit Discrimination, Alternative Data, and AI/ML

Our Nation's history of housing, employment, and credit discrimination has contributed to disparities in income, wealth, rates of home and small business ownership, and access to other important life opportunities. Existing credit markets reflect that history: studies have found evidence of racial disparities in credit scoring and factors on which traditional scores rely.[1] Disparities also exist with respect to populations without credit histories—those that are "credit invisible" or "thin file." Black and Hispanic Americans are more likely than white or Asian

[1] *See, e.g.*, Lisa Rice & Deidre Swesnik, *Discriminatory Effects of Credit Scoring on Communities of Color*, 46 Suffolk U. L. Rev. 935, 952-959 (2013); Aaron Klein, Brookings Institution, "Reducing bias in AI-based financial services" (July 10, 2020), https://www.brookings.edu/research/reducing-bias-in-ai-based-financial-services; Bd. of Governors of the Fed. Reserve Sys., Report to the Congress on Credit Scoring and Its Effects on the Availability and Affordability of Credit (2007); CFPB, "Analysis of Differences Between Consumer- and Creditor-Purchased Credit Scores" (Sept. 2012); National Consumer Law Center, Credit Discrimination § 6.4.1.1, "Studies Showing that Minorities Have Lower Credit Scores as a Group" (7th ed. 2018).

Americans to be credit invisible or to have unscored records.[2] These gaps are pernicious because it can be difficult to build a credit history without access to credit, limiting certain consumers' ability to improve their financial circumstances.

At the same time, many of these individuals may in fact be good credit risks. There is evidence that some "alternative data"—information not traditionally found in the credit files of the nationwide consumer reporting agencies and not commonly provided by consumers on credit applications—can help lenders make accurate underwriting and loan pricing decisions for these consumers.[3] Accordingly, some market participants supplement traditional lending strategies with alternative data, and regulatory agencies have worked to facilitate the responsible use of certain types of alternative data.[4]

Alternative data is distinct from, but often discussed in combination with, "alternative models" like ML models. The term "alternative models" can refer to methods for constructing predictive models from historical data without requiring human modelers to explicitly specify relationships among the variables that can be used in the model.[5] Like with alternative data, there is some evidence that the use of alternative models (and, in particular ML algorithms), has the potential to expand credit access by allowing lenders to evaluate the creditworthiness of consumers who are difficult to score accurately using traditional techniques.[6]

The use of alternative data and alternative models can also raise real concerns, including serious risks related to accuracy, completeness, explainability, validity, barriers to improving credit standing, and discrimination.[7] Moreover, fair lending concerns are not resolved solely because a practice increases access to credit; increases in access to even favorable credit or financial products can drive persistent inequality—and disparate impact—if distributed

---

[2] Kenneth P. Brevoort, et al., CFPB Office of Research, "Data Point: Credit Invisibles," at 6 (May 2015), https://files.consumerfinance.gov/f/201505_cfpb_data-point-credit-invisibles.pdf.

[3] *See, e.g.*, CFPB, Request for Information Regarding Use of Alternative Data and Modeling Techniques in the Credit Process, 82 Fed. Reg. 11183, 11185 (Feb. 21, 2017).

[4] *See, e.g.*, FinRegLab, "The Use of Cash-Flow Data in Underwriting Credit" at 8 (July 2019), https://finreglab.org/wp-content/uploads/2019/07/FRL_Research-Report_Final.pdf.; Bd. of Governors of the Fed. Rsrv. Sys., CFPB, FDIC, NCUA, OCC, "Interagency Statement on the Use of Alternative Data in Credit Underwriting," at 1 n.1 (Jan. 2019), https://www.federalreserve.gov/newsevents/pressreleases/files/bcreg20191203b1.pdf.

[5] *See* Nicholas Schmidt & Bryce Stephens, "An Introduction to Artificial Intelligence and Solutions to the Problems of Algorithmic Discrimination," Consumer Finance Law Quarterly Report, Vol. 72, No. 2, 131, at 133 (2019), https://arxiv.org/ftp/arxiv/papers/1911/1911.05755.pdf.

[6] Patrice Alexander Ficklin, et al., CFPB Blog, "Innovation spotlight: Providing adverse action notices when using AI/ML models" (July 7, 2020), https://www.consumerfinance.gov/about-us/blog/innovation-spotlight-providing-adverse-action-notices-when-using-ai-ml-models/; Fed. Rsrv. Bd. Governor Lael Brainard, *Supporting Responsible Use of AI and Equitable Outcomes in Financial Services*, Speech at the AI Academic Symposium hosted by the Bd. of Governors of the Fed. Rsrv. Sys. (Jan. 21, 2021), https://www.federalreserve.gov/newsevents/speech/brainard20210112a.htm.

[7] Brian Kreiswirth, et al., CFPB Blog, "Using alternative data to evaluate creditworthiness," (Feb. 16, 2017), https://www.consumerfinance.gov/about-us/blog/using-alternative-data-evaluate-creditworthiness/; *see also, e.g.*, Solon Barocas and Andrew Selbst, "Big Data's Disparate Impact," 104 Cal. L. Rev. 671 (June 2016); FinRegLab, "AI in Financial Services: Key Concepts" (May 2020), https://finreglab.org/wp-content/uploads/2020/05/FinRegLab_AIFAQ_Key-Concepts_AI-in-Financial-Services.pdf; Caroline Wang et al., "In Pursuit of Interpretable, Fair and Accurate Machine Learning for Criminal Recidivism Prediction" (May 2020), https://arxiv.org/abs/2005.04176.

unequally. That said, whether alternative data and alternative models will fairly increase access to credit in any given situation depends on a number of criteria, such as the type of alternative data and models at issue, how those tools are deployed, and characteristics of applicant pools.

B. Legal framework

Two primary federal antidiscrimination laws—the Equal Credit Opportunity Act ("ECOA") and the Fair Housing Act ("FHA")—prohibit institutions from discriminating in lending on the basis of characteristics such as race, national origin, religion, and sex.[8] ECOA applies to nearly all lending, including lending to businesses. The FHA applies to housing discrimination, including lending related to residential real estate transactions.

These laws prohibit intentional and overt discrimination, sometimes called "disparate treatment," as well as an unintentional type of discrimination called "disparate impact." Disparate treatment occurs when an entity explicitly or intentionally treats people differently based on prohibited factors, such as race, national origin, or sex. Unlike disparate treatment, disparate impact does not require any showing of intent to discriminate or that the protected characteristic was considered at all. Instead, the focus of disparate impact is on outcomes. Generally, unlawful disparate impact occurs when a (1) facially neutral policy or practice disproportionately adversely impacts members of protected classes, and either (2) the policy or practice does not advance a legitimate interest, or (3) is not the least discriminatory way to serve that interest.[9] These frameworks translate well to assessments of lending models, including AI/ML models.

C. Testing Models for Discrimination

Questions regarding how to ensure that algorithms are not discriminatory have received a significant amount of attention in recent years, particularly with the increased use of alternative data and alternative models. At the same time, these questions are not being written on a blank slate, either legally or with respect to institutions' internal compliance practices. Certain companies, including many financial services companies, have been testing models for discrimination for years and have systems in place guiding those assessments. Of course, even the most robust existing systems can be improved, and disparities in credit markets remain; although essential, improving model fairness alone will not solve those disparities. At the same time, existing programs demonstrate that model discrimination testing is both possible and effective, and—even if it may not make sense to incorporate these frameworks wholesale into other markets—these systems can serve as guides for markets where testing is nonexistent or nascent.

The methodologies that institutions use for fair lending testing their models vary, but as a general matter the most effective systems are designed to align with regulatory expectations and traditional principles gleaned from antidiscrimination jurisprudence. These systems often include: (1) ensuring that models do not include protected classes or close proxies for protected

---

[8] 15 U.S.C. § 1691(a); 12 C.F.R. § 1002.2(z); 42 U.S.C. § 3605.
[9] *See, e.g.*, 12 C.F.R. part 1002, Supp. I, ¶ 6(a)-2 (ECOA articulation); 24 C.F.R. § 100.500(c)(1) (FHA articulation); 42 U.S.C. § 2000e-2(k) (Title VII articulation).

classes, for example as variables or segmentations; and (2) assessing whether facially-neutral models are likely to disproportionately lead to negative outcomes for a protected class, and if such negative impacts exist, ensuring the models serve legitimate business needs and evaluating whether changes to the models would result in less of a disparate impact while maintaining model performance.[10]

This final step—identifying whether less discriminatory alternatives exist—is key. In the case of traditional statistical models, identifying less discriminatory alternatives often involves a process of adding, dropping, or substituting variables in the model, with the goal of identifying variations that maintain reasonable performance but that have less disparate impact on protected classes.[11] Newer methods exist that can improve upon that process for ML models. Advancements in computing power, along with sophisticated algorithms, can help analyze the impact of many different sub-combinations of variables, which allows institutions to explore numerous iterations of variable combinations and adjustments to hyperparameter settings. Other techniques also exist, such as training models to optimize for performance and metrics of fairness. In short, not only can these searches work for ML models, they can be more effective than traditional methods because there are more options for model adjustments.

As noted, disparities in the current credit system are stark. That said, disparate-impact testing—including the adoption of less discriminatory alternatives—has proven critical in reducing credit inequalities. It has caused lenders to search for and implement model variations that predict accurately and reduce disparate outcomes. This process can benefit both businesses and consumers. Among other things, identifying less discriminatory practices can help businesses responsibly diversify their borrower pools, which can lead to more representative training samples and increases in access to credit over time. Less discriminatory models also help mitigate disparities, counteract the legacies of historic credit discrimination, and close unnecessary credit gaps.

In robust programs, these quantitative statistical tests are paired with more holistic compliance measures: fair lending training for relevant staff, including modelers; ensuring teams have diverse backgrounds and are empowered to identify and remedy issues; reviewing policies and procedures within which models operate; and assessing areas of discretion to ensure that the potential for judgmental bias is mitigated.

D. Why Do Some Institutions Test Models But Others Do Not?

Although some companies routinely test their models for discrimination risks, many do not. Several legal and structural characteristics contribute to this unevenness.

---

[10] *See* Initial Report of the Independent Monitor, Fair Lending Monitorship of Upstart Network's Lending Model at 7 (April 14, 2021) ("Initial Upstart Report"), https://www.relmanlaw.com/media/cases/1086_Upstart%20Initial%20Report%20-%20Final.pdf; Nicholas Schmidt & Bryce Stephens, "An Introduction to Artificial Intelligence and Solutions to the Problems of Algorithmic Discrimination," Consumer Finance Law Quarterly Report, Vol. 72, No. 2, 131, at 141–142 (2019), https://arxiv.org/ftp/arxiv/papers/1911/1911.05755.pdf; David Skanderson & Dubravka Ritter, Federal Reserve Bank of Philadelphia, "Fair Lending Analysis of Credit Cards" 38–40 (2014).

[11] *See* Initial Upstart Report at 11.

First, agency supervision spurs internal testing. Many lenders are supervised by the CFPB, and have a long history of supervision by banking regulators such as the Federal Deposit Insurance Corporation, Office of the Comptroller of the Currency, the Federal Reserve Board, and the National Credit Union Administration. Supervisory expectations with respect to models have not always been consistent, but in general supervisory authority means that lenders' models can be (and at times have been) subject to scrutiny and therefore are expected to be fair and transparent in ways that are not true for non-supervised companies.

Second, lenders have been on notice for years that, via ECOA and the FHA, disparate treatment and disparate impact apply to their credit models, and so many lenders have developed methodologies to address these risks.[12] However, discrimination and disparate impact is not clearly prohibited in all markets.[13] For example, financial regulatory agencies focus on credit discrimination but historically have not regulated discrimination related to other core consumer financial activities like acquiring checking accounts, credit reporting, or third-party debt collection—all areas in which models may be used. To the extent some antidiscrimination laws explicitly apply in these areas, they are generally not enforced by federal agencies and prohibit only disparate treatment. Disparate treatment, standing alone, is unlikely to ensure that models are non-discriminatory. Simply removing explicit protected class information from models will not eliminate bias and discrimination. Disparate impact, on the other hand, has historically proven effective for increasing equitable access to credit and housing.

Third, general regulatory model risk management ("MRM") expectations that are not directly related to discrimination can nonetheless complement and encourage appropriate fair lending model compliance.[14] MRM principles are articulated through guidance meant to help supervised banks avoid adverse consequences like financial loss and safety and soundness risks that can occur because of inaccurate or misused models. These MRM expectations foster responsible model development, accuracy, validation, use, and monitoring. They can facilitate fair lending testing because, among other reasons, they can help create an "audit trail" for models. Effective MRM programs ensure that modelers catalog models, assess the representativeness of data, validate that models work across populations, ensure that models are only used as intended, and establish a routine cadence for reviewing models.

These characteristics make it more likely that supervised entities providing services like credit will have systems for addressing discrimination risks arising from model usage. Importing similar principals and methodologies to model use in other markets could help advance equity and ensure models do not perpetuate historic disparities unnecessarily. At the same time, there is significant variation with respect to testing models for discrimination even within supervised financial institutions. This variation could be addressed through clear regulatory expectations, a

[12] HUD, DOJ, OCC, OTS, Fed. Rsrv. Bd., FDIC, FHFB, FTC, NCUA, Policy Statement on Discrimination in Lending, 59 Fed. Reg. 18266 (Apr. 15, 1994) ("Joint Policy Statement on Lending Discrimination"); HUD Final rule, Implementation of the Fair Housing Act's Discriminatory Effects Standard, 78 Fed. Reg. 114460, 11476 (Feb. 15, 2013).

[13] Stephen Hayes & Kali Schellenberg, "Discrimination is 'Unfair,'" Student Borrower Protection Center at 10-13 (April 2021), https://protectborrowers.org/wp-content/uploads/2021/04/Discrimination_is_Unfair.pdf.

[14] *See, e.g.*, Bd. of Governors of the Fed. Rsrv. Sys., OCC, SR Letter 11-7, "Supervisory Guidance on Model Risk Management" (Apr. 4, 2011).

need for which has become increasingly important given the rising interest in and use of alternative data and alternative models.

E. Moving Forward to Ensure Alternative Model Use is Fair and Equitable

Policymakers can take concrete steps to ensure models, including AI/ML models, are built in a transparent and accountable manner and result in fair and equitable outcomes. If appropriately tested, these models can serve as important tools to spur innovation, improve customer experiences, promote inclusiveness, and help overcome historic disparities experienced by protected classes.

1. <u>Regulatory agencies should routinely test models for discrimination, including assessing disparate impact and identifying less discriminatory alternatives.</u>

There is an uneven landscape with respect to how or whether institutions assess their models for discrimination. The CFPB and other agencies with supervisory authority can promote uniform internal testing by making clear that the agencies will review institutions' model testing results. Where internal testing is insufficient, the agencies should conduct independent testing, including assessing whether less discriminatory alternatives exist.

2. <u>Regulatory agencies should announce the methodologies they use to test models for discrimination.</u>

Even among financial institutions that conduct rigorous fair model testing, questions exist as to acceptable methodologies. This leaves institutions in a precarious situation, with some dedicating resources towards compliance without a clear picture of regulatory expectations. Institutions might also be overly cautious about using promising alternative data or modeling techniques that could benefit consumers for fear of regulatory risk. The CFPB should explain what methodologies it will use in supervision and enforcement so that entities can align their internal systems accordingly. The CFPB should also tailor MRM-like guidance specifically to fair lending assessments, and encourage more widespread adoption of these techniques by entities beyond those that are directly supervised by the federal banking agencies.

3. <u>Regulatory agencies should clarify that discrimination, including unnecessary disparate impact, is illegal across markets. To the extent ambiguities exist, Congress can explicitly codify coverage.</u>

Regulatory agencies should articulate that discrimination—including unnecessary disparate impact—is prohibited across markets outside of just credit and housing. This clarification would spur internal antidiscrimination measures, including testing, in areas that have historically been unregulated. Agencies like the CFPB and FTC currently have the tools to regulate many of these activities.[15] To the extent there are statutory authority concerns, Congress should consider explicitly codifying coverage.

---

[15] Hayes & Schellenberg, "Discrimination is 'Unfair,'" at 20-21.

FinRegLab

House Financial Services AI Task Force

"Equitable Algorithms: How Human-Centered AI Can Address

Systemic Racism and Racial Justice in Housing and Financial Services"

Melissa Koide, FinRegLab

Good afternoon. Thank you for the opportunity to testify before you today at the AI Task Force hearing, "Equitable Algorithms: How Human-Centered AI Can Address Systemic Racism and Racial Justice in Housing and Financial Services."

I am the founder and CEO of FinRegLab. FinRegLab is a DC-based independent, nonpartisan research organization that evaluates the use of new technologies and data to drive the financial services sector toward a responsible and inclusive marketplace. Through our research and policy discourse, we facilitate collaboration across the financial ecosystem to inform public policy and market practices.

FinRegLab has focused much of our work on the use of alternative financial data and machine learning algorithms in credit underwriting because credit plays such a critical role in borrowers' long-term financial health and economic participation. Credit can not only help bridge short-term gaps, but fund long-term investments in housing, transportation, education, and small business formation. The credit system thus both reflects and influences the ability of families, small businesses, and communities to participate in the broader economy.

The initial shift in consumer credit markets from subjective decision-making toward greater use of data and automated underwriting began more than 50 years ago. Research suggests that these changes have tended to lower underwriting costs and default losses, improve consistency of treatment of similarly situated applicants, and increase competition for borrowers.[1] Yet for all of these benefits, traditional scoring methods and underwriting models have limitations because they are dependent on data which is often not available for people in marginalized communities. We know, for instance, that about 20 percent of U.S. adults lack sufficient credit history to be

[1] Board of Governors of the Federal Reserve System, Report to Congress on Credit Scoring and Its Effects on the Availability and Affordability of Credit S-2 to S-4, O-2 to O-4, 32-49 (2007); Allen N. Berger & W. Scott Frame, Small Business Credit Scoring and Credit Availability, 47 J. of Small Bus. Mgmt. 5 (2007); Susan Wharton Gates *et al*., Automated Underwriting in Mortgage Lending: Good News for the Underserved?, 13 Housing Policy Debate 369 (2002); FinRegLab, The Use of Cash-Flow Data in Underwriting Credit: Market Context & Policy Analysis 11 n.16 (2020).

scored under the most widely used models.[2] Prior to the pandemic, another 30 percent may have struggled to access affordable credit because their scores were considered to be "non-prime."[3] And small business owners often struggle to access commercial credit in part because of information gaps that discourage banks from serving the small end of the market.[4]

In each of these cases, communities of color and low-income populations are substantially more likely to be affected by these information barriers than other applicants. For example, nearly 30 percent of African-Americans and Hispanics cannot be scored using the most widely adopted credit scoring models, compared to about 16 percent of whites and Asians. Racial disparities regarding access to credit are far greater than for more basic transaction accounts, for instance.[5]

FinRegLab's first empirical evaluation focused on the use of cash-flow data from bank accounts and other sources, concluding that the data have substantial potential to increase inclusion in consumer and small business credit markets.[6] We have recently announced a new project that is studying the market and policy questions raised with the use of machine learning underwriting models in consumer credit. In particular, we are empirically evaluating the capability and

---

[2] Consumer Financial Protection Bureau, Data Point, Credit Invisibles 4-6, 17 (2015); FinRegLab, The Use of Cash-Flow Data in Underwriting Credit: Market Context & Policy Analysis § 2.2.

[3] In lower score bands, the majority of applicants may be likely to repay but lenders cannot determine which particular applicants are lower risk without additional information. Lenders may choose not to lend to that cohort or may impose higher prices because default risks for the group as a whole are relatively high. For instance, depending on interest rates, consumers with scores near the typical minimums for approval may pay $7500 more over the life of a $20,000 auto loan and $86,000 more over the life of a $250,000 mortgage than peers with high scores. FinRegLab, Market Context & Policy Analysis §§ 2.1, 2.2; Lyle Daly, Here's How Much Money Bad Credit Will Really Cost You, The Ascent (Apr. 8, 2019). Since the pandemic, average credit scores have risen due to changes in reporting practices, the effect of stimulus payments, and other factors, but there is also some evidence that lenders are relying on them less due to uncertainty. FinRegLab, Research Brief, Covid-19 Credit Reporting and Scoring Update 2 & nn. 7, 10 (2020); Elisabeth Buchwald, A Pandemic Paradox: Americans' Credit Scores Continue to Rise as Economy Struggles — Here's Why, MarketWatch (updated Feb. 20, 2021).

[4] FinRegLab,The Use of Cash-Flow Data in Underwriting Credit: Small Business Spotlight §§ 2.1, 2.2 (2019).

[5] For example, a 2017 Federal Deposit Insurance Corporation survey found that about 10% of black and Hispanic households lacked bank and/or prepaid accounts, while more than 30% of both groups reported not having mainstream credit accounts of the type that are likely to be reported to credit bureaus. FinRegLab, Market Context & Policy Analysis § 2.2; FDIC, 2017 National Survey of Unbanked and Underbanked Households (2018).

[6] Specifically, the study analyzed data from six companies to evaluate the potential effects of cash-flow information on predictiveness, inclusion, and fair lending. The results suggested that cash-flow information could not only be used to predict default risk in situations in which traditional credit report information is not available, but that it also added somewhat different insights with regard to borrowers who did have traditional credit reports and scores. The analysis also found evidence that the participating companies were extending credit to applicants who may have faced constraints in accessing credit historically, and that the degree to which the information was predictive of credit risk appeared to be relatively consistent across borrowers who likely belong to different demographic groups. FinRegLab, The Use of Cash-Flow Data in Underwriting Credit: Empirical Research Findings (2019).

performance of tools to explain and manage machine learning underwriting models with respect to reliability, fairness, and transparency, among other concerns.

Our work directly intersects with the Task Force's inquiry into ways to harness the power of AI to increase opportunity, equity, and inclusiveness for those who have been historically disadvantaged through traditional lending models and analyses. Our research on machine learning and alternative financial data is designed to address many of the concerns that this Task Force has focused on in this hearing and elsewhere. In particular, we believe our latest project will yield important empirical and analytical insights about the tools available to financial services firms and regulators to diagnose and manage machine learning underwriting models, enhance market intelligence on the state of adoption in credit underwriting, and inform policy analyses about the need to adjust market practices and regulatory frameworks to promote efficiency, fairness, and inclusion when using machine learning models.

While our projects focus separately on the use of non-traditional data and more advanced analytical models to isolate and understand better the effects of each change, lenders often consider data and models in tandem to improve their ability to predict credit risk, including among minority populations. My testimony today will therefore discuss both the use of more advanced analytical models and non-traditional data, since the risks and rewards often intersect and overlap.

The Promise and Risks of AI and Machine Learning for Financial Inclusion

Financial services providers have begun to use machine learning models in a variety of business and operational contexts because they offer potential increases in the accuracy of predictions relative to incumbent models that have been used for decades. Depending on the context, such predictiveness improvements may allow providers to reduce losses due to default and fraud, cut processing times and costs, tailor their products, and/or expand their customer bases.

Such improvements could potentially benefit consumers and small businesses in a variety of ways, but could be transformational where information gaps and other obstacles currently increase the cost or risk of serving particular populations using traditional models and data. For example, more nuanced models have the potential to assess default risks among consumers who lack extensive credit history and relatively new small businesses. Machine learning may also help analyze changes in economic conditions and detect more quickly nuanced signs of improvement in the financial capabilities of millions of people and small businesses, which may be especially useful for post-Covid recovery.

But the greater predictive power of machine learning models can increase risks as well as potential benefits, due to the models' greater complexity and to their potential to exacerbate historical disparities and other flaws in underlying data. Because the models are more complex, they are often more difficult for lenders and regulators to understand, adjust, and monitor over time. These concerns are especially important for two reasons. First, AI and machine learning models can be brittle because they have been overfitted to the data used in initial development and testing. This means that their performance can deteriorate when the conditions in which they are used differ from training and testing. Second, AI and machine learning models might amplify patterns of historical discrimination and financial exclusion due to reliance on flawed data or mistakes made in development and deployment. The greater complexity also makes it more challenging to explain to individual applicants why they were rejected or charged higher prices, and how they might improve their risk profiles over time.

Publicly available research supports the general predictiveness benefits of machine learning models, but provides only limited insights on these more complicated questions about reliability, fairness, and inclusion effects. For example, multiple academic studies have found substantial predictiveness gains from machine learning models relative to conventional credit card algorithms,[7] and a survey of five recent studies on the use of AI and machine learning in commercial lending reported gains of 2-3 percent on average, with one study reporting gains over 15 percent and another reporting 3-4 percent gains independent of alternative data.[8] But limited information on the effects of machine learning underwriting models for populations that have historically struggled to access affordable credit creates a more complicated picture:

- VantageScore reports that its use of machine learning to develop scorecard models for consumers who are not scorable under some third-party models because their credit histories have not had an update in the prior six months resulted in a performance improvement of 16.6 percent for bank card originations and 12.5 percent improvement for auto loan originations.[9]
- An academic study of machine learning models using conventional data in the mortgage context concluded that such models would likely lead to modest improvements in application approvals among black and Hispanic applicants, but would increase pricing

[7] Amir E. Khandani et al., Consumer Credit Risk Models via Machine-Learning Algorithms (2010); Florentin Butaru, *et al.*, Risk and Risk Management in the Credit Card Industry, J. of Banking and Finance (2016); Anastasios Petropoulos *et al.,* A Robust Machine Learning Approach for Credit Risk Analysis of Large Loan Level Datasets Using Deep Learning and Extreme Gradient Boosting, Bank for International Settlements (2018).
[8] Dinesh Bacham & Janet Zhao, Building AI in Credit Risk: A Commercial Lending Perspective, Moody's Analytics Risk Perspectives at fig. 6 (July 2017).
[9] VantageScore, Our Models (undated), https://vantagescore.com/lenders/our-models#vantage-score-4.

differentials between different demographic groups due to many minority applicants being evaluated as higher risk than under conventional approaches.[10]

- Another academic study shows that credit scores for minority groups generally reflect significantly more signal noise than other potential borrowers due to thin credit files, which may undercut inclusion effects of using machine learning models with traditional credit data.[11]

Although more research is needed, this suggests that the known inclusion benefits of using alternative financial data, such as cash flow data,[12] may be enhanced when that data is used to develop machine learning underwriting models.[13]

FinRegLab decided to launch its new research project on machine learning in credit underwriting specifically because of this heightened potential for both benefits and risks relative to traditional models and data. We also believed that credit underwriting could serve as a particularly useful case study with regard to questions about AI adoption in financial services and other contexts more generally because credit decisions are subject to a relatively robust set of federal regulatory requirements concerning model reliability, fairness, and transparency. These existing requirements provide a useful starting point for evaluating algorithmic decision-making and available tools for managing these models, although they themselves may need to be adjusted in response to evolution in modeling techniques and data. Our research is designed to provide the first empirical data measuring available model diagnostic tools against these requirements. This research will help inform lenders' decisions about whether they can trust machine learning underwriting models and policymakers' decisions about how protections and oversight processes need to be adapted to foster fair and responsible use of machine learning underwriting models.

The balance of my testimony will provide an overview of the state of machine learning adoption for credit underwriting, the core conceptual and regulatory questions that are creating uncertainty regarding its use, and our research plan. I will conclude with some broader thoughts about the way that adoption of machine learning and non-traditional data in credit underwriting fits into the broader quest to make the U.S. financial system and economy more inclusive and fairer for all participants.

[10] Andreas Fuster et al., Predictably Unequal? The Effects of Machine Learning on Credit Markets (Oct. 2020).
[11] Laura Blattner & Scott Nelson, How Costly Is Noise? Data and Disparities in the U.S. Mortgage Market (Jan. 2021).
[12] FinRegLab, The Use of Cash-Flow Data in Underwriting Credit: Empirical Research Findings.
[13] Blattner & Nelson, How Costly is Noise? (suggesting that the combination of machine learning underwriting models and alternative data, such as cash flow data, is required for greater financial inclusion).

The State of Adoption of Machine Learning Credit Underwriting Models

Machine learning can be used in a variety of ways in the credit context that could have implications for fairness and inclusion, including marketing, customer onboarding, fraud and illicit activities detection, originations, servicing, and collections. FinRegLab's immediate focus is on the use of machine learning models in underwriting, since such activities are at the heart of the lending process and are subject to the most federal regulatory scrutiny.

Thus, as a precursor to our empirical research on this use of AI and machine learning, we are conducting outreach to financial services stakeholders to understand the state of adoption of machine learning underwriting models and areas of greatest concern and uncertainty for stakeholders. This outreach suggests that some firms are only using machine learning models to help them develop traditional logistic regression models and scorecards by using them to identify variables and relationships that are particularly predictive of credit risk. Other firms, however, are beginning to explore using machine learning algorithms directly in their underwriting models to evaluate individual applications. With regard to this latter group:

- Banks and nonbank lenders are interested in using machine learning underwriting models to make credit decisions due to their potential to improve the accuracy of credit risk assessment and reduce losses, to speed up the process of updating and refitting models, and to keep pace with market competitors. Many also cite the ability of machine learning models to leverage large, diverse data sets as a motivation. Nonbank usage is likely more established due to a number of factors, including reliance on digital business models, newer lending platforms, and differences in the nature and maturity of risk management processes.
- Credit cards and unsecured personal loans are the asset classes in which use of machine learning models to make credit decisions is most advanced. This reflects the historical position of credit cards as being at the analytical forefront of consumer finance and the dominance of digital lending in unsecured personal loans. Auto lending and small business lending are also areas where machine learning underwriting models are in use.
- Individual decisions about whether to use machine learning models to make credit decisions, what forms of machine learning to use, and how to enable appropriate oversight of such models varies based on firm culture and strategy and competitive dynamics in specific asset classes.

- Forms of machine learning used to make credit decisions range from gradient boosted trees and neural networks to ensembles combining multiple machine learning models.[14] Firms frequently introduce constraints to improve model transparency, even if those constraints impose performance tradeoffs. Common constraints include:
  - **Monotonicity constraints:** these constraints make it easier to understand the relationship between input data and predictions by ensuring one-directional relationships between the two;
  - **Sparsity terms:** these terms limit the number of features that a model uses to make a prediction; and
  - **Simulatability:** this approach renders the operation of a complex model in a list of if-then rules in order to make it easy to understand how a decision was reached.
- Decisions about whether to develop machine learning underwriting models in-house or to rely on third-party service providers are most likely to depend on the overall size of the lender and the importance of specific consumer asset classes to the institution. Many firms are likely to lack the resources – foremost among them personnel with appropriate data science and credit expertise – to develop and operate such models on their own.[15] To meet this need, a number of potential third-party providers have entered this market, including score providers, technology firms, and consulting firms.

Prior to the pandemic, surveys of industry executives conducted by other organizations found that the respondents viewed AI/machine learning as a "major differentiator" in their businesses, and about half of the participating institutions lacked AI and machine learning capabilities in some or all of their platforms.[16] The pandemic has likely accelerated interest in and use of machine learning for making credit decisions, as it has accelerated adoption of all forms of AI across financial services and in other industries.[17] Nevertheless, our research suggests that broad conceptual questions about the trustworthiness of AI models, uncertainty about the net performance benefits of operating AI models, concerns about compliance with existing federal regulatory requirements, and policy initiatives in other jurisdictions are shaping many firms' decisions in whether and how to use machine learning in connection with predicting default risk.

---

[14] Office of the Comptroller of the Currency, Credit Card Lending Handbook Version 2.0 at 17 (April 2021).
[15] Across both large and small institutions, approximately 20% of institutions have no in-house staff for credit modeling and rely on third parties to conduct such activities. Even large institutions with credit modeling teams do not devote significant resources, as just 16% of large institutions had four or more full time modelers. Cornerstone Advisors, Credit Monitoring and the Need for Speed: The Case for Advanced Technologies at 4, fig. 4 (Q2 2020).
[16] Leslie Parrish, Risky Business: The State of Play for Risk Executives in the Analytics Ecosystem, Aite at 14, fig. 9 (2019).
[17] KPMG, Report: Thriving in an AI World at 6-7 (April 2021).

Core Questions Regarding the Use of Machine Learning Underwriting Models

Our ability to realize the potential of AI and machine learning to significantly enhance the efficiency, fairness, and inclusiveness of credit decisions depends on resolving uncertainty about when we can trust a specific model and when we cannot. Some machine learning underwriting models are by their nature more transparent than others, although a range of techniques has emerged that are designed to explain some of the more opaque models. In the context of credit decisions, this means we need to understand how a model used in credit underwriting makes decisions and reaches particular outcomes such as denying certain credit applicants or charging those applicants more for their loan. Specifically, we need to be able to assess the model's stability and robustness, to provide an explanation for why a particular credit decision was made and a particular price was offered, and to understand if the model is treating people fairly. In this context, the significant jump in financial services leaders who report AI is being adopted "too fast for comfort" warrants close attention.[18]

Questions about the trustworthiness of AI and machine learning models pose a core challenge for all stakeholders in consumer and small business lending: how to enhance our ability to understand and rely on these technologies without unnecessarily diluting their ability to improve predictive power. This challenge is all the more complex in the financial sector, where extensive policy frameworks force consideration of questions about machine learning's trustworthiness more holistically and at an earlier stage than may occur elsewhere. Indeed, implementing AI in the financial system often requires meeting exacting requirements focused on securing the financial system from illicit activity, promoting responsible risk-taking, and providing broad, non-discriminatory access to credit.

There are a number of key issues that need to be addressed by policymakers and firms to foster responsible adoption and use of machine learning for credit underwriting:

Reliability

Some firms continue to explore whether machine learning models that meet governance and compliance requirements deliver sufficient performance gains in the short- and the long-term to make the operational costs of implementing and operating those models worthwhile. Others report significant improvement in lending outcomes. One key area of concern is our ability to understand and manage what happens to a model's predictions when data conditions in deployment differ from the data on which the model was trained. Current machine learning technologies are not generally well equipped for responding to such changes and may not even

[18] This study reported a 20% increase in financial services executives saying adoption of AI is moving too fast for comfort between 2020 and 2021. *Id.* at 8, chart 3.

do well in recognizing them. Emerging tools to monitor this drift in the quality of a model's predictions have not been independently evaluated in the context of financial services.

Transparency

Lenders, as well as academics and other stakeholders, are debating whether certain "black box" technologies are ever appropriate for credit underwriting or lenders should only use interpretable or self-explanatory models.[19] Surveyed industry executives reported that explaining model results, and specifically adverse actions, was the most significant challenge for using AI and machine learning in decisioning applications for credit.[20]

The need to respond to a range of specific regulatory requirements, such as demonstrating general model performance and explaining specific loan decisions, requires different kinds of information and has led to the development of an array of explainability techniques. However, these techniques – which often involve the use of secondary machine learning models to explain the underwriting model – introduces a second layer of questions and concerns about the trustworthiness of information provided about the model's decision-making. We must ask whether we can trust the information produced by the secondary model and whether that information serves important oversight and governance needs.[21] Further, these explainability techniques are relatively new and are rapidly evolving. That means there is little independent evidence on their capability and performance in applications like credit underwriting where the stakes for consumers, communities, and firms are high and where exacting legal and regulatory requirements apply. As a result, there is considerable interest in understanding and evaluating the extent to which lenders have appropriate tools to explain machine learning models – from the simplest to the most opaque and complex – with sufficient confidence to satisfy themselves, their regulators, and potential critics.

Fairness

Ensuring that machine learning underwriting models do more than simply replicating historical lending patterns depends heavily on two important factors: (1) the articulation of expectations about what constitutes fair and responsible conduct for those who develop, operate, and

---

[19] Cynthia Rudin, Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead, Nature Machine Intelligence (Sep. 2019).
[20] Leslie Parrish, Alternative Data and Advanced Analytics: Table Stakes for Unsecured Personal Loans, Aite at 16, fig. 12 (Nov. 2019).
[21] Agus Sudjianto, What We Need Is Interpretable and Not Explainable Machine Learning, presentation at Cogilytica Machine Learning Lifecycle Conference (Jan. 2021), available at https://events.cognilytica.com/wp-content/uploads/2020/12/What-we-need-is-interpretable-and-not-explainable-machine-learning-Agus-Sudjianto-ML-Lifecycle-Slides-.pdf.

monitor such technologies and (2) the choices that individual firms make in the model development and data selection processes, how they test and monitor model performance, and what customers their business and product strategies aim to serve. The efficacy of emerging approaches to adapting technologies to minimize bias throughout the model lifecycle need to be better understood. Those approaches include data diversification, reweighting or preprocessing data to reduce bias; improving techniques to manage bias in model training, such as adversarial debiasing or learning fair representation; and enhancing reject option classification and other post-processing techniques to reduce discrimination.

More broadly, stakeholders are currently debating what fairness requires in the context of lending, and many are proposing alternative methods for defining, measuring, and evaluating fairness. For example, the advent of AI and ML use more broadly has driven data scientists to proliferate metrics for measuring fairness,[22] and some academics and advocates have begun to ask if including protected class information in underwriting models might actually improve credit access for those in protected classes. At the same time, there is also a recognition that fairness is not just a mathematical problem and that it is one area where policy processes are needed to promulgate broadly applicable approaches to what fairness should mean. In considering this question, as well as others on model transparency, it is likely that the standards that emerge will apply even where machine learning models are not used and can thus drive the broader financial system to enhanced fairness and inclusiveness.

Privacy

The potential for machine learning underwriting models to analyze large, diverse data sets raises significant questions about privacy and consumer data controls and protections. Those questions include what data can be used for underwriting, how data are acquired, what constitutes informed consumer consent to a lender's acquisition and use of data, what constraints exist on firms' ability to use and retain that information, and whether and in what circumstances certain applicants should have to provide more access to information to facilitate credit risk assessment.

These broad conceptual questions are also implicated in a number of specific existing federal regulatory requirements that apply to credit underwriting activities regardless of the type of underwriting model that is used. In particular, areas of uncertainty about the application of lending-specific laws and regulations include:

- **Fair lending:** Concern about the state of our ability to identify and control appropriately proxies for protected class information is paramount for firms using machine learning for

[22] Sahil Verma & Julia Rubin, Fairness Definitions Explained, 2018 IEEE/ACM International Workshop on Software Fairness (FairWare), IEEE at 1-7 (May 2018).

credit underwriting and those considering such use. This concern is heightened in the context of more complex models and larger, more diverse data sets and includes consideration of whether traditional fair lending analysis and the prohibition on protected class data are suitable for use in the context of machine learning underwriting models.[23]

- **Adverse action reporting:** Required disclosures for applicants whose applications are denied or granted on less favorable terms based on information in a credit report make it necessary for lenders to be able to identify up to four primary bases of their decision. In models with higher numbers of variables and features, as well as models with more complex structures, producing accurate and reliable information for these notices can be a challenge.[24]
- **Model risk management:** The principles-based framework that governs banks' use of models requires model users to demonstrate the conceptual soundness and fitness for purpose of their models and to implement appropriate oversight and risk management measures. Some firms have reported navigating review of machine learning underwriting models in their model risk management programs, applying risk-based controls and governance appropriate to the type of model being used, the product being underwritten, and the customer base being served.

<u>**FinRegLab's Research on the Explainability and Fairness of Machine Learning Underwriting Models**</u>

FinRegLab has begun empirical and qualitative research to document how machine learning is being used for credit underwriting and evaluate available model diagnostic tools designed to improve the transparency of complex models. In partnership with researchers from the Stanford Graduate School of Business, we are empirically evaluating the performance and capabilities of emerging tools designed to help lenders develop, monitor, and manage machine learning underwriting models. FinRegLab is also conducting research on the market and policy implications of the use of machine learning underwriting models. This will be the first public research shaped by input from key stakeholders – including bank and fintech executives, data scientists, and consumer advocates – to address the questions about explainability and fairness

[23] Talia Gillis, The Input Fallacy, Minn. L. Rev. (forthcoming 2022), April 2021 draft available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3571266.

[24] Leslie Parrish, Alternative Data and Advanced Analytics: Table Stakes for Unsecured Personal Loans, Aite at 16, fig. 12 (Nov. 2019). Surveyed industry executives cited explaining model decisions and adverse actions in particular as the most important challenge for using AI and machine learning.

that can promote responsible, fair, and inclusive adoption and use of machine learning in consumer credit.

Over the past year, FinRegLab has engaged these stakeholders to help inform our empirical assessment of a set of proprietary and open-source model diagnostic tools. We are evaluating how well these tools help lenders using machine learning models:

- Demonstrate the conceptual soundness, performance, and reliability of the models at the portfolio or lender level to satisfy prudential regulators and investors;
- Identify, measure, and enable mitigation of fair lending risks, particularly whether models have a disparate impact on protected classes; and
- Provide applicants with individualized "adverse action" notices explaining why they were denied credit or offered less favorable terms where required by law and regulation.

These represent a set of diverse requirements that apply to consumer lending regardless of the type of model being used to make credit decisions. Each one focuses attention on important threshold questions of model transparency related to the shift from incumbent automated underwriting models to machine learning models.

Drawing on nationally representative traditional credit data, we will evaluate the performance and capabilities of a set of model diagnostic tools using benchmark underwriting models developed by the research team. We expect to investigate proprietary tools provided by Fiddler Labs; H2O.ai, Relational AI, SolasAI/BLDS, LLC; Stratyfy, and Zest AI, as well as open source tools such as measures of feature importance, surrogate models, and plots of trained model predictions. We will assess the model diagnostic tools across a variety of dimensions:

- **Type of machine learning model:** benchmark underwriting models will range from logistic regression and boosted trees to neural networks and ensemble models to identify whether the type of underwriting model being explained affects the accuracy and utility of information produced by the model diagnostic tools;
- **Model complexity:** each form of machine learning being evaluated will have simple and complex forms to help us identify the tradeoffs, if any, between performance and transparency and between performance and fairness;
- **Changes in economic conditions:** test data sets will simulate different economic environments, such as data from 2009-2010, to help assess whether the model diagnostic

tools can help lenders identify changes in data conditions and model performance once in operation; and

- **Shifts in applicant distribution:** test data sets will encompass different kinds of borrowers with respect to geographic location and socio-economic status to help us evaluate how well these tools detect fair lending and other risks.

Our set of benchmark models have generally been designed to approximate the models that lenders might use to estimate the risk of default associated with an application. In this evaluation, we will assess how a set of alternative definitions of algorithmic fairness that have emerged in academic literature work in the context of the underwriting models and model diagnostic tools used in our research.

In addition to empirical findings, we expect to put forth a framework that will help all stakeholders – model developers, risk and compliance personnel, and regulators – assess the accuracy and utility of accessible information about a machine learning underwriting model's decision-making. This framework will provide a substantive, thoughtful contribution to the current oversight approach about model transparency – defining the questions we should all be asking about the information that currently available model diagnostic tools produce. Those questions will help us assess whether those tools produce information that is necessary for assessing compliance with legal and regulatory requirements and policy goals. Our aim is that this framework will stimulate debate and evolution of a framework for promoting responsible, fair, and inclusive use of machine learning underwriting models.

Concluding Thoughts

We believe that our empirical and market research will help to inform a broad dialogue among financial services firms, policymakers, advocates, and others to enable adoption of machine learning in a responsible, fair, and inclusive way. As machine learning uses continue to spread in credit and other financial services contexts, determining how to refine and strengthen market practices and federal regulatory frameworks is a top priority going forward.

We also expect that this project and other FinRegLab research initiatives will shed additional light on a range of issues concerning data governance protocols for the use of non-traditional financial information in credit underwriting. As discussed above, those issues are closely intertwined with machine learning underwriting models, and may be a bigger driver of financial inclusion when the two are used in tandem. But data bias and governance issues also arise in contexts that do not involve machine learning in the first instance, and therefore require continuing direct attention in their own right to further refine best practices and regulatory expectations.

Beyond answering these basic questions about the tools and processes for managing machine learning models and non-traditional data, business and market considerations will also play a critical role in determining if these innovations in fact improve the inclusiveness and fairness of the credit system or of financial services more generally. Individual firms' decisions about strategy and market segmentation, operational barriers to technology and data adoption particularly among smaller firms, and acceptance by investors and secondary market actors of loans originated using non-traditional methods and data will all shape the extent to which model innovations are used to increase access to historically underserved populations or merely to target existing market segments with greater precision.

It is also critical to note that while filling information gaps and adopting more predictive models could help substantial numbers of consumers and small business owners access more affordable credit, such actions will not by themselves erase longstanding disparities in income and assets or recent hardships imposed by the pandemic. These factors are likely to result in many individual applicants being assessed as presenting significant risk of default, which will continue to affect whether they are granted credit and at what price. This underscores the importance of using many initiatives and policy levers to address the deep racial disparities in income and assets at the same time that stakeholders in the credit system continue to explore and implement promising credit and modeling technique innovations. While there is reason to believe that the financial system can enhance its ability to provide fair and inclusive products and services, relying solely on it to address these cumulative, structural issues would produce too little change too slowly.

In addition to focusing on data governance issues and addressing long-term income and asset disparities from other angles, Congress could play a critical role in expanding the resources and data needed to support technical research and policy analysis on AI and machine learning use in financial services. Earmarking some portion of the funding the U.S. government is committing for AI research at the National Science Foundation for financial services research is one possible way to facilitate more research by academics and public-mission research organizations. Access to data is critical for research in these areas as well. The creation of publicly available datasets, with anonymized and perhaps synthetic data, would enable a range of research questions to be probed.

As our work progresses this year, I would be pleased to share insights from our efforts. Our empirical study, market review, and policy analysis should help to inform the extent to which current laws and regulations are able to be satisfied in light of the emergence of more complex underwriting models. It should also help to show how well diagnostic tools are able to monitor and explain complex underwriting models, and it should generate policy insights for the safe, inclusive, and nondiscriminatory adoption of machine learning.

Thank you again for the opportunity to speak with you today.

NFHA NATIONAL FAIR HOUSING ALLIANCE

Congressional Testimony
for the
House Financial Services Task Force on
Artificial Intelligence's Hearing

Equitable Algorithms: How Human-Centered AI Can Address Systemic Racism and Racial Justice in Housing and Financial Services
May 7, 2021 @ 12:00pm ET

Lisa Rice, President and CEO
National Fair Housing Alliance

Introduction

"Equality is not a concept. It's not something we should be striving for. It's a necessity. Equality is like gravity. We need it to stand on this earth as men and women…We need equality. Kinda now."
— Joss Whedon, co-founder of Bellwether Pictures

Algorithmic-based systems impact every area and aspect of our lives either providing access to key services that can open the doors of opportunity or blocking our ability to take advantage of critical amenities that we need to survive and live successful lives. Algorithms can determine whether consumers will have access to housing, get a living wage job, access quality credit, get released on bail after an arrest, or serve a prison sentence. Algorithms even determine whether a sick patient will receive needed healthcare or even whether a homeowner will get a refinance loan.

The math and science behind the development of algorithms are neither good nor bad. However, how these systems are designed, the data used to build them, the subjective renderings applied by the scientists creating the models, and other components of the systems can create or further entrench structural racism and other forms of inequality.

It is imperative that we hastily work to eliminate bias from these systems. Studies reveal that structural inequality, including the harms perpetuated by unfair tech, are not only having a deleterious impact on individuals and communities, but it is stifling the nation's economic progress.

Many innovations have been made in the use of algorithms and Artificial Intelligence (AI) such that they can be used to mitigate against biases innate in legacy tech systems. Much as scientists used the coronavirus, a deadly germ that has killed millions of people in the world, to develop live-saving vaccines, we can use AI to detect, diagnose, and cure harmful technologies that are extremely harmful to people and communities.

Part I – History of Housing/Banking Bias

Algorithmic Systems Can Perpetuate Injustice and Discriminatory Outcomes

Algorithmic systems can create, manifest, amplify, and systemize bias creating harmful impacts for millions of people. This is largely because the data used to build models is deeply flawed, scientists and mathematicians developing the systems are not educated about how technology can render discriminatory outcomes, and regulators are not equipped to sufficiently handle the myriad manifestations of bias generated by the technologies we use in financial services and housing.

While AI, including Machine Learning ("ML") systems, may be relatively new innovations, the building blocks for the models these tools create are tainted with historical bias. Algorithmic-based systems are not developed in a vacuum. They are crafted in a polluted environment that embeds particles of inequality into systems that appear to be facially neutral and innocuous. They carry with them and are imbued with a centuries-long legacy of discriminatory actions and unfair policies that still impact our society.

Throughout the entire history of the U.S., our housing and lending policies were written or implemented in ways that were intentionally discriminatory. In fact, many of our laws – Indian Removal Acts, Slave Codes, Fugitive Slave Acts, Repatriation Acts, Chinese Removal Act, Black Codes, Sundown Ordinances, Japanese Internment Act, Racially Restrictive Covenants, and much more - were explicitly and purposefully designed to provide opportunities to Whites and to simultaneously deny opportunities to people of color.

Even laws that appeared to be racially neutral were implemented with racialized policies. For example, the Home Owners Loan Corporation ("HOLC") Act was passed during the Great Depression for the purpose of saving homeowners from foreclosure, but in implementing the law, the federal government institutionalized a structure for redlining communities of color that was widely adopted by the financial services and real estate industries.[1] The HOLC systemized the association between race and risk, a connection that still exists today.

The National Housing Act of 1934 created the Federal Housing Administration ("FHA"). However, the FHA, building off of the HOLC's racialized method of redlining communities of color, developed race-based underwriting guidelines that not only promoted residential segregation but described people of color as "incompatible racial elements" and "inharmonious

[1] Gregory D. Squires, The Fight for Fair Housing: Causes, Consequences, and Future Implications of the 1968 Federal Fair Housing Act (2018). For an in-depth discussion of the myriad ways the federal government institutionalized redlining and lending discrimination see Chapter 6, entitled *The Fair Housing Act: A Tool for Expanding Access to Quality Credit*.

racial groups."[2] The FHA encouraged the use of racially restrictive covenants and required them in exchange for supporting the bevy of new housing developments built throughout the nation's suburban communities. Even after the Supreme Court declared that racially restrictive covenants were not enforceable, the FHA gave preferential treatment to developers that adopted them.[3] From 1934 to 1962, the federal government backed over $120 billion in mortgages but the FHA's race-based policies meant that less than 2 percent of loans went to people of color.

Many other laws, seemingly racially neutral, were implemented with the use of discriminatory policies including the National Highway Acts, Fair Labor Standards Act, Tax Codes, Housing Act of 1949, Social Security Act, Anti-Drug Abuse Act of 1986, and local zoning ordinances. Moreover, hundreds more laws have been passed with no outright ill-intention, but because the laws were implemented with no consideration for the deep levels of inequality in our society, they produced disparate outcomes. The CARES Act, passed in response to the COVID-19 pandemic, is a prime example. The Paycheck Protection Program when initially rolled out excluded roughly 95% of Black-owned, 91% of Latino-owned businesses, 91% of Native Hawaiian or Pacific Island-owned businesses, and 75% of Asian-owned businesses. Business owners who were already well-connected with mainstream banks and business and financial experts were much more likely to access PPP loans, even if they did not have dire need for assistance.[4]

This bevy of laws, regulations and policies created structural inequities and systemic bias that is still being manifest in our society. Residential and school segregation, the inextricable link between place and opportunity, the dual credit market, the inequitable health ecosystem, the patchwork of exclusive and restrictive zoning systems, and additional structurally unfair systems all stem from a long stream of laws that were either explicitly racist, implemented with racialized policies, or produced disparate impacts on communities of color. The effect of these policies was to steepen the racial wealth, income, and homeownership gaps.

These systems are still performing their originally intended function, perpetuating disparate outcomes and generating tainted, bias-laden data that serves as the building blocks for algorithmic-based utilities like tenant screening selection, credit scoring, insurance rating, risk-based pricing, digital marketing, and automated underwriting systems. The scalability power and reinforcement effect of AI algorithms could make them bad agents that amplify discriminatory outcomes if they are not controlled.

While we have passed civil rights statutes designed to stop discrimination; we have not designed laws to dismantle the systems of inequality that are still producing biased impacts. Laws like the Fair Housing Act of 1968 or the Equal Credit Opportunity Act of 1974 prohibit housing and financial services providers from considering race, national origin or gender when making a housing related decision. But we have done little to nothing to remedy or rectify the

[2] Lisa Rice, Missing Credit: How the U.S. Credit System Restricts Access to Consumers of Color, Testimony before the U.S. House Committee on Financial Services (Feb. 26, 2019).
[3] Richard Rothstein, The Color of Law: A Forgotten History of How Our Government Segregated America (2017).
[4] Brian Thompson, Getting Help for Minority-Owned Businesses Shut Out of PPP Loan Relief, Forbes (May 12, 2020)

discriminatory structures that we created from centuries of discriminatory laws. For example, though the Fair Housing Act does contain a provision for dismantling systemic inequality – the Affirmatively Furthering Fair Housing mandate – it has never been enforced.

Part II - How Algorithms Can Manifest Bias

Data Risks

Introduction: Data and Technology are Not Innocuous

Data is tainted. Computers and technology are not color- or gender-blind. In fact, much of the data used to build algorithmic systems is covered in a patina of bias. We all know the adage that bad inputs equal bad outputs. Well, the same holds true here. Biased data in equals biased outcomes. All the technologies we use in housing, employment, health, credit, law enforcement, advertising, and other sectors contain bias because the systems were created with tainted data.

The Data Can Be Under-Inclusive

Building fair AI systems requires the use of quality, reliable, robust data that truly reflects the patterns and behaviors of the people the models are designed to assess. For example, one challenge is that a disproportionate amount of data used to build models in the housing and financial services space is generated from information housed with the credit repositories. However, credit repository data can be very limiting because not all information about consumer behavior is reported to the credit reporting agencies. Moreover, the data that is reported is reflective of the structural biases replete throughout our society.

In many instances, BIPOC (Blacks, Indigenous, and People of Color) consumers are disproportionately missing from the data. AI systems can only see the patterns that are existent in the data. Because people of color disproportionately access data outside of the financial mainstream, they are underrepresented in datasets used to build financial services systems. Moreover, because BIPOC consumers are disproportionately rejected for credit, their consumer patterns are under-represented in the data. For example, many BIPOC consumers live in credit deserts and disproportionately access financial services from non-traditional, alternative credit providers such as payday lenders, check cashers and title money lenders. These non-traditional credit providers do not report consumers' timely payments to the credit repository system. Thus, consumers who are accessing credit outside of the financial mainstream and who pay their obligations as they should are not reaping the benefit of their good behavior simply because it is not reported. These consumers are essentially invisible to most scoring systems used in the housing and financial services space. This in no way means that these consumers are poor risks or are not responsible. It simply means that the data used to build traditional algorithmic financial services models is not representative of underserved groups.

As a result, AI systems built using unrepresentative data will not be able to score underserved consumers at all since these consumers register as credit invisible, or the systems will inaccurately score underserved consumers likely assessing them as more risky than they really are.

Finally, AI systems are sometimes built with data sets that are over-weighted with certain features or lack critical information that can better inform the algorithm. The data collection itself might be biased. An example of this is when Amazon's recruitment AI system disadvantaged women. The system was built with Amazon's own database of senior executives who were disproportionately White men. The system learned that men were preferable applicants. Rather than solely relying on a candidate's qualifications, the system penalized applicants whose resumes contained the word "women" and downgraded graduates of all-women's universities.[5] Another example is when facial recognition technology mis-reads women or people of certain racial or ethnic groups because the data used to train the system did not include enough examples of women and people of color.[6]

The Data Can Reflect Historical Bias

Discrimination in the marketplace taints the data collected by credit repositories thus data can be extremely harmful. Discrimination in the employment, housing, credit, health and other sectors impacts the type and quality of data reflected in our credit repository system. How that data is ultimately used by credit modelling agencies can exacerbate disparities. Although discrimination is a common occurrence[7], it is not accounted for in the way credit data is collected or utilized. When credit repositories gather data, they do not simultaneously ascertain if a consumer has obtained credit from a predatory, discriminatory or abusive debtor for the purposes of ameliorating any negative fallout. Data is captured as if it is innocuous and benign when the opposite is the case. Data is infused with the discrimination replete throughout our society. When credit repositories collect data, without any assessment of the quality or legitimacy of that data, they help perpetuate the inequities that harm under-served consumers.

Some have attempted to mitigate bias in our markets by moving toward automated systems lulled by the myth that data is blind. Data is not blind, nor is it harmless. It can be dangerous and toxic particularly when it manifests the discrimination inherent in our systems. For example, researchers at Berkeley have found that fintech lenders that rely on algorithms to generate decisions on loan pricing discriminate against borrowers of color because their systems "have not removed discrimination, but may have shifted the mode."[8] It is estimated that borrowers of color are being overcharged by $765 million per year. Similarly, concerns have been raised about AI systems based on appraisal data, which may reflect historical biases due to the HOLC maps and other forms of discrimination. A 2018 Brookings Institution study found that homes in

[5] David Meyer, Amazon Reportedly Killed an AI Recruitment System Because It Couldn't Stop the Tool from Discriminating Against Women, Fortune (Oct. 10, 2018).
[6] James Vincent, Gender and Racial Bias Found in Amazon's Facial-Recognition Technology (Again), The Verge (Jan. 25, 2019).
[7] There are over 4 million instances of housing discrimination each year. *See* National Fair Housing Alliance, Defending Against Unprecedented Attacks on Fair Housing: 2019 Fair Housing Trends Report (2020).
[8] Robert P. Bartlett, Adair Morse Richard H. Stanton, and Nancy E. Wallace, Consumer Lending Discrimination in the FinTech Era, UC Berkeley Public Law Research Paper (Sept. 2019).

NFHA NATIONAL FAIR HOUSING ALLIANCE

majority Black neighborhoods were appraised for 23 percent less than properties in mostly White neighborhoods, even after controlling for home features and neighborhood amenities, which raises questions about the appropriateness of the data. Finally, the data gleaned from credit reporting agencies that go into the credit scoring, risk-based pricing, and automated underwriting models do not exist in isolation. Each piece of information has appended to it other bits of data that is inherently connecting risk to race. In essence, these data systems manifest systemic and institutional racism.

The Data Can Inappropriately Exclude Race/Gender Data Needed for Testing Outcomes

Confusion exists regarding how to collect and use race or other protected class data or proxies. As a result, the data used to develop an AI system may not include the information needed to test outcomes based on race or other protected characteristics. However, while race or other protected class data may not be appropriate to use in the model, it may be critical to later evaluating the impact of the model's outcomes.

Model Risks

The Model Can Be Flawed and Discriminatory

AI systems can be designed in a way that encourages biased outcomes. For example, systems that allow users to exclude certain racial or ethnic groups can cause discrimination against protected groups and even enhance the different ways in which users can discriminate against people. The National Fair Housing Alliance and several of its member organizations filed a legal challenge against Facebook over such an issue.[9] The company used to allow entities placing ads for housing, employment, and credit on Facebook's platform to target audiences based on protected class characteristics like gender, race, and national origin. Resolution of this case involved Facebook making eight meaningful and structural changes to its advertising platform including:

- Establishing a separate advertising portal for creating housing, employment, and credit ("HEC") ads on Facebook, Instagram, and Messenger that will have limited targeting options, to prevent discrimination.
- Creating a page where Facebook users can search for and view all housing ads that have been placed by advertisers for the rental, sale, or finance of housing or for real estate related transactions (such as appraisals and insurance), regardless of whether users have received the housing ads on their News Feeds.
- Requiring advertisers to certify that they are complying with Facebook's policies prohibiting discrimination and all applicable anti-discrimination laws.
- Providing educational materials and features to inform advertisers about Facebook's policies prohibiting discrimination and anti-discrimination laws.
- Meeting regularly with the Plaintiffs and their counsel to report on and discuss the implementation of the terms of the settlements.

[9] See National Fair Housing Alliance, Facebook Settlement (March, 2019).

- Permitting the Plaintiffs to engage in testing of Facebook's ad platform to ensure the reforms established under the settlements are implemented effectively.
- Working with NFHA to develop a training program for Facebook's employees on fair housing and fair lending laws.
- Engaging academics, researchers, civil society experts, and civil rights/liberties and privacy advocates (including plaintiffs) to study the potential for unintended bias in algorithmic modeling used by social media platforms.

AI systems that use a scoring system to determine ad placement can also generate bias. Such might have been the case with a research project conducted using Google's platform. A Harvard researcher found that Google searches for people with Black-identifying names turned up more ads suggestive of arrest records and/or criminal backgrounds than did ad searches using White-identifying names. Researchers recommended that by changing the quality score of ads to discount for unwanted bias, Google might be able to minimize bias on its platform. By measuring real-time unwanted discrimination in the way an ad is delivered, and then adjusting the score at auction, bias can be eliminated or minimized.[10]

The Model Can Result in a Biased Feedback Loop

If not carefully designed, AI systems can unduly amplify discriminatory information. For example, if an ad features an African American man, a digital platform registering the content of the ad might skew the ad's delivery to men. As more men click on the ad, because they were historically more likely to see the ad, the digital platform might mis-perceive that men are more likely to be interested in seeing the ad than women and continue to over-skew the ad's delivery to even more men.

As another example, predictive policing systems have been shown to discriminate against Black residents because of feedback loops that, because of historical discrimination in the criminal justice system, result in the targeting of people of color for heightened policing activity, even when no crime has been committed. The U.S. criminal justice system is notoriously biased, particularly when it comes to the area of substance abuse. The FBI's criminal database shows that Blacks, Asian Americans, Latinos and Whites use and sell illegal substances to the same degree. Yet, Blacks are 3-4 times more likely than Whites to be arrested and almost 6 times more likely to be incarcerated for drug-related charges.[11] AI systems that rely on tainted data from the law enforcement system will reinforce discriminatory patterns.

Biased feedback loops exist in models used in financial services as well. The Berkeley study on bias in fintech offers a prime example. The study shows that risk-based pricing systems are likely overpricing Black and Latino borrowers to the tune of $765M annually. Researchers posit that the systems may be optimized for profit and might be picking up on reduced shopping activity among Black and Latino borrowers. However, reduced levels of mortgage loan shopping among Black and Latino borrowers can be linked to the fact that these borrowers disproportionately live

[10] Latanya Sweeney, Discrimination in Online Ad Delivery, ACM Queue (Apr. 12, 2013).
[11] NAACP, Criminal Justice Fact Sheet.

in credit deserts and have less access to banks. In this way, the structural inequities linked to residential segregation and the dual credit market serve as a biased feedback loop that results in borrowers of color being charged more for credit when they pose no greater level of risk.

There Can be Failures in Adequately Testing Models for Discriminatory Outcomes

If systems are not tested for bias, companies can use algorithms that unknowingly manifest discrimination. In other words, modelers may not see bias in an algorithm if they are not looking for it or have not been sufficiently trained to look for it. This is why testing is so important. For example, algorithms might have the incorrect optimization or unseen correlations that perpetuate or amplify an unintended bias. Data scientists must perfect the design of the algorithm to ensure that systems don't treat people unfairly. CoreLogic has been challenged on its CrimSafe tenant-screening system, which contains arrest information. The system can penalize people who have an arrest record but no convictions. This feature, of course, disproportionately discriminates against Blacks and Latinos[12] and no or insufficient testing of models for discriminatory impacts will result in reduced housing opportunities for underserved groups.

Part III - Recommendations for Mitigating the Risk of Algorithmic Bias

There are significant risks of bias and discrimination in AI systems, but the risks are not insurmountable. Following are recommendations as to how lawmakers, regulators, housing providers, financial institutions, and tech companies can mitigate the risk of algorithmic bias.

Integrate the Review of Racial and Other Bias into Every Phase of the Algorithm's Lifecycle

Given the systemic discrimination that exists in almost every aspect of American life, there is a high risk that the data and models used for AI systems will reflect that systemic bias. Accordingly, it is imperative that equity and non-discrimination be top of mind at every phase of the algorithm's lifecycle. It is not enough to merely consider discrimination risk once the AI system is built or even deployed. Instead, the risk of bias must be considered and mitigated at every phase, from data selection to development to deployment to monitoring. Unfortunately, in many instances, regulators in the United States seem to view fair housing and fair lending risk as separate and apart from other AI risks. For example, the federal financial regulators recently issued a Request for Information regarding AI.[13] The section requesting comment on fair lending is relegated to the end of the questions, separate and apart from other AI concepts, such as explainability. Time and again, we see U.S. regulators considering fair housing and fair lending risk as somehow distinct from other risks, rather than as an integral and important part of all discussions of AI risk.

---

[12] National Fair Housing Alliance, Defending Against Unprecedented Attacks on Fair Housing: 2019 Fair Housing Trends Report (2020).

[13] Board of Governors of the Federal Reserve System, Consumer Financial Protection Bureau, Federal Deposit Insurance Corporation, National Credit Union Administration, and Office of the Comptroller of the Currency, Request for Information and Comment on Financial Institutions' Use of Artificial Intelligence, including Machine Learning, 86 Fed. Reg. 16837 (March 31, 2021).

NFHA NATIONAL FAIR HOUSING ALLIANCE

By contrast, the European Union's newly-released proposed regulation for AI ("EU Proposed Regulation") clearly recognizes that AI systems that impact the evaluation of creditworthiness pose a high risk to fundamental rights, including the right to non-discrimination.[14] The proposed regulation creates a risk-based framework of three categories: (i) unacceptable risk, where the practices are prohibited (e.g., social scoring by public authorities); (ii) high-risk AI systems, which would need to comply with new requirements; and (iii) non-high-risk AI systems, which are encouraged to adopt voluntary codes of conduct. The appendix to the proposed regulations lists several high-risk AI systems, most notably, AI system that relate to the access to and enjoyment of essential private services and public services and benefits, including:

- **AI systems intended to be used to evaluate the creditworthiness of natural persons or establish their credit score.**

Importantly, the EU made this determination based on explicit recognition of (i) the importance of this benefit to fully participate in society or improve one's standard of living and (ii) the high risk of discrimination. The preamble to the proposed regulation states:

> Another area in which the use of AI systems deserves special consideration is the access to and enjoyment of certain essential private and public services and benefits **necessary for people to fully participate in society or to improve one's standard of living**. In particular, AI systems used to evaluate the credit score or creditworthiness of natural persons should be classified as high-risk AI systems, since they determine those persons' access to financial resources or essential services such as housing, electricity, and telecommunication services. AI systems used for this purpose may lead to discrimination of persons or groups and **perpetuate historical patterns of discrimination,** for example based on racial or ethnic origins, disabilities, age, sexual orientation, or **create new forms of discriminatory impacts**.

Thus, the EU recognizes that not all AI is the same and that AI systems that evaluate creditworthiness should be held to a higher standard given the far-reaching impact on consumers' life options and the high risk of discrimination. The proposed regulation reflects this key premise by incorporating a review for discrimination risk in all aspects of the proposed requirements, from data governance to post-market monitoring.

[14] European Commission, Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence (aka, the "Artificial Intelligence Act") (April 21, 2021). It may also be instructive to review recent actions by the federal Food and Drug Administration ("FDA") and the state of Virginia, both of whom have considered the use of AI with respect to high-risk scenarios. *See* FDA, AI/ML Action Plan for AI/ML-based Software as a Medical Device (Jan. 12, 2021); Virginia Consumer Data Protection Act, Title 59.1, Ch. 52 (2021) (requiring data protection assessments for the processing of any personal data that is to be used for the purpose of profiling where there is a reasonable risk of unlawful disparate impact on consumers).

NATIONAL FAIR HOUSING ALLIANCE

Use Reliable Methods for Mitigating the Risk of Data Bias

The Reject Inference Pool Can Be Used for Equitable Credit Access

Diverse peer-reviewed research works have shown that basing credit scoring solutions solely on the behaviours of approved customers or performances of approved loans can be detrimental to future loan applicants, especially the historically under-approved BIPOC applicants. As AI algorithms may only learn from patterns present in a dataset, declined applicants may never be scored fairly by an AI-credit scoring solution because their patterns are either missing or almost invisible in the data being used to train such scoring solutions. AI solutions are not magical; they can only see or detect what is already existent in the data. Thus, a credit scoring solution that has been historically trained on data exclusive of applicants that are thin-file, (undocumented) immigrants, or renters may continue to classify such borrowers as high risks since its training data lacks sufficient signals from these categories of applicants.

The Reject Inference (RI) is an inclusive method that augments data of approved loan applicants with data of declined applicants so that an AI algorithm trained on such inclusive data would be unbiased or less discriminatory towards under-approved applicants. RI is a collection statistical technique that tries to simulate what the reality could look like if declined loan applications were approved.

While it may be difficult to rigorously justify the fitness of counterfactual RI techniques such as fuzzy augmentation, simple augmentation, or any of their variants for credit scoring solutions, an (experimental) pool may be created for a fraction of the declined applicants so that the credit risks in this pool are shared (with some formula) by all lenders. Such a pool would provide real quality data that could be used to evaluate the accuracy of the original reject decision; augment training data on approved applicants without a need for theoretical, uncertain RI techniques; and, more importantly, present inclusive signals from underserved borrowers to AI algorithms.

Representative and Robust Datasets Should be Developed

One way to address challenges with insufficient data is to augment more exclusive datasets with information from non-traditional sources as a means of building a more representative and robust dataset. Community Development Financial Institutions and state Housing Finance Agencies may be two sources of obtaining data that are more reflective of the practices of BIPOC and other underserved consumers.

Another means of building more robust dataset is to capture rental housing payment data. The Urban Institute conducted important research[15] regarding the efficacy of using rental housing payment information in financial services automated underwriting systems. Traditional credit scoring systems do not incorporate the use of rental housing payment information and this can be

[15] Laurie Goodman and Jun Zhu, Rental Pay History Should Be Used to Assess the Creditworthiness of Mortgage Borrowers, Urban Institute (Apr. 16, 2018).

harmful for consumers who access credit outside of the financial mainstream. But rental housing payment information may be able to significantly improve the ability of models to expand access to credit. The Urban Institute's research found that borrowers who did not miss a housing payment for two years made on-time mortgage housing payments for the next three years. The analysis reveals that rental housing payment data would be a very strong predictor of mortgage risk.

Protected Class Data Should be Collected and Used to Appropriately Build and Test Fairer Tech

Although protected class data should not be used to create disparate treatment or disparate impacts, such data can be used responsibly to build and test AI systems. Here, the EU Proposed Regulation's approach to data governance may be instructive. The preamble to the proposed regulation clearly states the importance of robust data governance with respect to fair AI systems: "High data quality is essential for the performance of many AI systems, especially when techniques involving the training of models are used, with a view to ensure that the high-risk AI system performs as intended and safely and it does not become the source of **discrimination** prohibited by [European] Union law."[16] More specifically, the proposed regulation would require the review of data sets in view of possible bias.[17] In addition, the proposed regulation would allow the providers of high-risk AI systems to process special categories of personal data based on protected characteristics in order to protect the right of others from the discrimination that might result from the bias in AI systems.[18] Similarly, here in the U.S., protected class data should be used responsibly to build equitable AI systems and test for potentially discriminatory outcomes.

A Publicly-available Dataset Should be Released for Research Purposes

Congress should encourage and support public research that analyzes the impact of AI in housing and financial services for consumers of color and other protected classes. In particular, Congress should encourage the Consumer Financial Protection Bureau, the Federal Housing Finance Agency, Fannie Mae, Freddie Mac, and the Federal Housing Administration to release more loan-level data from the national mortgage survey and the national mortgage databases so researchers, advocacy groups, and the public can study bias in the housing and finance markets, including as that bias may relate to the use of AI.

Ensure Models Undergo Robust Testing for Potential Discriminatory Outcomes

We must develop methods to analyze and test our systems to understand better how multi-variate interactions in AI models might be manifesting bias and affecting consumers' ability to fairly access products and services. For example, we can use AI to test the data we use in our systems

[16] EU Proposed Regulation at Recital 44.
[17] *Id.* at Title III, Ch. 2, Art. 10.
[18] *Id.*

NFHA NATIONAL FAIR HOUSING ALLIANCE

to determine if there are any discriminatory associations and then mitigate against them. We can also set the bar high for model validation with an eye toward diminishing bias in the systems.

Here again, the EU Proposed Regulation may be helpful. The EU's Proposed Regulation provides a robust regulatory framework for high-risk AI systems, which includes those systems that evaluate creditworthiness. In addition to the data governance requirements noted above, the proposed regulation would require providers to implement controls related to the following:

- Transparency,
- Human oversight,
- Risk and quality management systems,
- Security, and
- Post-market monitoring.[19]

Moreover, a provider of a high-risk AI system would need to conduct a conformity assessment and certify the system's conformity with the regulation *before* the system is released to the market to avoid consumer harm and the proliferation of discriminatory systems.[20] Penalties by regulators for non-compliance would be as high as 6% of the entity's total global earnings (before costs).[21] Although the EU's Proposed Regulation has been subject to criticism by some advocates for the over-reliance on provider self assessments and the lack of a private right of action,[22] it does provide a useful example of a robust regulatory framework. In particular, it is notable that the proposed regulation shows a clear commitment to fundamental rights, including the right to non-discrimination, that is integrated throughout the proposal.

Ensure Relevant Staff Receive Appropriate Fair Housing/Fair Lending Training and Reflect the Diversity of America

Educate AI Stakeholders about Racial Inequality and Structural Racism

All AI stakeholders – including regulators, housing providers, financial institutions, and tech companies - should be committed to ensuring that all of their staff receive fair housing and racial equity training. Trained professionals are better able to identify and recognize issues that may raise red flags; they are also better able to design solutions for debiasing tech and building fairer systems. In fact, recent innovations in developing mechanisms for debiasing tech has come from data scientists and engineers who were trained on issues of fairness. For example, employees at Google developed What-If[23], a diagnostic tool for detecting various types of bias and ML-fairness-gym[24] a simulation tool to test the impacts of machine learning systems in different

[19] EU Proposed Regulation at Titles III and VIII.
[20] *Id.* at Title III, Ch. 3 and 5.
[21] *Id.* at Title X, Art. 71.
[22] *See, e.g.,* Adam Satariano, Europe Proposed Strict Rules for Artificial Intelligence, N.Y. Times (April 21, 2021).

[23] Google, https://pair-code.github.io/what-if-tool/
[24] Google, https://github.com/google/ml-fairness-gym

social environments. Employees at Microsoft developed Fairlearn[25], a tool for diagnosing and debiasing machine learning systems. The more the field is educated about fairness and equity issues, the better tools will be created to expand opportunities for consumers.

Increase Diversity in the Tech Field

Increasing diversity will lead to better outcomes for consumers. Research shows that diverse teams are more innovative and productive.[26] Companies with more diversity are more profitable.[27] Diverse teams can help bring broader ideas and solutions to the workplace and enhance morale. Moreover, in several instances, it has been the people of color who were able to able to identify potentially discriminatory AI systems.[28]

Ask the General Accounting Office ("GAO") to Review Federal Oversight of AI Bias

Given the rapid proliferation of AI systems in the critically-important areas of housing and financial services, Congress should ask the GAO to immediately review federal supervision and enforcement of fair lending laws, particularly with respect to oversight of AI systems used by housing providers and financial institutions. The GAO last conducted this type of review 25 years ago (in 1996), which resulted in significant policy changes and renewed efforts for robust fair lending supervision and enforcement.[29] The time is right to conduct a new review of the federal banking regulators' fair lending approaches and methodologies.

Conclusion

We can all agree that discriminatory policies like the federal HOLC's discriminatory redlining system and the FHA's biased practices created a housing finance structure that had a long-lasting and detrimental effect on American society, limiting the life choices of millions of people of color for generations up through the present time. Right now, America is at a similar crossroads in determining whether to develop equitable AI systems that serve and uplift the whole of the national financial services market, or one that perpetuates and amplifies old discriminatory patterns. The time to act is now as the use of AI in financial services proliferates in every aspect of housing and consumer credit and has the potential for far-reaching adverse impacts for people of color that could overshadow even the devastation caused by the HOLC, FHA, and other entities that perpetuated discriminatory practices. Government, industry, and advocacy groups should work together to envision and create AI systems that support equitable, non-discriminatory housing and finance markets. Doing so will not just benefit individual consumers, it will advantage our whole society. Citigroup issued an analysis revealing that if racial

[25] Microsoft, https://fairlearn.org/
[26] John Rampton, Why You Need Diversity on Your Team, and 8 Ways to Build It, Entrepreneur (Sept. 26, 2019).
[27] David Rock and Heidi Grant, Why Diverse Teams Are Smarter, Harvard Business Review (Nov. 4, 2016).
[28] Steve Lohr, Facial Recognition is Accurate, if You're a White Guy, N.Y. Times (Feb. 9, 2018) (explaining how Joy Buolamwini, a Black computer scientist, discovered that facial recognition worked well for her White friends but not for her).
[29] GAO, Fair Lending: Federal Oversight and Enforcement Improved but Some Challenges Remain, GGD-96-145, Aug. 13, 1996).

NFHA NATIONAL FAIR HOUSING ALLIANCE

inequality was eliminated, the U.S. GDP would increase by $5 trillion over a 5-year period.[30] Advancing equitable algorithmic systems would lead to increased productivity and improve people's quality of life.

In some respects, the U.S. is behind the ball in advancing fair tech. If we want to retain our competitive edge in the global society, we should hasten to remove bias from existing technologies and take the necessary steps to ensure all systems going forward are fair and equitable.

[30] Dana Peterson and Catherine Mann, Closing the Racial Inequality Gaps, Citigroup (September, 2020).

Remarks To The Task Force on Artificial Intelligence of the House Financial Services Committee

Kareem Saleh
Founder and CEO, FairPlay
May 6, 2021

Thank you Chairman Foster, Ranking Member Gonzales and other members of the task force for the opportunity to testify today.

My name is Kareem Saleh and I'm the founder and CEO of FairPlay, the world's first fairness-as-a-service company. I have witnessed firsthand the extraordinary potential of AI algorithms to increase access to credit and opportunity. But I have also seen the risks these algorithms pose to many Americans.

If we are to fully harness the benefits of AI, we must commit to building infrastructure that embeds fairness into every step of the algorithmic decisioning process. AI is speeding like a train to power the decisions of companies and governments and we should be laying down fairness tracks to guide its route.

Despite the passage of the fair lending laws almost 50 years ago, people of color, women and other historically underprivileged groups are still denied loans at an alarming rate. The result is a persistent racial wealth gap and fewer opportunities for minority families and communities to create a stable and prosperous future.

Why are we still so deeply unfair?

The truth is, the current methods of bias detection in lending are unsuited to the AI era. Even though lending has become AI-powered and automated, fair lending compliance is stuck in the analog, paper-based past.

So how can we bring fair lending compliance into the 21st century?

Here are three ways:

First, we must do better at debiasing data and identifying variables that interact in ways that proxy for protected status.

Currently, fair lending compliance starts with a human review of variables for discrimination. While this is an important step, no human can discern the complex interactions between seemingly fair variables, where bias often hides. Today there are increasingly good methods for locating and counteracting sources of data bias before they result in discriminatory decisions, and these methods ought to be more widely used.

Second, we could require that AI models be validated in ways that enhance their fairness rather than legitimate their unfairness.

A key step in the fair lending compliance process is reviewing an algorithm for unfair outcomes. If, say, Black applicants are approved at materially lower rates than whites, lenders are required to investigate whether this disparity could be justified by a business necessity or whether the lender's business objectives could be met through a less discriminatory means.

It's at this stage -- the search for less discriminatory alternatives and the invocation of business justifications -- where our current fair lending system has the greatest potential to evolve.

The way most lenders search for less discriminatory models involves taking credit scores out of an algorithm, re-running it, and evaluating the differences in outcomes for protected groups. This method almost always results in a fairer model, but also a less profitable one.

This puts lenders in a Catch 22: they'd like to be fair but they'd also like to stay in business. Thus, most lenders end up trying to justify use of the biased model as a business necessity because they could not find a less discriminatory algorithm with the same predictive power.

Today a better, fairer option exists: use AI fairness tools to debias algorithms without sacrificing profitability. Several AI techniques allow lenders to take a variable like credit score and disentangle its predictive power from its disparity-driving effects. The predictive power remains, while fairness increases.

In many instances, these AI fairness tools have improved outcomes for protected groups anywhere from 10-30% without increasing risk.

Third, we must give the people charged with fair lending compliance the tools and training they need to succeed in the AI era.

At FairPlay, we make AI fairness software that allows every fair lending compliance officer to answer five key questions about their algorithm:

1. Is it biased?
2. If so, why?
3. Could it have been fairer?
4. Does being fairer have a cost?
5. Did we give rejected customers a pathway to being approved in the future?

Finally, there are policy measures Congress and regulators could take to enhance the fairness of automated decisioning systems, including mandating that fairness testing:

- Be done by more lenders,
- More often,
- To their marketing, underwriting, pricing and collections models
- And include a rigorous search for less discriminatory alternatives.

In addition, policymakers should ease the fear of liability for lenders who commit to thoroughly searching for disparities, to reward rather than punish those who proactively look for fairer decisioning systems.

To bring fairness to AI decisions we must build the fairness infrastructure of the future, not justify the discrimination of the past. Industry, urged on by policymakers and regulators, has the opportunity to update fair lending compliance for the AI era. Using AI debiasing tools, we can embed fairness into algorithmic decisions in a way that promotes opportunity for all Americans while allowing financial institutions to reap the rewards of a safe and inclusive approach.

AI Fairness will not happen on its own. It requires attention and action. If we prioritize fairness, the machines we build will follow.

Thank you for allowing me to address this body. I’m happy to answer any questions.

July 14, 2021

Representative Sylvia Garcia
1620 Longworth HOB
Washington, DC 20515

Re: Follow Up Questions Regarding FSC Artificial Intelligence Task Force Hearing– Equitable Algorithms: How Human-Centered AI can Address Systemic Racism and Racial Justice in Housing and Financial Services

Dear Representative Garcia;

Thank you for your interest in my testimony before the Artificial Intelligence Task Force at its hearing held on May 7, 2021. Below are responses to the questions you posed to me for the hearing record.

The FDIC found that 60 percent of the decline in unbanked households over the last 5 years was because of a change in income. Ms. Rice, how would you characterize the growing opportunity gap related to overall economic well-being? In your opinion, would technology alone increase economic well-being? Or is this change a larger, more systemic need?

Income, Technology, and Opportunity Gap

America's long history of discriminatory housing and finance policies[1] has created distinct advantages for White families and disadvantages for families of Color, leading to massive wealth, homeownership, and credit gaps that persist today. White wealth has soared while Black wealth has not kept pace. In 2016, the typical middle-class Black household had $13,024 in wealth versus $149,703 for the median White household.[2]

[1] *See* Lisa Rice, "The Fair Housing Act: A Tool for Expanding Access to Quality Credit," The Fight for Fair Housing: Causes, Consequences, and Future Implications of the 1968 Federal Fair Housing Act (Gregory Squires, 1st ed. 2017) (providing a detailed explanation of how federal race-based housing and credit policies promoted inequality). *See also*, K. Steven Brown et al., Confronting Structural Racism in Research and Policy Analysis, The Urban Institute (Feb. 2019); Richard Rothstein, The Color of Law: A Forgotten History of How Our Government Segregated America (2017).

[2] Heather Long and Andrew Van Dam, The Black-White Economic Divide Is as Wide as It Was in 1968, Washington Post (June 4, 2020).

NFHA NATIONAL FAIR HOUSING ALLIANCE

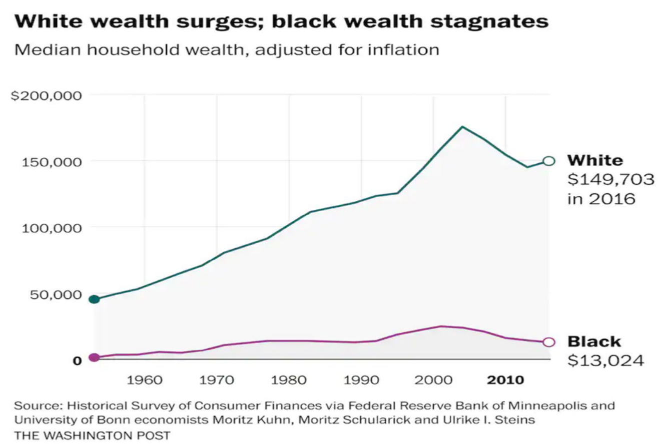

In 2019, White family wealth sat at $188,200 (median) and $983,400 (mean).[3] In contrast, Black families' median and mean net worth were $24,100 and $142,500, respectively. These wealth disparities, in turn, reflected intergenerational transfer disparities: 29.9 percent of White families have received an inheritance, compared with only 10.1 percent of Black families.[4]

There are racial barriers to economic opportunities and these barriers are more prominent for people of Color, Blacks, Native Americans, and Latinx in particular.[5] These barriers are further magnified by the U.S.'s discriminatory and broken credit scoring system[6] which is based on noisy data that underrepresents people of Color.[7]

---

[3] Neil Bhutta, Jesse Bricker, Andrew Chang, et al., Changes in U.S. Family Finances from 2016 to 2019: Evidence from the Survey of Consumer Finances, 106(5) Fed. Reserve Bulletin (Sept. 2020).
[4] Neil Bhutta, et al., Disparities in Wealth by Race and Ethnicity in the 2019 Survey of Consumer Finances, FEDS Notes, Board of Governors of the Federal Reserve System (Sept. 2020).
[5] Will Douglas Heaven, "Bias isn't the only problem with credit scores—and no, AI can't help", June 17, 2021. URL: https://www.technologyreview.com/2021/06/17/1026519/racial-bias-noisy-data-credit-scores-mortgage-loans-fairness-machine-learning/?truid=&utm_source=weekend_reads&utm_medium=email&utm_campaign=weekend_reads.unpaid.engagement&utm_content=06.26.21.subs&mc_cid=aad1663503&mc_eid=eead1c58a0

[6] Megan Leonhardt, "Democrats and Republicans in Congress agree: The system that determines credit scores is 'broken'". Accessed 07/08/2021. URL: https://www.cnbc.com/2019/02/27/american-consumer-credit-rating-system-is-broken.html
[7] Laura Blattner and Scott Nelson, "How costly is noise? Data and disparities in consumer credit", May 5, 2021. URL: https://arxiv.org/pdf/2105.07554.pdf

Credit scores are almost a necessity to access basic economic opportunities such as employment, mortgages, auto loans, and credit cards. A good deal of research suggests that credit scores highly correlate with income, and as people of color are overrepresented in the low- and moderate-income categories it is not unreasonable to observe a dense cluster of them on the lower tail of credit score distribution thereby revealing economic inequalities that are racially ubiquitous.[8] In addition, a high correlation between income and credit scores may imply that historical income inequalities would widen disparities in credit access, which in turn may worsen inequalities in consumption, lending, housing access, and healthcare.[9]

The "60 percent of the decline in unbanked households over the last 5 years" observed by the FDIC, driven by changes in income, will undoubtedly exacerbate racially disparate effects on credit scores.

In addition, people of color face a high rejection rate when they apply for loans and as such using observations on approved loans to develop credit scores exacerbate the under-representativeness issues that Blacks and Latinx consumers face.[10] These marginalized patterns, whereby people of Color are under-represented in data sets used to build algorithmic models, translates to limited consumption, housing, employment, welfare, and healthcare opportunities thereby widening economic disparities.[11]

These data issues cannot be fixed by technology, i.e. AI and machine learning.[12] Conscientious policy actions are required to ensure that credit bureaus have permissioned access to consumer data that could be predictive of credit risk or economic progress.[13] While there are tested statistical techniques to account for

---

[8] Andre Dua et al., "Unequal America: Ten insights on the state of economic opportunity", McKinsey & Company May 26, 2021. URL: https://www.mckinsey.com/about-us/covid-response-center/inclusive-economy/unequal-america-ten-insights-on-the-state-of-economic-opportunity

[9] Rachael Beer, Felicia Ionescu, and Geng Li, "Are Income and Credit Scores Highly Correlated?", Aug 13, 2018. URL: https://www.federalreserve.gov/econres/notes/feds-notes/are-income-and-credit-scores-highly-correlated-20180813.htm

[10] Rashida Richardson, "Racial Segregation and the Data-Driven Society: How Our Failure to Reckon with Root Causes Perpetuates Separate and Unequal Realities", Berkeley Technology Law Journal, Vol. 36, No. 3, 2022. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3850317

[11] Andre Dua et al., "Unequal America: Ten insights on the state of economic opportunity", McKinsey & Company May 26, 2021. URL: https://www.mckinsey.com/about-us/covid-response-center/inclusive-economy/unequal-america-ten-insights-on-the-state-of-economic-opportunity

[12] For a more in-depth discussion of the risks to consumers of color of artificial intelligence and machine learning, *see* National Fair Housing Alliance and Advocate Letter to the Federal Financial Regulators regarding the Request for Information and Comment on Financial Institutions' Use of Artificial Intelligence, including Machine Learning (July 1, 2021). *See also*, National Fair Housing Alliance Press Release, Leading Civil Rights, Consumer, and Technology Advocates Urge the Federal Financial Regulators to Promote Equitable Artificial Intelligence in Financial Services (July 1, 2021).

[13] Rashida Richardson, "Racial Segregation and the Data-Driven Society: How Our Failure to Reckon with Root Causes Perpetuates Separate and Unequal Realities", Berkeley Technology Law Journal, Vol. 36, No. 3, 2022. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3850317

rejected applicants in credit score development, policymakers and regulators should ask credit score developers to disclose if or how they account for rejected applicants in their modeling. Moreover, guidance on how to incorporate information on rejected applicants in the model development process should be provided.

Technology alone cannot fix economic inequalities. The inequities in our society were brought about by centuries of race-based policies and programs that have created unfair systems, like segregation, the dual credit market, and exclusionary zoning policies, that still exist today. It will take a comprehensive suite of policies, programs, and resources to counter systemic inequality. But access to more inclusive data would minimize the inequality gap especially when alternative data such as bank transaction activities are used as complementary data that augment conventional credit data. More research should also be conducted to see what difference this additional data will make in practice. Regulators will need to ensure that access to data does not compromise consumer privacy or expose them to predatory lenders.

The ever-widening gap in overall economic well-being for people of Color cannot be solved through technology alone, but through much greater systemic change, beginning with urgently and systematically addressing the homeownership gap. Homeownership has long been the key to wealth-building for American families, but significant barriers remain for people of Color. We urge you and your colleagues to keep these impediments in mind as you consider the priorities for the Financial Services Committee, and in particular, the Artificial Intelligence Task Force.

According to analysis of Home Mortgage Disclosure Act (HMDA data), there are at least three key barriers to homeownership for borrowers of color:

- Collateral issues, often caused by homes in communities of color being undervalued or subject to disinvestment and disrepair;
- Debt-to-income ratio issues, often caused by down payment challenges as well as unfair and expensive Loan Level Pricing Adjustments (LLPAs) imposed by Fannie Mae and Freddie Mac (collectively, the Government-Sponsored Enterprises, or GSEs); and
- Credit score issues, often reflecting a history of discrimination and a dual housing and finance market (as described above).

For these reasons, we urge you and your colleagues on the House Financial Services Committee (the Committee) to support the following:

- Revise the Appraisal System: The Committee should work with key stakeholders to support efforts by the U.S. Department of Housing and Urban Development (HUD), Federal Housing Finance Agency, and others to revise the appraisal system to ensure a more equitable process that will fairly value homes in communities of Color when compared to similar homes in White communities.

- Support the Neighborhood Homes Investment Act: The Committee should support the Neighborhood Homes Investment Act, which would allow federal income tax credits for home rehabilitation.
- Support the CRL/NFHA Down Payment Assistance Framework:[14] The Committee should support the Center for Responsible Lending and National Fair Housing Alliance (CRL/NFHA) First Generation Down Payment Assistance Program, which would provide assistance for first generation homebuyers with incomes at or below 120% of the Area Median Income. It is estimated that this program could help as many as 12.2 million families, 72% of whom would be families of color. In addition, the Committee should support the recommendation that the U.S. Department of Justice (DOJ) and HUD conduct a study to determine how to implement race-conscious remedies for socially and economically-disadvantaged individuals.
- Eliminate the GSE LLPAs: The LLPAs were created during the foreclosure crisis to increase income for the GSEs, but they often drive up the cost of mortgages for borrowers of color and undermine the GSEs' mandate to provide liquidity for the whole of the national housing finance market. The Committee should work with the Federal Housing Finance Agency (FHFA) to review and ultimately eliminate discriminatory LLPAs.
- Support Guidance for Special Purpose Credit Programs[15]: Special Purpose Credit Programs have long been available under Regulation B (which implements the Equal Credit Opportunity Act), but lenders have been reluctant to use these programs without practical guidance from HUD and the prudential regulators. The Committee should work with the federal financial regulators, HUD, and FHFA to develop practical guidance as well as programs that ensure that the GSEs will buy these mortgages and provide additional liquidity for this market.

In addition to these recommendations, the National Fair Housing Alliance along with the National Housing Council, NAACP, Mortgage Bankers Association, National Association of Real Estate Brokers, National Association of Realtors, and the National Urban League serve on the Steering Committee for the "3 by 30" initiative, which seeks to

[14] Center for Responsible Lending and National Fair Housing Alliance, First Generation: Criteria for a Targeted Down Payment Assistance Program (May 21, 2021)
[15] Lisa Rice, "Using Special Purpose Credit Programs to Expand Equality", National Fair Housing Alliance (November 4, 2020)

NATIONAL FAIR HOUSING ALLIANCE

create 3 million net new Black homeowners by 2030. The plan identifies the following seven areas that will make it possible to address the homeownership gap on a systemic basis:

- Homeownership counseling,
- Down payment assistance,
- Housing production,
- Credit and lending,
- Civil and consumer rights,
- Homeownership sustainability, and
- Marketing and outreach.

We urge you and your colleagues to keep these areas in mind as you consider ways to enact systemic change to close the homeownership and wealth gap for communities of color.

Lisa Rice

Lisa Rice
President and Chief Executive Officer