

**ARTIFICIAL INTELLIGENCE AND HOUSING:
EXPLORING PROMISE AND PERIL**

HEARING

BEFORE THE

SUBCOMMITTEE ON
HOUSING, TRANSPORTATION, AND COMMUNITY
DEVELOPMENT

OF THE

COMMITTEE ON
BANKING, HOUSING, AND URBAN AFFAIRS
UNITED STATES SENATE

ONE HUNDRED EIGHTEENTH CONGRESS

SECOND SESSION

ON

EXAMINING ARTIFICIAL INTELLIGENCE AND HOUSING

—————
JANUARY 31, 2024
—————

Printed for the use of the Committee on Banking, Housing, and Urban Affairs



Available at: <https://www.govinfo.gov/>

—————
U.S. GOVERNMENT PUBLISHING OFFICE

COMMITTEE ON BANKING, HOUSING, AND URBAN AFFAIRS

SHERROD BROWN, Ohio, *Chair*

JACK REED, Rhode Island	TIM SCOTT, South Carolina
ROBERT MENENDEZ, New Jersey	MIKE CRAPO, Idaho
JON TESTER, Montana	MIKE ROUNDS, South Dakota
MARK R. WARNER, Virginia	THOM TILLIS, North Carolina
ELIZABETH WARREN, Massachusetts	JOHN KENNEDY, Louisiana
CHRIS VAN HOLLEN, Maryland	BILL HAGERTY, Tennessee
CATHERINE CORTEZ MASTO, Nevada	CYNTHIA M. LUMMIS, Wyoming
TINA SMITH, Minnesota	J.D. VANCE, Ohio
RAPHAEL G. WARNOCK, Georgia	KATIE BOYD BRITT, Alabama
JOHN FETTERMAN, Pennsylvania	KEVIN CRAMER, North Dakota
LAPHONZA R. BUTLER, California	STEVE DAINES, Montana

LAURA SWANSON, *Staff Director*

LILA NIEVES-LEE, *Republican Staff Director*

ELISHA TUKU, *Chief Counsel*

AMBER BECK, *Republican Chief Counsel*

CAMERON RICKER, *Chief Clerk*

SHELVIN SIMMONS, *IT Director*

PAT LALLY, *Assistant Clerk*

SUBCOMMITTEE ON HOUSING, TRANSPORTATION, AND COMMUNITY DEVELOPMENT

TINA SMITH, Minnesota, *Chair*

CYNTHIA M. LUMMIS, Wyoming, *Ranking Member*

JACK REED, Rhode Island	MIKE CRAPO, Idaho
ROBERT MENENDEZ, New Jersey	MIKE ROUNDS, South Dakota
JON TESTER, Montana	JOHN KENNEDY, Louisiana
CATHERINE CORTEZ MASTO, Nevada	BILL HAGERTY, Tennessee
KYRSTEN SINEMA, Arizona	J.D. VANCE, Ohio
RAPHAEL G. WARNOCK, Georgia	KATIE BOYD BRITT, Alabama
JOHN FETTERMAN, Pennsylvania	

TIM EVERETT, *Subcommittee Staff Director*

KELSEY PRISTACH, *Republican Subcommittee Staff Director*

C O N T E N T S

WEDNESDAY, JANUARY 31, 2024

	Page
Opening statement of Chair Smith	1
Prepared statement	23
Opening statements, comments, or prepared statements of:	
Senator Lummis	3

WITNESSES

Lisa Rice, President and Chief Executive Officer, National Fair Housing Alliance	5
Prepared statement	25
Responses to written questions of:	
Chair Smith	55
Senator Warnock	55
Vanessa Perry, Interim Dean and Professor, The George Washington University School of Business; Nonresident Fellow, Housing Finance Policy Center, Urban Institute	6
Prepared statement	46
Responses to written questions of:	
Chair Smith	55
Senator Warnock	60
Nicholas Schmidt, Partner and Artificial Intelligence Practice Leader, BLDS; Founder and CTO, SolasAI	7
Prepared statement	47
Responses to written questions of:	
Chair Smith	61
Senator Warnock	63

ADDITIONAL MATERIAL SUPPLIED FOR THE RECORD

Letter submitted by ACU	68
Concept Paper submitted by AI	71
Letter submitted by the Center for AI and Digital Policy	88
Letter submitted by Family Housing Fund	94
Letter submitted by Urban Institute	97
Letter submitted by NMHC/NAA	160
“To Err Is Automated: Have Technological Advances in the Mortgage Market Increased Opportunities for Black Homeownership?” Working Paper	163
Letter submitted by Zillow	185

ARTIFICIAL INTELLIGENCE AND HOUSING: EXPLORING PROMISE AND PERIL

WEDNESDAY, JANUARY 31, 2024

U.S. SENATE,
COMMITTEE ON BANKING, HOUSING, AND URBAN AFFAIRS,
SUBCOMMITTEE ON HOUSING, TRANSPORTATION, AND
COMMUNITY DEVELOPMENT,
Washington, DC.

The Subcommittee met at 10:03 a.m., in room SD-538, Dirksen Senate Office Building, Hon. Tina Smith, Chair of the Subcommittee, presiding.

OPENING STATEMENT OF CHAIR TINA SMITH

Chair SMITH. Good morning. The Subcommittee on Housing, Transportation, and Community Development will come to order. Today's hearing will focus on both the promise and the threats that artificial intelligence poses in the housing sector, and I'm very much looking forward to our witness' testimony and to this conversation.

I want to thank Ranking Member Lummis and her staff for your ongoing bipartisan work as we put together this hearing. We both share, I believe, a deep interest in how we can develop Federal policy that supports innovation, and expands opportunity for everyone to have a safe, decent, affordable place to live. And one of the most consequential innovations in recent years is artificial intelligence.

Leader Schumer, Senator Rounds, Senator Young, and Senator Heinrich are leading a bipartisan effort to explore the impacts, opportunities, and threats that AI poses. And they have asked Senate committees to engage in our areas of expertise, which leads us to this Committee hearing today, examining what AI means for housing.

So without a safe, decent, affordable place to live, nothing in your life works. Not your job, your family, your education, your health. So, a foundation question is how AI can help or hinder this goal. We know that some aspects of artificial intelligence have been around for a long time, and we also know that major advances are fueling the use of AI in finance and housing in ways that we need to understand.

Consumers find AI when they encounter chatbots, when they shop online, or digital helpers that seem to be ubiquitous. And AI plays a role when a prospective tenant is looking to rent an apartment, or a renter submits a maintenance request to a management company, or a family tries to qualify for a home loan, or when a person experiencing homelessness is connected to services. These

are powerful tools that hold great potential to cut costs and to target services, to reduce wait times, and to even reduce bias.

But they also have the potential to bake in existing inequities and to reduce accountability and even limit opportunity. Today, AI is being used actively in every part of the housing continuum, from emergency homelessness services to mortgage financing. And as I was preparing for this hearing, I found endless applications.

AI is being deployed, for example, to help connect people experiencing homelessness with health and housing resources. AI is helping to forecast more precisely and accurately where families are at risk of eviction so that we can better target assistance. Academics and advocates are using AI and machine learning to help understand, and map the country's zoning laws and codes, spanning 30,000 different localities.

And these insights could help us to understand the dense and complicated rules that govern where, and how, and what types of housing are being built so that we can make better decisions about how to boost housing supply and lower costs. So there are many opportunities, and there are also some very real concerns about the threats that AI poses.

In Minnesota, some landlords are reportedly using AI-generated tenant-screening reports that include incorrect and sometimes illegal or off-limits information. The result, it's even harder for people to find a place to rent in some circumstances, and they may never know why they were declined or be able to correct the record. For landlords, maybe it's just easier to move on to the next applicant rather than consider additional information, another example of how AI used in a bad way can be quite harmful.

There's a lawsuit in Minnesota right now against a law firm that has allegedly automated the process of filing evictions for landlords. In 1 month, the firm filed 400 eviction complaints. These eviction filings lacked much detail about why the eviction was happening, and seemed to routinely lack basic information about lease terms or to include significant errors around lease states, and rental amounts, and payment information.

So the fact that a firm allegedly leaned on AI to generate a large number of eviction filings with false information, apparently without any meaningful review by an attorney, that's a big problem. Not only is the eviction illegal, but that eviction will live on in public records and hurt the tenant into the future.

AI is also increasingly part of how people buy homes. It is used in credit scoring models and automated valuation models which determine the value of a home. How AI is deployed has major implications for a person's credit scores, their mortgage rates, and whether home ownership and wealth building is even within reach.

We know that we have historic, systemic challenges with fairness and equity in this country. My own hometown of Minneapolis, has some of the greatest disparities in home ownership between Black and White families of anywhere in the country. So we need to carefully explore whether AI is extending and reinforcing these biases, and how it has the potential to correct them.

Our excellent witnesses here today have an unenviable task in your opening statements to ground us in both these opportunities and threats in AI and housing in 5 minutes. I look forward to hear-

ing from you, and I also look very much forward to the questions from my colleagues as we follow-up after your testimony.

With any innovation there are both opportunities and challenges that we need to balance, and our job is to think about these complex issues so that we can develop the best public policy. So I very much look forward to this conversation, and I now turn to Senator Lummis for her opening statement.

OPENING STATEMENT OF SENATOR CYNTHIA M. LUMMIS

Senator LUMMIS. Well, thank you Madam Chair, and thanks to our witnesses for being here today.

Yesterday, I was on an airplane sitting next to a woman who was flying into D.C. so the health care business could have their first workshop on AI and health care in some public health fields, and how to define AI and use guidelines in a way that will make the sideboards between existing technology and AI more clear.

You are way ahead of the game that health care was. And Ms. Rice, it's so nice to see you again because you are the one who came to my office and explained to me how AI can be used in a nefarious way. And before that, I never even would've considered that AI could be used in a way that reinforces bad behavior, and so I'm so pleased to see you here today. And all of you, thank you for coming. So now we're going to explore the range of promises and perils that artificial intelligence can offer to housing in the United States.

So there are many potential benefits from the use of AI in housing. Incorporating more information into housing decisions might expand access to credit. Automated tools may expedite the process of approving new housing developments. In rural areas, AI can ease the appraisals process and accelerate supply chains. Now, I'm the broken record on the Banking Committee for my support establishing a regulatory framework for digital assets. But I see some parallels where we have a new technology rapidly changing in a space, and a need for thoughtful consideration of how existing regulations may apply.

So from today's hearing, I'd like to better understand what parts of AI fit in today's regulatory framework and what needs to evolve. We've already seen straying into some areas that create perils because they're poorly thought out from the regulatory aspect.

This summer, in response to the potential use of AI when providing financial services to customers, the Securities and Exchange Commission proposed a rule called "conflicts of interest" in predictive data analysis. It was an overreaching rule, and the scope got into even using basic spreadsheets as covered technology. So what that did was put unneeded burdens on a new use of an old technology.

So we need rules and regulations governing AI to be scoped appropriately, not using it as an excuse to force decades-old technology to comply with a new and unnecessary set of impractical regulations. So I'm delighted to see you here to help us parse the difference. And this is a really qualified panel, so this should be a really informative morning.

Concerns about AI merit discussion. We need a financial system that gives families access to building wealth and create strong com-

munities through home ownership. Credit decisions should be based on who can repay not on the color of their skin or the community they call home. We should recognize the risk that AI can, if not trained on complete data and within safe guards, reinforce bias and discrimination. But aggressive action prohibiting innovative AI use is perilous, and we need to find that balance.

Today's hearing presents an opportunity to explore the application of AI within the housing space, and I look forward to hearing the testimonies of our witnesses. I'm optimistic about the prospect of harnessing AI to uplift and expand the American housing experience. So let's leverage innovation to improve access to housing in Wyoming and all over the United States. Thank you for being here. Thank you, Madam Chair.

Chair SMITH. Thank you so much, Senator Lummis. And I also am just really grateful that Senator Rounds is here. I know that you have done a lot of work in this space across the board and your leadership of that bipartisan task force that you and Senator Schumer are leading. So thank you very much. I'm glad you're here. Thank you.

So we have three witnesses with us today. I'll introduce each of them. Dr. Vanessa Perry is interim dean and professor at the George Washington University School of Business and nonresident fellow at the Urban Institute's Housing Finance Policy Center. Dr. Perry has extensive experience in the use of AI in mortgage finance and home ownership.

She previously held positions at the Department of Housing and Urban Development, the Consumer Finance Protection Bureau, and Freddie Mac. Dr. Perry earned her Ph.D. from the University of North Carolina at Chapel Hill and an MBA from Washington University in St. Louis.

Ms. Lisa Rice is the president and CEO of the National Fair Housing Alliance. Ms. Rice leads the Alliance's work on eliminating housing discrimination, and addressing the evolving role of artificial intelligence in housing. She has provided expert testimony to the Banking and Housing Committee, the bipartisan Senate AI Forum, and the House Financial Services Committee. Ms. Rice previously served as a CEO of the Toledo Fair Housing Center and the Northwest Ohio Development Agency.

And we have Nick Schmidt with us today, who is a partner and artificial intelligence practice leader at BLDS, and the founder and CTO of SolasAI—did I say that correctly? Close enough, he says. BLDS is a consulting firm that specializes in statistics and economics.

Mr. Schmidt's work focuses on algorithmic fairness, explainable AI, and model governance practices. He founded SolasAI, a compliance-focused AI platform that works to identify and mitigate bias in algorithmic decision making. He earned his MBA from the University of Chicago's Booth School of Business.

We'll begin with Ms. Rice, and just go down the line. You each have 5 minutes for your opening statements, and each of you have a clock in front of you which will count down the time. And your full written statement will be made a part of the record. Ms. Rice, you can begin.

**STATEMENT OF LISA RICE, PRESIDENT AND CHIEF
EXECUTIVE OFFICER, NATIONAL FAIR HOUSING ALLIANCE**

Ms. RICE. Chairwoman Smith, Ranking Member Lummis, and other distinguished Members of the Senate Subcommittee on Housing, Transportation, and Community Development, thank you for the opportunity to testify during this hearing on artificial intelligence.

The National Fair Housing Alliance works to eliminate all forms of housing and lending discrimination, and ensure equal opportunities for all people. We are the trade association for over 170 Fair housing organizations throughout the U.S. and its territories. We know that technology is the new civil and human rights frontier. That is why we are committed to creating automated systems that are fair, explainable, and trustworthy.

It is why we also create and promote policies, frameworks, and other tools to advance responsible AI principles. Artificial intelligence holds great promise for improving systems, democratizing opportunities, lowering costs, and increasing productivity. Yet it also holds great dangers for perpetuating bias, spreading misinformation, excluding people from necessary services, and generating other harms.

AI is used extensively in the housing and financial sectors, including in credit scoring, tenant screening, automated underwriting risk-based pricing, dynamic rental, pricing marketing, and many other areas. All these processes can manifest bias and extensive harm to consumers, communities, and our economy. Large language models like ChatGPT can also present challenges like data privacy, security bias, and accuracy concerns.

Through our work, we have uncovered discrimination in all these systems. For example, our investigation against a major insurance company in which we examined a third-party scoring system used by the insurer, showed the system was overcharging Black consumers at rates that exceeded their commensurate level of risk. Our findings in that case are similar to findings by researchers at Berkeley that found algorithmic pricing models discriminate against Black and Latino consumers by overcharging them to the tune of \$765 million per year.

Consumers who live in communities with few mainstream banks, like rural areas and communities of color, can often be negatively impacted by AI systems that rely heavily on data contained in the credit reporting agencies. For example, the CFPB reports that consumers in southern rural areas of the U.S. disproportionately have lower credit scores because, in part, they have lower access to physical and online mainstream financial services and products.

While AI can present serious threats, it can also be used for good. For example, NFHA uses AI tools to conduct research that is helping expand credit and housing opportunities for underserved groups and communities. Just this month, we released important research completed in conjunction with FairPlay AI, documenting a novel approach for de-biasing algorithmic models that simultaneously optimizes for high quality model performance.

Most previous research on fairness techniques have shown a tradeoff lower moderate, a model accuracy for increased fairness. The novel approach at the center of our research, which is based

on distribution matching, presents the ability to optimize models for both fairness and accuracy.

In addition to de-biasing systems, AI can be used to detect discrimination risks, build systems that can expand access to credit, provide financial services to consumers more efficiently, optimize privacy, identify barriers to fair and affordable housing and zoning ordinances, and other positive activities.

To ensure AI can be optimized for good, Congress must first do everything possible to ensure Federal agencies, researchers, and NGO's can apply existing laws and standards to the use of automated systems. Congress must ensure agencies have sufficient resources to tackle the issues, and Congress must pass legislation that both supports innovation and protects our society to ensure the U.S. can continue to lead the world in the development of the responsible AI. Thank you.

Chair SMITH. Thank you very much. Dr. Perry.

STATEMENT OF VANESSA PERRY, INTERIM DEAN AND PROFESSOR, THE GEORGE WASHINGTON UNIVERSITY SCHOOL OF BUSINESS; NONRESIDENT FELLOW, HOUSING FINANCE POLICY CENTER, URBAN INSTITUTE

Ms. PERRY. Thank you. Good morning. My name is Vanessa Perry. I'm a professor and interim dean at the George Washington University School of Business, and a nonresident fellow at the Housing Finance Policy Center at the Urban Institute focusing on consumers housing and financial markets. I want to thank Chairwoman Smith and Ranking Member Lummis for inviting me to testify on the impact of artificial intelligence, which is being employed increasingly throughout the housing and mortgage industry.

Based on my three decades of research in this field, regulation of AI warrants urgent attention. While these models can enhance efficiency, they can have unintended impacts on fairness and equity. Compared to traditional models, AI relies on a wider range of data inputs and more complex combinations thereof.

Although intricate multivariate algorithms have been used in the mortgage industry for years, AI models have the potential to incorporate new data sources due to their complexity. It is difficult, but not impossible for anyone other than AI developers to scrutinize and monitor their inputs.

AI is already widespread in the mortgage market. AI digital marketing models target prospective home buyers, and communications with customers are intermediated by AI chatbots. Credit scoring companies and mortgage underwriting systems use AI to evaluate credit risk. AI models are used for property valuation, loan servicing, and loss mitigation. Because these models rely on historical data, there is the potential for them to systematize and amplify discrimination and inequality. For example, due to the legacy of redlining and segregation, and their effects on present day neighborhood conditions and home values, why should we expect AI to produce estimates that are both accurate and fair? And absent guardrails, how would we know if AI models were to incorporate data elements such as GPS location that serve as a proxy for race, gender, or ability?

AI models are not subject to human error, and they enable more accurate and consistent decisions. Depending on how they are developed, their enhanced capabilities could expand access to home ownership for households currently underrepresented in the mortgage market. For example, AI can produce faster and less subjective estimates than human property appraisals, and can devise credit scores for those who lack a traditional credit history.

To address concerns about AI's impact on access to the housing finance system for underrepresented communities, my colleagues and I have proposed five factors summarizing the societal, ethical, legal, and practical issues that should be considered in the development and implementation of AI. They form a memorable acronym, SCALE.

First, societal values. Algorithms tell us what factors the developers think are important in what order, and to what degree. AI models should consider the socioeconomic and historical context such as past discrimination, and should align with prevailing legal and ethical paradigms, such as disparate impact law, individual freedom, privacy, and racial equity.

Next, is contextual integrity. Model inputs should be relevant to the mortgage and housing domain, and may differ substantively from those used for other or less consequential context.

Then, accuracy. Models should be reliable, error-free, unbiased, and representative of all demographic and income groups across varying macroeconomic conditions.

Next, legality. The model and its inputs should not incorporate characteristics protected by fair lending laws, or generate unjustified disparate impacts based on these characteristics.

Finally, expanded opportunity AI models should significantly increase access to credit in addition to offering greater cost efficiency or risk assessment benefits. This criterion represents a higher bar than existing regulatory frameworks, and has perhaps the most promising impact on the economy and communities.

This framework could inform new or expanded regulations such as guidance for the use of certain types of data, for example, an individual's social media profile for certain purposes, such as mortgage lending decisions. If designed to do so, AI models can increase access to home ownership and eradicate the effects of systemic discrimination while increasing accuracy and efficiency in the mortgage value chain.

We need new laws on the Federal level that utilize the scale factors previously enumerated at every stage of the AI life cycle. Thank you, again, for the opportunity to testify here today, and I look forward to answering your questions.

Chair SMITH. Thank you very much. Now we'll hear from Mr. Schmidt.

STATEMENT OF NICHOLAS SCHMIDT, PARTNER AND ARTIFICIAL INTELLIGENCE PRACTICE LEADER, BLDS; FOUNDER AND CTO, SOLASAI

Mr. SCHMIDT. Thank you, Chair Smith, Ranking Member Lummis, Senator Rounds, and the esteemed Members of the Subcommittee. Thank you for this opportunity to share my experience

and insights into how AI is revolutionizing the housing sector, and how we can make it safer, fairer, and more accountable.

I'm Nicholas Schmidt, the founder and CTO of SolasAI. I'm also the practice leader at BLDS. I've spent nearly 25 years working to ensure that data, algorithms, and statistics are applied fairly. And I've spent the last decade devoted to the responsible and ethical application of AI and machine learning, and with a significant portion of my work focused on credit and housing.

Examples of AI, including the amazing ability of generative AI, and the horror of deep fakes and what they might do to our democracy, drive our fear, our fascination speculation. But really, as amazing, and as scary as those truly are, the more mundane examples of AI in machine learning are just as important, especially in housing.

We must recognize that these technologies are already shaping many aspects of the industry and getting distracted by those sort of shinier aspects of AI, may take our eyes off the ball at a time that is really critical.

So what we should consider is that while an AI algorithm could help enable people to get a home loan for the first time, it could also drive a decision to evict someone from their home or apartment without sign sufficient justification. Because of the former, we want to encourage the adoption of the technology, because of the latter, we want to make sure that it's only used responsibly. This is only possible through effective regulation.

When considering this, what is really important to remember is that these systems are driven by human decisions. Unfortunately, this human-centric nature of AI development is overlooked, but AI is not self-operating, and claiming so is a dangerous notion. It absolves us of the responsibility, and also ignores the numerous opportunities we have to shape AI for the better.

Just to illustrate this, let's consider a mortgage underwriting algorithm. I'm going to go through some of the ways in which a human impacts how that algorithm ultimately operates determining what data to use. Should you use something like education or purchase history, how do you define delinquency? That's going to make a very large impact on people who have income insecurity, but are ultimately able to pay their mortgages over time. That'll make a big difference over whether or not they'll actually get a loan.

What algorithm should the lender use? That affects not only discrimination, fairness, accountability, but also potentially the safety and soundness of the lender? No algorithm is perfect. We have to make a decision. Is the algorithm good enough to put into production? And finally, someone has to decide how to use the algorithm. All of these decision points offer an opportunity for smart intervention.

And when considering such regulation that comes out of this intervention, we should remember that these technologies are not new to housing. We do not need to reinvent the wheel. Instead, we should learn from the practices and frameworks to create a strong and reasonable regulatory environment.

In particular, there are three things I'll, I'll mention. The Federal Reserve's SR11-7 Policy, the NIST AI Risk Management Frame-

work, and SP 1270 guidance and the disparate impact and disparate treatment framework from the Fair Housing Act, and potentially E-COA—we shall see what the courts say about that. But out of these frameworks and practices, I want to distill this down to four principles. Fairness, transparency, accountability, and materiality.

Fairness, disparate impact and disparate treatment. And regulators should set expectations that bias and discrimination should be measured and mitigated in AI systems. But really, importantly, fairness is not just about antidiscrimination, and it is an issue that everyone here has a self-interest in.

What is being put into some, not all, but some AI systems is not fair to anyone regardless of their race, ethnicity, sex, or age. We need to empower regulators to set and enforce principles-based standards for dating inclusion.

Transparency. If a person receives a negative outcome from an AI-based decision, they should have the right to understand why.

Accountability. If you are negatively impacted, you should have a right to appeal if it's a sufficiently high-stakes decision. Additionally, people or entities that deploy AI systems irresponsibly should face the appropriate consequences.

And materiality. Businesses need to understand and make risk-based decisions. You can't have, for example, an Excel spreadsheet subject to the same regulations and, you know, requirements and oversight as a credit underwriting algorithm that is scoring millions of people. Doing so, ultimately, is ineffective for all of our goals.

So by considering these four sets of principles, we really can develop a balanced regulatory approach. And by fostering collaboration among all of the stakeholders, we really do have a potential to harness AI's value in housing. With this, I really do believe that AI can serve as a force for good. Thank you very much for the opportunity to share my insight.

Chair SMITH. Great. Thank you so much to all of our panelists. And we'll now begin a round of questioning from colleagues. And I will start.

So each of you, it seems to me, are highlighting both the opportunities and also the risks of AI. But I'm struck by how—I guess I want to dive in a little bit more into how you—we know that AI used in housing could pose significant risks. But I'm wondering if you could talk a bit more about what the beneficial applications for AI might be.

Let's say just take for example a goal of figuring out how to increase home ownership amongst, traditionally, you know, people who've been traditionally unable to buy their own home. AI is a tool. How could AI, just as an example, be used as a tool to advance that goal? And anybody, I'd love to hear what you have to say about that, if it's implemented correctly is right.

Ms. RICE. I can start. As I shared, we have just released research this month using AI to help expand access to credit for people who are creditworthy and have been kept out of the financial mainstream. There are many organizations, and our research shows that we can increase fairness and access by about 13 percent for underserved groups.

There are many organizations—well, I take that back, not many, but there are some organizations that are working on using AI and other forms of technologies to incorporate factors that are currently not included in the underwriting systems and the credit scoring systems that we utilize today in order to determine if people can access credit opportunity.

So I'll give you an example. Rental housing payment information. This is a very viable factor. It is an important and highly predictive variable that is not included in the underscoring systems that we utilize today. That missing factor, that missing information, that missing data could really open up the aperture for opportunity, particularly for people who live in rural communities.

As I stated before, the CFPB's research and other research shows that people in rural areas have disproportionately lower credit scores, not because they're less creditworthy, but because they don't have access to the mainstream financial services that regularly report positive data to the credit repositories. So if you access credit outside the financial mainstream, using rental housing payment information can really increase access to affordable credit opportunities.

Chair SMITH. So if AI is including sort of better information, less biased information, potentially information that isn't included in regular decision making, you could actually have a net positive impact on the fairness of the system.

Ms. RICE. That's right.

Chair SMITH. Right. Thank you. Let me ask kind of a different question. Like, how do we address this so-called black box phenomenon in AI where, you know, you have an AI tool that includes machine learning, so it is constantly learning and changing the way it's, you know, improving supposedly the way that it's making decision-making, but it seems to sort of spit out an answer at the end.

And the question is, like, where did that answer come from? Could maybe Mr. Schmidt or Dr. Perry, could you just address that? How do we actually get accountability when you have that kind of a technological tool?

Mr. SCHMIDT. So there are two issues here. One is directly your question and how do we deal with these incredibly opaque algorithms? But then there's another important piece, which is, should we even be worrying about this in high-risk situations, or should we be using more interpretable—requiring more interpretable, more understandable algorithms?

And the way that this is first sort of surfaced over the last 5 or 10 years is been this idea that in order to have accuracy, you had to have an opaque black box. By being able to incorporate lots of interactions and nonlinear relationships and all sorts of difficult math, you would get a more accurate measure of, say, credit outcomes.

What has happened over the recent years is that people have realized that you don't need to have these incredibly opaque algorithms. You can actually develop ones that are pretty understandable. And so I believe that those are going to become much more commonly used in these high-rate, high-stakes decisions.

Chair SMITH. Interesting. Thank you. Dr. Perry, I'm out of time, if there's anything that you'd like to add, I'd love to hear from you.

Ms. PERRY. I just wanted to add that because these models are dynamic, they change often. The question is, which version of a model even would you scrutinize or monitor?

Chair SMITH. Yeah. Thank you. Senator Lummis.

Senator LUMMIS. Thank you, Madam Chair. Mr. Schmidt, welcome. You have an interesting job and an important one in terms of helping us understand the four criteria you mentioned in your testimony. With regard to transparency, what level of explanation do you think companies and organizations should provide to consumers?

Mr. SCHMIDT. So I think that there are two principles that, that need to apply. One of them is transparency into the process that they have undergone, and that will allow people to understand whether or not the data that was used was fair as well as the process, and whether or not the data was right. So were they using some factor, you know, education in a situation where education really didn't belong, or maybe they got the education of the person wrong. Being able to understand and appeal those decisions is very important.

The other area where this is really important is having some sort of understanding of how to improve one's outcome. So you got rejected and you got told that, you know, had, I don't know, number of delinquencies in the past 60 days. Maybe that one's easy to figure out how to rectify, but in these complex systems it could be much more difficult than that. And giving people a clear framework to move from being rejected to being accepted, I think could be a very helpful public policy goal.

Senator LUMMIS. Another question, this one about data privacy. How do you think companies that use AI models for housing should be thinking about data privacy?

Mr. SCHMIDT. I would like to preface this by saying that I'm not a data privacy expert, but what I have seen in my experience, particularly with large lenders and other organizations that are really leading the financial system, is that data privacy is very high priority, and making sure that there are not data leakages as people spend inordinate amounts of time on. So I think that that is something that is very well recognized and something that the industry is really striving toward.

Senator LUMMIS. So what is the weakest link here? You know, when you named those four criteria, which one is crying in need of improvement?

Mr. SCHMIDT. The one that keeps me up at night is fairness. And this is both from, again, antidiscrimination in the traditional sense of the word, as well as fairness for everybody in this room. The quality of models being built across all industries is low. That is absolutely not universal. There are brilliant people doing very brilliant work, but there is a lot of stuff going on there that is extremely low quality.

And what we need is to have the same kind of, or similar, model governance standards that were put forward by the Fed and the OCC, put into lower tier lending institutions and places like the health care industry. Because you are absolutely right, that the

health, the financial system, financial industry, is light years ahead of every other industry. And the potential for damage if strong model governance is not put into place, astronomical.

Senator LUMMIS. Can you help with that? Have you studied this enough that you're beginning to see how to break through those areas where the data is being used in a very unsophisticated way?

Mr. SCHMIDT. These are mostly people problems. The problem with data scientists and data science is that there are—colloquially, I'd say they're drunk on algorithms. They get very excited by the math, and they get very excited by the next new thing, and they don't take a step away and say, "What is it I'm doing?"

And so, I think that the first step is making people answer questions, making the model builders answer questions. What is it I'm putting in here? Why am I putting it in this model? What purpose does it serve? And the way that that can happen effectively is by requiring effective model governance.

And if I may, the other thing that is really important is having lawyers in the middle of the room. And in the questions, I always tell my clients that lawyers are the best data scientists because they're the most skeptical and negative people in a room, and that's really what is needed in data science. So we need lawyers, and compliance, and people like all the ones who are laughing in this room, because you are the people who are going to question the data scientists and get them to start doing better work.

Senator LUMMIS. Well, thank you. As a lawyer, I resemble that remark.

[Laughter.]

Senator LUMMIS. So thank you very much. I yield back. Thank you, Madam Chair.

Chair SMITH. Thank you, Senator Lummis. Senator Menendez.

Senator MENENDEZ. Thank you, Madam Chair. And as a lawyer, I'm not going to speak as a lawyer today, but as a legislator, so I can avoid that.

Because of the way AI learns by studying large sets of data, if there is a bias in the underlying data, it can become encoded in AI's decision-making process. And when it comes to housing, the historical data reflects years of redlining, disparities in wealth, access to credit, appraisal bias, among others. So Ms. Rice, Dr. Perry, how can we work to ensure that AI doesn't stratify or even expand existing disparities in housing?

Ms. RICE. Sure. I could take a first crack at that. There are multiple ways that we have to work to make sure that AI doesn't magnify existing inequities. So one way is to take steps to audit and monitor AI, to examine it preproduction, predeployment, and also post-deployment.

We have developed, actually, what we believe is a state-of-the-art framework. It's called purpose process and monitoring. We believe that it's a state-of-the-art auditing model framework to help ensure that algorithms are fair, and equitable, and that they do not pose harm to consumers. Unfortunately, we do know, Senator Menendez, that many models that have been deployed are exacting harms on consumers. We have clear evidence of that. So what the purpose monitoring process and monitoring framework does, is it helps model developers answer all of the questions that Nick just talked

about. Right? Is this a good use for whatever problem it is that you're trying to solve? Does it need an algorithmic or an AI solution? Because there are many problems that that need to be solved that don't necessarily need an AI solution.

It examines the data, it examines the model development, and the efficacy and accuracy of that model. And then, it also looks at the potential and examines the potential for harms. And it identifies, or I should say, compels model developers to look at less discriminatory alternatives while they are designing the model.

Models can also change post-deployment, right, their effectiveness. Just because a model is effective on day one doesn't mean it is going to be still be highly effective and accurate on day 200. So you have to continually be auditing and monitoring that model to make sure that it is still efficacious. Thank you.

Senator MENENDEZ. Dr. Perry, anything you want to add to that?

Ms. PERRY. Yes. I would just like to add that in order to, sort of, prevent these models from producing, reproducing existing discrimination or past discrimination, the one thing to do is to intentionally design models that explicitly remove those effects and reduce those effects. And actually, that's one of the great promises of these tools; is that they can be programmed explicitly to remove the effects of past discrimination because they are more capable of identifying them than the sort of traditional models that we've relied on in this industry for so long.

Senator MENENDEZ. Yeah. Let me follow up this line of questioning. That problem is further magnified by the challenge of explainability. That is the model—the results of AI models and the steps that led them to where they are can be difficult to parse due to the complexity of the algorithms and the underlying data. How should we think about issues of accountability and oversight when it comes to these AI models and the data that they use? And I welcome anybody's.

Ms. PERRY. I'll start. My opinion is to effectively monitor, we need to examine, one, the guiding principles that are used to create these models, the inputs which data could potentially be used and which are prohibited by the models. And then finally, the effects; so the entire process with the understanding that regulatory agencies may not necessarily have the resources to do the volume of monitoring and enforcement that we are accustomed to. So we definitely need stronger protections and processes in place to make this happen.

Ms. RICE. And Senator Menendez, I'll just add that a couple of weeks ago, we held the first our inaugural responsible AI symposium. And I had the privilege of hosting a segment with the heads of all of the Federal regulatory agencies. And I posed this question to them, and they each said that we currently have existing laws on the books that effectively govern many aspects of AI development and usage. And they are working to make sure that their agencies are adept and can enforce and exercise effective oversight over the use of AI, particularly in financial services.

Now, one of the challenges is getting staff that are educated and trained on these systems and usages. That is a huge lift. I will tell you, when we stood up our responsible AI division, we had to comb the globe to find someone to head it up. And we did find someone

in Canada, not in the United States. So we have to really do a good job of making sure that our citizens, our students, are well educated in this area.

The other thing that they mentioned is that they need resources. They need effective resources, and they need to be using AI tools themselves in order to effectively police this area.

Chair SMITH. Thank you very much. Senator Rounds.

Senator ROUNDS. Thank you, Madam Chair. And let me just, first of all, say thank you for doing this hearing today for both of you.

The AI working group, when we started the program, one of the things that we talked about was is we wanted every single committee to be able to participate because of the expertise on the committees and the individuals who participated in meetings like this to share their points of view and their recommendations. So I really appreciate your jumping on as early as you have. You're probably one of the first subcommittees to actually have a significant opportunity to talk about a specific area in which AI does make a difference.

And I also want to say, Ms. Rice, what you've indicated in terms of the personnel issue is something that we have to take into account. And we're never going to have enough experts on this, and so to be able to share those experts is one of the reasons why NIST, I think, is going to be very important, being a location where we can find individuals from industry who can come in and participate and help to bring that expertise to the committees, as well as to the organizations responsible for overseeing AI development.

I also wanted to—Mr. Schmidt, I think you would agree with me on this, but I want to get your thoughts. I think that a company is responsible and has to abide by the Fair Credit Reporting Act, the Fair Housing Act, the Equal Credit Opportunity Act, regardless of what technology that they're using, AI or otherwise. Correct?

Mr. SCHMIDT. Correct.

Senator ROUNDS. So really, one of the challenges we've got is how do we make sure that within an algorithm or within a decision-making process, because AI really is simply a very, very fast decision-making process that is now available to us, but one in which a computer system and a use of formulas actually learns from itself, whether or not it was making correct decisions and modifies it, as Dr. Perry indicated, can modify itself along the way.

And so what I'm curious about is, I don't think we can be afraid of the technology, but I think we've got to be able to leverage the best of it while at the same time being able to provide the regulators the ability to ask the questions and to make modifications or request modifications, should there be biases identified. Can you think of any program that we've ever made that doesn't have some biases built into it?

Mr. SCHMIDT. No.

Senator ROUNDS. So really, the challenge for us is how do we identify those biases? And in fact, any adult who is responsible for making decisions today, is there a human that doesn't have biases built into their decision making?

Mr. SCHMIDT. I've yet to meet one.

Senator ROUNDS. Yeah. So really, the question for us is how do we overcome those human attributes and allow a machine to do a better job of making decisions that eliminates the biases that we're trying to get out of our decision-making processes. Fair enough?

Mr. SCHMIDT. Yes.

Senator ROUNDS. OK. I know that we're going to have a lot of discussions about the problems and the challenges that we're going to have anytime we bring in a new type of technology. But I wanted to close with this, and it's going to come right back down to what Ms. Rice had indicated. Suitable AI talent is going to be in high demand.

JPMorgan said in May of 2023 that it had hired 900 data scientists, 600 machine learning engineers, and 200 AI researchers to execute its technology initiatives. Google has hired thousands of researchers and engineers to work on machine learning and AI.

However, I believe that it's critical that AI solutions reach beyond our largest companies. We learned at the end of the last year that between 2022 and 2023, the number of people experiencing homelessness on a given night increased in the United States by 12 percent to 653,104 people. The highest number of people recorded since the inception of the annual count in 2007. Researchers have used AI to identify individuals most at risk for losing their housing to target with homelessness prevention assistance.

Now, I think this is one of the things that we should do a better job on, and if we can eliminate some of the challenges we have, just in terms of the number of individuals that we have to impact, AI could actually help us identify those individuals before they become homeless.

Mr. Schmidt, could you discuss what steps we need to take now to provide the AI workforce of tomorrow, and how we make certain that AI talent is reaching areas like homelessness prevention?

Mr. SCHMIDT. To address the latter half of your question, I think that it is very important that we encourage organizations like Ms. Rice's, as well as academia, and the type of work that Ms. Perry is doing, because quite frankly, industry is not likely to jump on those sorts of things. And we have to put the incentives in the right places and give the right people the right money and ability to put forward those sorts of programs. So that's what I would say about that.

And then in terms of the talent gap, what we really need to do is concentrate on both STEM and liberal arts. What we need are well-rounded people. I've worked in tech-type companies my entire career, and I've always liked hiring English majors and history majors, and teaching them how to do statistics, because they're the ones who are really thoughtful.

And so what we want to do is we want to encourage that. We want to encourage an educational system that is focused on questioning things, really getting into why are we doing things? And then the technology is secondary. My company, actually, so I founded it out 4 years ago. The first Python program I ever wrote became my company, and so I learned that. You know, I guess you can teach an old dog new tricks. I learned it later, and I think that many people can learn those kinds of things if they're given the right education to start with.

Thank you. And thank you for your leniency on the time today, Madam Chair.

Chair SMITH. Absolutely. Thank you, Senator Rounds. Senator Warnock.

Senator WARNOCK. Thank you, Madam Chair. And amen to Mr. Nicholas Schmidt on hiring folks with education and the humanities who think critically about these issues. Georgia is in a housing crisis, and it is not getting better. Since 2019, rent in Georgia has increased 13.7 percent. This is one of the highest increases in the country, but Georgia's not alone. According to Moody's Analytics, the amount that Americans are spending on rent is at a record high. For the first time last year, the average apartment cost over 30 percent, over 30 percent of a family's monthly income.

One factor driving rents up may be the use of AI software by some of the nation's largest and most powerful rental property management companies. As a *ProPublica* report details, this AI software looks at data across multiple companies and multiple buildings in the same area, and then recommends a monthly rental price. I am concerned that these tools may enable property management to coordinate pricing, coordinate escalation, inflate rents, and stifle competition.

Ms. Rice, folks across Georgia are feeling the rental pinch. How can we ensure that AI driven rental software is not being weaponized to artificially inflate the rental market?

Ms. RICE. Senator Warnock, thank you for that question. And you're talking about dynamic pricing systems, and you're absolutely right. We are very concerned about these systems and believe that they are contributing to the increase in rental housing prices.

And because the systems are operated by a few companies that amass data amongst a number of different landlords throughout a region, there is a great propensity for reducing competition because those systems are conveying to all of those different landlords, essentially, the same dynamic, a kind of dynamic pricing.

So a land an apartment complex that may not be getting a lot of inquiries can still elevate their prices because their competitors are getting a lot of inquiries. Right? So it does really work sort of cross purposes when you think about fair and affordable housing principles.

The other danger is that there is very little transparency around these systems. So if a consumer wants to know why did the price increase from day one to day two the only information the consumer gets back is, "Well, this is what the algorithm is saying the price is today." There's no explainability or transparency around why that price increase happened. So unfortunately, this is an area where we do need increased regulation.

We're trying to use current fair housing laws to tackle this problem. We're also trying to use current antitrust laws to tackle this process, and we're not getting where we really need to be. So we do need increased regulation here. The only thing that I will add is that local jurisdictions may be able to use tools that they have to tamp down on the utilization of dynamic pricing algorithms, but that remains to be seen.

Senator WARNOCK. So in an ironic way, and AI is here to stay, which is why we're having a conversation, but in a way, it sort of

mitigates against natural competition in the market that this is falsely derived through these algorithms, and they're opaque. It's difficult to get through to ask the kinds of questions that we need to ask.

So this is a unique threat that Congress has to address. This is affecting renters, but home ownership and the impact of these algorithms are impacting communities of color. Buying a home access to credit has been a long significant barrier. Techniques like machine learning and artificial intelligence can improve credit access by incorporating more types of data, like paying utility bills and rental history to demonstrate that someone is creditworthy, but they can recreate in code historical disparities and discrimination in the housing market.

So, I'd like to quickly ask Dr. Perry. Dr. Perry, as credit scoring models become increasingly sophisticated, how do we ensure that these models do not perpetuate through the algorithm's historical discrimination?

Ms. PERRY. Thanks for the question. In order to ensure that these models, many of which operate in a sort of black box which is one of the key issues that we are concerned with, is really to require that we adopt ways to monitor and scrutinize every aspect of the model development process.

I can't say it enough how important it is to have a set of guiding principles or goals that drive the development from the outset. Then we need to know exactly what inputs are currently being used and what inputs could potentially be used. Then we need to look at the way those inputs are combined. And finally, we need to do effects testing, but we can't just look at effects the way that we currently look at, say, test for disparate impact, because this is far more complicated than that and requires far more information.

And, you know, these models take, you know, a long time to develop, but they also are changing on a dynamic basis, and that's why it's really important to know what the goals of the model are upfront.

That said, I think there's really great potential because for everything you can program the AI to optimize, you can also program it to optimize opportunity and fairness. You can tell it explicitly, do not look at factors that are correlated with race, ethnicity, ability, or any other kinds of protected factors or prohibited factors.

Chair SMITH. Thank you. Thank you very much.

Senator WARNOCK. Thank you very much, Madam Chair.

Chair SMITH. Yeah. Senator Cortez Masto.

Senator CORTEZ MASTO. Thank you, Madam Chair and our ranking member for this great conversation. I'm going to follow-up on this, but I'm going to ask Dr. Schmidt. If you would, Mr. Schmidt, talk a little bit about automated valuation models and home appraisals. Right? There's positives because it should lower costs. It should hopefully decrease bias, which we have been talking about. But you also note that the data used to build an automated evaluation model will be biased, could be biased. Right? And you have software program to address this issue. So, talk to us how you're decreasing that discrimination as the initial data going into these models.

Mr. SCHMIDT. Sure. So the problem with any kind of pricing model, automated valuation model is that it is necessarily going to be looking at location, location, location. You know, that's what drives price. Well, what do we know about locations and housing in America? It's got a history that goes up to today of discrimination. You know, that has not ended.

And so if you are building a model that incorporates information about location, which you have to do to get an accurate model, you are building a discriminatory model. But there are two things that kind of make it better. One of them is relative to the alternative. If you think about the potential bias, not necessarily, but potential bias of a human appraiser, and there's been a lot of investigations and work on that. It's shown that there's a really high risk of human appraisal bias, very unlikely that that will be explicitly built into an automated valuation model.

So what that means is that within a particular geography that you're controlling for in one of these models, you're likely to remove that idiosyncratic bias from the appraiser. But the fact that you are controlling from one geography to another means that you're incorporating that geographic level discrimination.

So I think of AVMs as having—they're a step, maybe a small step in the right direction. They are also much cheaper in rural settings. My understanding is in rural areas, it can be impossible to get an appraisal, a human appraisal. If there's good enough data, then an AVM could provide one. So there are many reasons why an AVM would represent just a sort of uniform improvement relative to current not ideal practices.

Senator CORTEZ MASTO. So let me ask you this, because this is a concern. I see where we're going, and I think my colleague's absolutely right. AI is here, and people are going to start using it, and they'll start using an appraisal process.

I know in Las Vegas, in the metropolitan area, home values in a majority of our Black neighborhoods were devalued by at least 12 percent. That's over \$22,000 difference compared to neighborhoods with fewer than 1 percent of Black residents. How do we ensure that that bias doesn't get put into these models?

Mr. SCHMIDT. So it will to some—

Senator CORTEZ MASTO. So then how do we adjust for these differences?

Mr. SCHMIDT. Yes. But there is yet another aspect to it, which is—and Ms. Rice was talking about it, and it's actually the center of my work—called Searching for Less Discriminatory Alternative Models. And it actually works better in AI and machine learning.

And what it turns out in machine learning models is that you can have thousands, really an infinite number of ways to specify these models, including certain information, excluding other information, changing what are called hyperparameters, which are kind of the knobs and dials that slightly change the algorithm. All of these can have, or changing the algorithm itself, all of these can have an effect on the bias as well as the quality of the model.

And machine learning, it turns out, there are just way more opportunities. And so what I've developed, and what Ms. Rice's organization has developed, and other very good organizations are

working on, are finding those fairer models, and that's where things can continue to improve.

Senator CORTEZ MASTO. OK. Any other comments, Ms. Rice?

Ms. RICE. Thank you, Senator Cortez Masto. I will say and remind us that not every issue warrants a technological response. So there is still—let's still remember that we have human beings that are highly trained to perform these types of appraisals, and we have to use them.

Now, that said, AVMs, while they can be a useful tool, they are not as accurate in every community, so they don't have the same level of performance. And so we do find that AVM models are less accurate. In fact, in some communities, they have a 15 percent degree of accuracy.

So I think you'd agree you would not want to use an AVM model in a community that has a 15 percent degree of accuracy. That we're finding that in rural areas, and we're finding it in predominantly Black and Latino areas. AVM models are much more predictive in areas that have a lot of new housing development, and that are predominantly White.

So, we are working on tools to make AVM models fairer. What will make AVM models fairer is if we change the construct or the approach that we currently use to determine property values. Right now, we use a sales comparison approach, which as Nick has already said, requires that I look at what the house across the street from me sold, or the house around the block from me sold.

We have worked extensively with the insurance sector to change their property valuation models to make them fairer by considering the cost of reconstruction, and making sure that we have accurate data that can assess what is that reconstruction cost, and that gets embedded or enfolded into the valuation process.

Senator CORTEZ MASTO. Thank you. Thank you very much.

Chair SMITH. Thank you so much. Senator Lummis. We're going to move to a second round of questions for anybody who has them.

Senator LUMMIS. Thank you, Madam Chair. And I'd like to start just by saying, Mr. Schmidt, you finally gave me an explanation for what I can tell people. My son-in-law does—and I've wondered how to explain what he does for years. Now I know. He manages teams of people who are drunk on algorithms, to make a product that's useful by sober people.

[Laughter.]

Senator LUMMIS. So thank you for that. I finally know how to explain what he does. Panel, thank you so much. I have the same question for each one of you to conclude my questioning. And I'd like to start with you, Dr. Perry. What regulatory approach would you advocate for this Committee to pursue? What guiding principles can you point us toward? How can you help us be relevant in this discussion?

Ms. PERRY. Thank you, Senator Lummis, for the question. I think that of the most important consideration is holding AI accountable for expanding opportunity. More so than anything, I think that is of critical importance. It's also important to move really very quickly because the industry is already adopting and applying these tools in ways we are not even aware.

Senator LUMMIS. Thank you. Ms. Rice.

Ms. RICE. Senator Lummis, thank you so much for the question. First, I'd say I'd echo what I said earlier, that we already have a bevy of existing laws and regulations that apply in this sector. And so, what we can do is help you understand the extent of that. And one of the things that we are thinking about is pulling together sort of a matrix, you know, some agencies that explain what are the existing laws and regulations that are in place, and then understanding what we need to do to enforce those existing laws and standards.

But the second, I will note, is that there are gaps. There are clearly gaps, and we're finding that—we talked about a gap with the dynamic rent, the rental dynamic pricing schemes. And so, listing where there are gaps, and then helping legislators understand what it would really take to fill those gaps.

And here I do advocate for an intersectoral approach, because you're going to need academicians, you're going to need developers, you're going to need industry, and you're going to need civil rights and human rights organizations to understand, you know, how to design the legislation so that we are addressing those gaps that do exist.

Senator LUMMIS. Thank you. Mr. Schmidt.

Mr. SCHMIDT. So, I would really repeat what my copanelists said. I think we need to move fast. I think it's extremely important to recognize that existing regulations already cover 90 percent of it. What I have seen the Consumer Financial Protection Bureau, CFPP, and the OCC doing, I cannot say enough for. I think they're really moving in the right direction. They're doing excellent work, and they should be empowered to keep going in those directions. What I see is that industry is responding in a way that is really right.

The other thing I would say is just that we need to make sure that our regulators have the education, they have the experience that they need. What the bureau has done from 2016 to today, I think they've done an amazing amount of self-education and bringing on really good people. And they're doing a great job considering the resources they have, but they could use more.

Senator LUMMIS. Thank you. And Ms. Rice.

Ms. RICE. Senator Lummis, can I just mention one thing? Section 230, the Communications Decency Act. We are going to have to revisit it.

Senator LUMMIS. Thank you for that. This has been a great panel. Thanks for your work in this area. This is all new to me. So your guidance is deeply appreciated. Thanks, Madam Chair, for holding this hearing. I yield back.

Chair SMITH. Thank you so much. I just have a follow-up question, kind of along the lines of what Senator Lummis was asking about. There have been some efforts at the State level creating legislation and guidance around AI implementation. Mr. Schmidt or anybody, would you just comment, do you see anything out there at the State level that you think is good or something that we should be cautious about?

Mr. SCHMIDT. The New York Department of Financial Services just released a circular for insurance. And of course, I could quibble with a few things, but really, I could have written the thing myself

and been proud of it. I think that what they have done is excellent and really principle-based, and based in existing regulation and law, and could be very effective.

Chair SMITH. That's great. Anybody else have a comment about State-led efforts that you think are useful or something to not do?

Ms. PERRY. There a number of States that have done some really impressive things. And the concern that I'm have now is that they're going to be different from each other, and it's going to create havoc for companies, and in particular, it's going to be problematic for smaller businesses to comply when there are different State regulations, and something that is by definition, crosses State boundaries.

Chair SMITH. Yeah. Thank you. That's great. Ranking Member Lummis, I have a feeling that you and I could—we have easily another hour-and-a-half of conversation with these excellent panelists about this topic. And I think that there's a lot more work for us to do around this idea of what model governance standards might look like. The President's executive order on AI, the AI Bill of Rights, and you all have laid out some possibilities. There's lots of work there.

But I'm struggling to really understand like how that becomes implementable, you know, in industry as you're trying to figure out—you know, as industry is trying to figure out what the rules are. And I'm also really struck by how one of the things that happens in AI is there's an imbalance of information.

Like, information is power. And if you have more information as a seller than a buyer, then you're going to have more power. I'm seeing this right now in my own home State of Minnesota, where we have these big institutional investors who are able to within a matter of minutes, when a house comes on the market, decide whether it's something that they want to buy or not, while an individual home potential homeowner can't possibly compete with that. And so that shifts the balance of power in how real estate market has worked in the past in ways that I think are concerning.

So I always think that if you want to understand the outcome, you need to look at what the incentives are. And to me, the question is how do we build incentives into this system so that we are maximizing fairness, and accountability, and transparency as we as we try to set some guidelines both to encourage innovation, but also to make sure that these systems are working for the public benefit.

So, thank you very much to all of you for this. Senator Cortez Masto, do you have a follow-up question at all?

Senator CORTEZ MASTO. Just a comment, and I appreciate the conversation. Let me just say I heard the be thoughtful, reasonable, and rapid three things Congress is challenged at doing. And so we are doomed from the start.

But it really would be helpful to get your thoughts maybe on the top three things that we can focus on in the housing space, or the best model that a State has put forward. I understand the patchwork concern, but if there's a good model that can be brought to at a Federal level, I think that's worth looking at as well. And then finally, thank you, Ms. Rice, for talking about CDA and 230. I

couldn't agree more. We have to do something about that. So, thank you.

Chair SMITH. Great. Thank you so much. Well, thank you to our witnesses for being here today and for providing such excellent testimony. Before we adjourn, I would like to enter several statements into the record from the Urban Institute's Housing Finance Policy Center, and from Zillow, as well as reports and articles from Dr. Perry. Is there any objection. Without objection, those will be entered into the record.

For Senators who wish to submit questions for the record, those questions are due 1 week from today, which is Wednesday, February 7th. For our witnesses, you will have 45 days to respond to any questions for the record, and thank you again. And with that, this hearing is adjourned.

[Whereupon, at 11:21 a.m., the hearing was adjourned.]

[Prepared statements, responses to written questions, and additional material supplied for the record follow:]

PREPARED STATEMENT OF CHAIR TINA SMITH

Today's Hearing will focus on the promise and threats that Artificial Intelligence poses in the housing sector, and I am very much looking forward to our witnesses testimony and this conversation. I want to thank Ranking Member Lummis and her staff for our ongoing bipartisan work as we put together this hearing. We both share, I believe, a deep interest in how we can develop Federal policy that supports innovation and expands opportunity for everyone to have a safe, decent affordable place to live.

And one of the most consequential innovations in recent years is Artificial Intelligence. Leader Schumer, Senator Rounds, Senator Young, and Senator Heinrich are leading a bipartisan effort to explore the impacts, opportunities and threats that AI poses. And they have asked Senate Committees to engage in our areas of expertise, which leads us to this Committee hearing today, examining what AI means for housing.

Without a safe, decent, affordable place to live, nothing in your life works—not your job, your family, your education, or your health.

So, a foundational question is how AI can help and hinder this goal. We know that some aspects of Artificial Intelligence have been around for a long time, and we also know that major advances are fueling the use of AI in finance and housing in ways that we need to understand.

Consumers can find AI when they encounter chat bots when they shop online, or digital “helpers” that seem to be ubiquitous.

AI plays a role when:

- a prospective tenant is looking to rent an apartment,
- a renter submits a maintenance request to her management company,
- a family tries to qualify for a home loan, or when
- a person experiencing homelessness is connected to services.

These are powerful tools that hold great potential to cut costs, target services, reduce wait times, and even reduce bias.

But they also have the potential to bake in existing inequities, reduce accountability, and limit opportunity.

Today AI is being actively used in every part of the housing continuum, from emergency homelessness services to mortgage financing. As I was preparing for this hearing, I found endless applications.

AI is being deployed to help connect people experiencing homelessness to health and housing resources.

AI is helping to forecast more precisely and accurately where families are at risk of eviction to help better target assistance.

Academics and advocates are using AI and machine learning to help understand and map the country's zoning laws and codes, which span about 30,000 localities. These insights will help understand the dense and complicated rules that govern where, how and what type of housing is allowed to be built, so we can make better decisions about boosting housing supply and lowering costs.

So there are many opportunities. And, there are also very real concerns about the threats that artificial intelligence can pose to Americans. In Minnesota, some landlords are reportedly using AI-generated tenant screening reports that include incorrect and sometimes illegal and off-limits information. The result—it's even harder for people to find a place to rent, and they may never know why they were declined, or be able to correct the record. For landlords, it may be easier to just move onto the next applicant rather than considering additional information.

Another example of how AI, used in a bad way, can be quite harmful: There's a current lawsuit in Minnesota against a law firm that allegedly has “automated” the process of filing evictions for landlords. In one month, the firm filed 400 eviction complaints. These eviction filings lacked much detail about why the eviction was happening and seem to routinely lack basic information about lease terms and included significant errors regarding lease dates, rental amounts, and payment information.

The fact that a firm allegedly leaned on AI to generate a large number of eviction filings with false information, apparently without meaningful reviewing by an attorney is a big problem. Not only is the eviction illegal, but that eviction will live on in public records and hurt the tenant into the future.

AI is also increasingly part of how people buy homes. It's used in credit scoring models and automated valuation models (or AVMs), which determine the value of a home. How AI is deployed has major implications for a person's credit scores, their mortgage rates, and whether home ownership and wealth building is even within

reach. We know that we have historic, systemic challenges with fairness and equity in in this country—my own home town of Minneapolis has some of the greatest disparities in home ownership between Black and White families anywhere in the country. We need to carefully explore whether AI is extending and reinforcing these biases, and how it has the potential to correct them.

Our excellent witnesses have an unenviable task in your opening statements—to ground us in both the opportunities and threats of AI in housing, in 5 minutes each. I look forward to hearing from you, and I look forward to hearing the questions and conversation with my colleagues that will follow.

As with any innovation, there are both opportunities and challenges that must be balanced, and our job is to think through how these complex issues so we can develop the best public policy. I very much look forward to this conversation.

PREPARED STATEMENT OF LISA RICE
 PRESIDENT AND CHIEF EXECUTIVE OFFICER, NATIONAL FAIR HOUSING ALLIANCE
 JANUARY 31, 2024

Introduction

Chairwoman Smith, Ranking Member Lummis, and other distinguished members of the Senate Subcommittee on Housing, Transportation, and Community Development, thank you for the opportunity to testify during the Subcommittee's hearing on Artificial Intelligence and Housing: Exploring Promise and Peril. Artificial Intelligence holds great promise for improving systems, democratizing opportunities, lowering costs, and increasing productivity. Yet, it also holds great dangers for perpetuating bias, spreading mis-information, excluding people from necessary services, and generating other harms. It is critical that Congress understand these dichotomies and establishes sound rules and guardrails that can help ensure the U.S. remains the world leader in innovation and technological advancement and that the nation protects its residents against the perils AI can present. NFHA welcomes the Subcommittee's commitment to advancing on both these fronts.

The National Fair Housing Alliance® (NFHA™) is the country's only national civil rights organization dedicated solely to eliminating all forms of housing and lending discrimination and ensuring equal opportunities for all people. As the trade association for over 170 fair housing and justice-centered organizations throughout the U.S. and its territories, NFHA works to dismantle longstanding barriers to equity and build resilient, inclusive, well-resourced communities where everyone can thrive.

NFHA's Evolution in Addressing Algorithmic and AI Bias & Description of NFHA's Responsible AI Program

NFHA has addressed harms associated with AI and automated systems since its inception in 1988. We first concentrated our efforts on prohibiting or restricting the use of discriminatory automated systems such as credit and insurance scoring, underwriting, and pricing models, in housing and financial services. Early settlements with entities like Prudential, State Farm, Nationwide, and Allstate addressed these discriminatory systems. Several years ago, while litigating a major case against then-Facebook, it became even more clear that technology, including AI, was the new civil and human rights frontier and, as a civil rights organization, we had to be a leader in this sector. Thus, we established our Responsible AI division with an initial focus on Tech Equity. The division is comprised of researchers and engineers committed to civil and human rights principles and is headed by one of the world's premier AI Research Scientists, Dr. Michael Akinwumi. NFHA's Responsible AI division has five workstreams founded on each of the following technical and policy research pillars:

- **Tech Equity:** We focus on developing and advocating for methodologies that ensure automated systems offer equitable access to housing opportunities.
- **Data Privacy:** We strive to test and promote technologies that balance consumer privacy with the need for data access to eliminate bias in automated systems.
- **Explainability:** We advocate for consumers' right to explanations for automated decisions and work to test and promote methodologies that clarify the reasoning or design behind automated systems.
- **Reliability:** We focus on testing and advancing techniques to ensure only safe and valid automated systems are used in housing applications.
- **Human-Centered Alternative Systems:** We work on advancing technical and policy solutions to determine when human-centered alternatives should take precedence over automated

systems in housing decisions, particularly when data quality is poor, infrastructure is inadequate, or there is a lack of social awareness about harms of automated systems.

Since launching our Responsible AI work, NFHA has contributed to, advocated for, and created technical and policy solutions that advance responsible use of technologies in housing, risk management frameworks, and developed a state-of-the-art framework for auditing algorithmic systems,¹ and other policies, including the White House's AI Bill of Rights² and White House Executive Order on Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence.³

The Genesis of Artificial Intelligence

The term "artificial intelligence" was first used by Professor John McCarthy, a mathematics professor at Dartmouth, who convened a summer research workshop held at Dartmouth College in 1956.⁴ Professor McCarthy and three other colleagues penned "A Proposal For The Dartmouth Summer Research Project On Artificial Intelligence," that stated the purpose of the convening was to "proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it."⁵

Just two years after this pivotal conference, Fair, Isaac and Company developed its Credit Application Scoring Algorithm, a rule-based decision management system that simulated human functions and intelligence.⁶ Since then, developers have created many different complex rule-based, statistical, computational models [including expert systems as an attempt to abstract and replicate human intelligence with the objectives] to streamline and standardize decisioning, improve efficiencies, and save costs for a host of transactions and activities in the housing and financial services sector.

¹ Michael Akinwumi, Lisa Rice, and Snigdha Sharma, *Purpose Process, and Monitoring: A New Framework for Auditing Algorithmic Bias in Housing and Lending*, National Fair Housing Alliance (2022), https://nationalfairhousing.org/wp-content/uploads/2022/02/PPM_Framework_02_17_2022.pdf.

² See White House Office of Science and Technology Policy, *Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People* (Oct. 2022), <https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf>.

³ See White House Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (Oct. 30, 2023), <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>.

⁴ See Dartmouth College, *Artificial Intelligence Coined at Dartmouth*, <https://home.dartmouth.edu/about/artificial-intelligence-at-coined-dartmouth>.

⁵ J. McCarthy, M. L. Minsky, N. Rochester, and C. E. Shannon, *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*, Dartmouth College (Aug. 31, 1955), <http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf>.

⁶ See Fair, Isaac and Company Incorporated, *FICO Investor Relations*, Annual Report (Dec. 1998), <https://investors.fico.com/static-files/461a5d42-97be-45d8-a837-38cc25a0c267>.

The Definition of AI

AI focuses on creating machines capable of intelligent behavior. It involves the computational understanding and creation of artifacts that exhibit intelligent behavior.⁷ The scope of AI is broad, encompassing various aspects such as machine intelligence, cognitive functions, or intelligent agents. It also includes systems that mimic human behavior. To guarantee existing automated systems in the housing and lending sectors are not overlooked in the pursuit of implementing AI that is secure, reliable, and free from discrimination, NFHA characterizes AI as a computerized mechanism capable of performing one or more of the following functions:

- Discerning patterns in data;
- Conducting exploratory, predictive, prescriptive, or diagnostic analyses based on data, logical reasoning, or established rules; or
- Generating patterns utilizing data, logic, or rules.

NFHA views a system as an amalgamation of algorithm, model tech infrastructure, and human elements. Fundamentally, for NFHA, AI represents a sociotechnical system, integrating technical capabilities with societal and human components. This practical definition is consistent with how federal agencies characterized AI in a 2023 joint statement on enforcement efforts against algorithmic bias and discrimination.⁸

Use of AI in Housing and Financial Services and Potential Challenges

Tenant Screening Selection Systems

To evaluate applicants, landlords sometimes purchase tenant screening services and tools from companies that offer products that “adjudicate” or “score” a rental applicant based on AI or other automated tools. Tenant screening selection systems have been a focal point of federal regulators and housing advocates for several years. The Consumer Financial Protection Bureau’s (CFPB) November 2022 *Tenant Background Checks Market Report and Consumer Snapshot: Tenant Background Checks* effectively laid out the landscape of the tenant screening market.⁹ In May 2023, NFHA submitted extensive comments in response to the Federal Trade Commission’s (FTC) and CFPB’s Request for Information on Tenant Screening.¹⁰ In those comments, NFHA cited research supporting its assertion that tenant screening practices that screen for criminal history, credit history, and eviction history have a clear and outsized negative impact on historically underserved populations, including people of color, immigrants, public housing voucher recipients, and renters with disabilities. Additionally, we noted that landlords

⁷ S. Shapiro, *Artificial intelligence (AI)* (Jan. 2003), <https://dl.acm.org/doi/abs/10.5555/1074100.1074138>.

⁸ Consumer Financial Protection Bureau, U.S. Department of Justice, Equal Employment Opportunity Commission, and Federal Trade Commission, *Joint Statement on Enforcement Efforts against Discrimination and Bias in Automated Systems* (2023) <https://www.eeoc.gov/joint-statement-enforcement-efforts-against-discrimination-and-bias-automated-systems>.

⁹ CFPB, *Tenant Background Checks Market Report and Consumer Snapshot: Tenant Background Checks* (Nov. 2022), https://files.consumerfinance.gov/f/documents/cfpb_tenant-background-checks-market-report_2022-11.pdf.

¹⁰ NFHA, Comment to the FTC regarding the Request for Information on Tenant Screening (May 30, 2023), <https://www.regulations.gov/comment/FTC-2023-0024-0585>.

are increasingly relying on algorithmic models that may perpetuate bias and make it harder for consumers to understand why they are being denied a housing opportunity.¹¹ As it relates the tenant screening systems, fair housing concerns are magnified because of potentially erroneous data that can be used to develop the models and because of the lack of transparency regarding tenant screening algorithms' predictiveness, design, development, testing process, and data inputs. There are also significant concerns about the quality of risk management efforts pertaining to these systems. Concerns about these systems must be addressed as they will likely impact, in one form or another, nearly all of the 100 million+ renters in the U.S. The tenant screening industry receives billions of dollars in revenue each year despite being plagued by numerous complaints of privacy violations and bias related to these systems.

Dynamic Rental Pricing Systems

In the rental housing market, dynamic pricing systems based on AI have become a widespread feature. Rent prices can change weekly, daily, or more frequently, based on any number of factors that are hidden from consumers making it profoundly difficult for prospective tenants to understand why they are being charged a certain rate. Generally, the factors that go into establishing rent pricing are kept from public view. However, ProPublica has reported that the company that produces the leading rental pricing software (RealPage) uses its clients' leasing data in pricing formation, which is believed to effectively cause rents to increase.¹² Following this revelation, the U.S. Government Accountability Office reported that rental rates had increased 24% in the last three years¹³ and the U.S. Department of Justice (DOJ) filed a Statement of Interest in the anti-trust litigation against RealPage.¹⁴

Moreover, dynamic rental pricing powered by AI can be a barrier to housing for rental voucher recipients as rent prices potentially fluctuate above HUD fair market rent amounts. This can negatively impact low wealth groups who are disproportionately single, female-headed families with children and people with disabilities¹⁵ as well as people who live in rural communities.¹⁶ We need to fully understand the role dynamic pricing systems play in impacting rental housing

¹¹ Valerie Schneider, *Locked Out by Big Data: How Big Data, Algorithms and Machine Learning May Undermine Housing Justice*, 52 Colum. Hum. Rts. L. Rev. 251, 279 (2020), <https://hrtr.law.columbia.edu/hrtr/locked-out-by-big-data-how-big-data-algorithms-and-machine-learning-may-undermine-housing-justice/>.

¹² Heather Vogell, *Rent Going Up? One Company's Algorithm Could Be Why*, ProPublica (Oct. 15, 2022), <https://www.propublica.org/article/realpage-rent-increase-realpage-rent>.

¹³ U.S. Government Accountability Office, *The Affordable Housing Crisis Grows While Efforts to Increase Supply Fall Short* (Oct. 2023), <https://www.gao.gov/blog/affordable-housing-crisis-grows-while-efforts-increase-supply-fall-short>.

¹⁴ *In re Real Page: Rental Software Antitrust Litigation (II)*, Case No. 3:32-MD-3071, DOJ Statement of Interest (Nov. 15, 2023), <https://www.justice.gov/d9/2023-11/418053.pdf>.

¹⁵ See Claudia D. Solari et al., *Housing Insecurity in the District of Columbia*, Urban Institute (Nov. 16, 2023), <https://www.urban.org/research/publication/housing-insecurity-district-columbia>; Sharon Cornelissen, *The Pandemic Aggravated Racial Inequalities in Housing Insecurities: What Can It Teach Us about Housing Amidst Crisis?*, Harvard Joint Center for Housing Studies (July 12, 2023), <https://www.jchs.harvard.edu/blog/pandemic-aggravated-racial-inequalities-housing-insecurity-what-can-it-teach-us-about-housing>.

¹⁶ See Irina Ivanova, *Inflation is Hurting Rural Americans More Than City Folks - Here's Why*, CBS MoneyWatch (Dec. 2, 2021), <https://www.cbsnews.com/news/inflation-rural-households-non-college-grads-hardest/>.

affordability particularly when it impacts members of protected classes under the Fair Housing Act.

Credit Scoring Systems

Credit scoring systems are algorithmic models that attempt to predict a borrower's risk and how well that person is likely to repay their debt obligations. These systems typically generate a numerical score used to help players in the financial services system determine the creditworthiness of a consumer. They are often used as part of a lender's decisions on underwriting and pricing.

Concerns about the potential for credit scoring models to perpetuate bias against certain groups have been raised over the decades. For the most part, the data used to build credit scoring models is generated from data housed at the three main credit reporting agencies (CRAs): Equifax, Transunion, and Experian. Yet the information reported to the CRAs, particularly historical data, can often reflect bias within the credit and housing markets. In other words, when a person is discriminated against in their efforts to access the credit markets, the negative information generated from those discriminatory transactions will be reflected in the data reported to the CRAs and that tainted data can harm consumers. As Federal Reserve Vice Chair of Supervision Michael Barr stated, "Artificial Intelligence...relies on the data that is out there in the world and the data...is flawed. Some of it is just wrong. Some of it is deeply biased...Information we have on the Internet is imperfect...if you train a Machine Learning device, if you train a Large Language Model on imperfect data, you're going to get imperfect results."¹⁷

Moreover, because of the U.S. dual credit market,¹⁸ mainstream lenders and banks are hyper-concentrated in predominately White communities. Conversely, non-traditional lenders like subprime lenders, payday lenders and check cashers, are concentrated in communities of color. This dynamic is not a function of economics. In fact, banks are closing their branches at a higher rate in affluent, high-income Black neighborhoods than they are in low-income non-Black communities.¹⁹ This means borrowers of color, who disproportionately access credit with non-traditional lenders that often do not report positive financial data to the CRAs, will be negatively impacted because the information that demonstrates their ability to repay their financial obligations is not included in the dataset used to build credit scoring models. It also means that that people of color are disproportionately credit invisible or unscorable. Roughly one-third of Black and Latino people do not have credit scores because they disproportionately access credit outside of the financial mainstream.²⁰

¹⁷ See Federal Reserve Board of Governors Vice Chair Michael Barr, *Setting the Foundation for Effective Governance and Oversight: A Conversation with U.S. Regulators*, Responsible AI Symposium (January 19, 2024), https://www.youtube.com/watch?v=HbM_zD0esDo

¹⁸ For more information about the U.S. dual credit market, see the National Fair Housing Alliance's webpage on Access to Credit, available at <https://nationalfairhousing.org/issue/access-to-credit/>.

¹⁹ Zach Fox, Zain Tariq, Liz Thomas, Ciaralou Palicpic, *Bank Branch Closures Take Greatest Toll on Majority-Black Areas*, Standard and Poor Global (July 25, 2019), <https://www.spglobal.com/marketintelligence/en/news-insights/latest-news-headlines/bank-branch-closures-take-greatest-toll-on-majority-black-areas-52872925#:~:text=Since%202010%2C%20the%20branch%20footprint,disparity%20in%20net%20closure%20rates>.

²⁰ CFPB, *Who Are The Credit Invisibles?*, (Dec. 2016), https://files.consumerfinance.gov/f/documents/201612_cfbp_credit_invisible_policy_report.pdf.

According to the Federal Reserve, credit scoring models often include the following considerations:

- Number and type of credit accounts
- Timely payment of bills
- Amount/percentage of available credit
- Collection actions
- Amount of outstanding debt
- Age of accounts

Each of the above components present systemic challenges for underserved consumers like those living in urban and rural areas as those communities can often be credit deserts.²¹ Additionally, consumers impacted by historical discrimination are negatively impacted by these factors because they have less access to mainstream credit. For example, people of color disproportionately were steered to predatory subprime loans that carried abusive terms and conditions. As a result, these consumers disproportionately experienced collection actions not because these consumers were not creditworthy, but rather, because these consumers were targeted by discriminatory practices.²²

There are racial barriers to economic opportunities and these barriers are more prominent for people of color, particularly Black people, Native Americans, and Latinos.²³ These barriers are further magnified by the U.S.'s discriminatory and broken credit scoring system²⁴ which is based on noisy data that underrepresents people of color.²⁵

Insurance Scoring, Underwriting, Rating, and Claims Systems

Insurance scoring systems are algorithmic models built using actuarial data designed to predict the likelihood of a consumer experiencing a risk-related event such as filing a homeowners insurance claim or filing an auto claim. Insurance scoring systems can be built using data from the CRAs; however, model developers also rely on proprietary company data and information purchased from data providers that reflect consumers' claims experiences, weather-related data, and other information. Unlike credit scoring systems, insurance scoring models are not typically designed to determine whether a consumer will pay their insurance premium, but

²¹ Trulia and NFHA, *50 Years After the Fair Housing Act – Inequality Lingers* (April 19, 2018), <https://www.trulia.com/research/50-years-fair-housing/>.

²² Lisa Rice and Deidre Swesnik, *Discriminatory Effects of Credit Scoring on Communities of Color*, Suffolk University Press (June 2012), <https://nationalfairhousing.org/wp-content/uploads/2021/12/NFHA-credit-scoring-paper-for-Suffolk-NCLC-symposium-submitted-to-Suffolk-Law.pdf>.

²³ Will Douglas Heaven, *Bias Isn't the Only Problem with Credit Scores—and No, AI Can't Help* (June 17, 2021), 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.unp.aid.engagement&utm_content=06.26.21.subs&mc_cid=aad1663503&mc_eid=eead1c58a0

²⁴ Megan Leonhardt, *Democrats and Republicans in Congress Agree: The System That Determines Credit Scores Is "Broken"* (Feb. 27, 2019) <https://www.cnbc.com/2019/02/27/american-consumer-credit-rating-system-is-broken.html>.

²⁵ Laura Blattner and Scott Nelson, *How Costly Is Noise? Data and Disparities in Consumer Credit* (May 5, 2021) <https://arxiv.org/pdf/2105.07554.pdf>.

rather, the consumer's risk profile related to an event that would present a monetary exposure for the insurance company. The score can be used to help underwrite business and/or determine the price a consumer should pay for insurance.

Insurance scoring systems, like other AI systems can perpetuate bias in myriad ways.²⁶ NFHA has challenged discriminatory provisions in insurance scoring systems, most specifically in a matter brought against the Prudential Insurance Company.²⁷ NFHA's investigation, research, and expert witness analysis revealed the company's insurance scoring system presented a disproportionate discriminatory impact on Black consumers in the price they paid for homeowners insurance. The differences in premiums between Black and White insureds was not explainable, based on NFHA's analysis, by appropriate risk factors. Instead, we alleged the model the company utilized was contributing to disparate outcomes. Moreover, our expert was able to devise a way to yield a less discriminatory outcome than the one generated by the model used by Prudential.

Additionally, the Casualty Actuarial Society (CAS) acknowledged that algorithmic bias can manifest in systems used in the insurance sector including underwriting, pricing, and claims models.²⁸ The CAS issued reports highlighting examples of discrimination in the insurance market over the decades and various analyses of whether certain variables used to assess risk or affix insurance rates presented unfair discrimination. Factors such as zip code, address, educational level, credit scores, and occupation have raised serious concerns about their propensity for manifesting discrimination. CAS' research pointed out that "fully or semi-automated systems...are...inherently capable of introducing unfairness into the process and thus have direct consequences to individuals affected by these models. Bias is all around us, and it can creep into the decision-making paradigm in subtle ways, whether it is the subjectivity of human judgement, prejudice, historical inequities baked into the data, or faulty algorithms."²⁹

Automated Underwriting Systems and Risk-Based Pricing Systems

Automated underwriting systems and risk-based pricing systems manifest and perpetuate bias as well. These systems rely on and are built using data contained in the CRAs. The data captured by CRAs is under-representative as it is missing critical information, like rental housing payment data, that can accurately reflect a borrower's willingness and ability to pay financial obligations. Data captured by the CRAs also includes information tainted by bias against underserved groups. Unfortunately, redlining and housing discrimination are still everyday

²⁶ AI systems can manifest discrimination by using biased or non-representative data sets, and by other means. For a detailed discussion about the many ways AI can perpetuate discrimination, see Testimony of Lisa Rice, *Hearing on Equitable Algorithms: How Human-Centered AI Can Address Systemic Racism and Racial Justice in Housing and Financial Services* before the U.S. House Financial Services Task Force on Artificial Intelligence (May 7, 2021), <https://nationalfairhousing.org/wp-content/uploads/2022/01/Lisa-Rice-House-Testimony-on-AI-5-7-21.pdf>

²⁷ See *National Fair Housing Alli. v. Prudential Ins. Co.*, 208 F. Supp. 2d 46 (D.D.C. 2002), <https://casetext.com/case/national-fair-housing-alli-v-prudential-ins-co>

²⁸ Ronda Lee, *AI Can Perpetuate Racial Bias in Insurance Underwriting*, Yahoo!Money (Nov. 1, 2022).

²⁹ Roosevelt Mosley, FCAS, and Radost Wenman, FCAS, *CAS Research Paper Series on Race and Insurance Pricing: Methods for Quantifying Discriminatory Effects on Protected Classes in Insurance*, Casualty Actuarial Society (2022), https://www.casact.org/sites/default/files/2022-03/Research-Paper_Methods-for-Quantifying-Discriminatory-Effects.pdf?utm_source=Landing&utm_medium=Website&utm_campaign=RIP+Series.

occurrences³⁰ and when consumers experience discrimination, that bias is reflected in the data captured by the CRAs.

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 University of California-Berkeley 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.”³¹ The study revealed that Black and Latino borrowers were overcharged by \$765 million per year. That is, Black and Latino borrowers were disproportionately charged a rate that is higher than their commensurate level of risk because of biased risk-based pricing systems.

Automated Valuation Models

The use of Automated Valuation Models (AVMs) for some portion of the home valuation process has been proliferating, especially as a potential remedy to multiple complaints of consumers of color having to “whitewash” their home and remove all indications of their race and ethnicity to receive a fair value.³² By statute, an AVM is defined as a “computerized model used by mortgage originators and secondary market issuers to determine the collateral worth of a mortgage secured by a consumer’s principal dwelling.”³³

An AVM may use AI or some other automated system to serve various purposes. In some cases, a secondary market issuer or lender may use the AVM in place of a human appraiser; in some cases, the human appraiser may develop the opinion of value using an AVM; and in still other cases, a secondary market issuer or lender may use an AVM as a check on the human appraiser. The AVM may be faster, cheaper, more consistent, and less prone to bias, but AVMs pose at least two problems.

First, the AVM still relies on the traditional sales comparison approach in which the home valuation is based on recent sales of comparable homes in a comparable neighborhood. Because of America’s long history of redlining and segregation, the AVM is likely using historical data that undervalues homes in formerly redlined areas.³⁴ For example, research has shown that

³⁰ See, *Recent Accomplishments of the Housing and Civil Enforcement Section*, U.S. Department of Justice (Oct. 5, 2023), <https://www.justice.gov/crt/recent-accomplishments-housing-and-civil-enforcement-section>; *What Modern-Day Housing Discrimination Looks Like: A Conversation with the National Fair Housing Alliance* (Feb. 4, 2019), <https://www.zillow.com/research/modern-housing-discrimination-22898/>; *Fair Housing Solutions: Overcoming Real Estate Sales Discrimination*, National Fair Housing Alliance (Dec. 2019).

³¹ 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), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3063448.

³² See, e.g., Julian Glover and Mark Nichols, *Our America: Lowballed*, ABC (2022), <https://abc7news.com/feature/our-america-lowball-home-appraisal-racial-bias-discrimination/12325606/>.

³³ 12 U.S.C. § 3354(d).

³⁴ See, NFHA and National Consumer Law Center, Comment to FHFA, CFPB, Federal Reserve, FDIC, OCC, and NCUA regarding Quality Control Standards for Automated Valuation Models (Aug. 21, 2023), <https://www.fdic.gov/resources/regulations/federal-register-publications/2023/2023-quality-control-standards-for-automated-valuation-models-3064-ae68-c-010.pdf>.

in 2021, homes in White neighborhoods were appraised at values nearly 250% higher than similar homes in similar Black neighborhoods and nearly 278% higher than similar homes in similar Latino neighborhoods.³⁵ These disparities then generate data points used by the AVM, which can perpetuate the disparity in future transactions.

Second, using AVMs may create a bifurcated valuation system. AVMs tend to work best in neighborhoods with similar homes that generate multiple data points, which means the models excel in suburban subdivisions, which tend to be majority-White. AVMs do not work as well in neighborhoods of color, which tend to have older properties with varied repairs and upgrades. Also, consumers of color tend to stay in their homes longer, which means fewer sales data points for the AVM. Fannie Mae and Freddie Mac have begun accepting the AVM value and waiving the traditional appraisal in certain situations.³⁶ The risk is that a bifurcated valuation system will develop in which consumers of color are more likely to be burdened with the cost of a traditional appraisal, which in turn tends to undervalue their home.

Marketing Systems and Digital Redlining

The use of AI and other algorithms have changed the face of the billion-dollar U.S. advertising industry from a market where ad content was posted predominantly in general media as compared to the hyper-targeted, individualized ad placement that we currently navigate in the new tech frontier. This transition has brought new challenges for bias in marketing practices, ad delivery, and services.

Fair housing groups challenged Facebook's digital advertising systems, operated by algorithms, that manifest discrimination against protected groups. As a result of the litigation, Facebook changed its ad-platform to restrict the ways that housing, employment, and credit ads can be targeted to potential consumers.³⁷ Then, the U.S. Department of Justice entered into a consent decree with Meta (formerly Facebook), whereby the digital ad platform agreement to implement a variance reduction auditing system intended to reduce racial bias in the delivery of ad campaigns to Meta-users.³⁸

Algorithms may also impact how consumers can access housing products and services. Recently, NFHA pursued a federal lawsuit and entered into a settlement with Redfin Corporation for alleged "digital redlining," which is essentially the discriminatory offering of their services in

³⁵ Junia Howell and Elizabeth Korver-Glenn, *Appraised: The Persistent Evaluation of White Neighborhood as More Valuable Than Communities of Color*, Eruka (2022)

https://static1.squarespace.com/static/62e84d924d2d8e5dff96ae2f/t/6364707034ee737d19dc76da/1667526772835/Howell+and+Korver+Glenn+Appraised_11_03_22.pdf.

³⁶ See, e.g., Fannie Mae, *Delivering Effective, Efficient, and Impartial Home Valuations Across America*, <https://singlefamily.fanniemae.com/valuation-modernization>.

³⁷ Tracy Jan and Elizabeth Dwoskin, *Facebook Agrees to Overhaul Targeted Advertising System for Job, Housing, and Loan Ads After Discrimination Complaints*, The Washington Post (March 19, 2019), https://www.washingtonpost.com/business/economy/facebook-agrees-to-dismantle-targeted-advertising-system-for-job-housing-and-loan-ads-after-discrimination-complaints/2019/03/19/7dc9b5fa-4983-11e9-b79a-961983b7e0d_story.html

³⁸ See, DOJ Press Release, *Justice Department and Meta Platforms, Inc. Reach Key Agreement as They Implement Groundbreaking Resolution to Address Discriminatory Delivery of Housing Advertisements* (Jan. 9, 2023), <https://www.justice.gov/opa/pr/justice-department-and-meta-platforms-inc-reach-key-agreement-they-implement-groundbreaking>.

non-White areas.³⁹ In NFHA's investigation, fair housing groups found that Redfin offered "No Service" for homes in non-White areas at a greater rate than for homes in White areas. Under a settlement agreement, Redfin agreed to alter its minimum home price policy to be more inclusive in its servicing.⁴⁰

As AI continues to shape changes in the advertising markets and in how consumers broadly access housing and lending services, it is imperative that players across the tech field consider the potential for discriminatory outcomes of their systems and implement ongoing monitoring audits of these platforms.

Use of Large Language Models in Housing and Financial Services and Potential Challenges

Large Language Models (LLMs) are a significant advancement in natural language processing and AI. These models are characterized by their ability to process, understand, and generate human-like text. Some of their capabilities are:

1. Capabilities in Language Generation: LLMs offer unparalleled capabilities in language generation, which opens exciting opportunities for interaction design. They are highly context-dependent and can adapt to a wide range of linguistic styles and nuances.⁴¹
2. Importance in Psycholinguistics: While LLMs are not precise models of human linguistic processing, their success in language modeling makes them significant in the field of psycholinguistics. They serve as practical tools for exploring language and thought relationships.⁴²
3. Enhancing Creativity: AI LLMs have contributed to creative writing, including newspaper articles, novels, and poetry. These models can generate creative and original text, demonstrating their potential in diverse creative applications.⁴³

In summary, large language models represent a significant leap in AI's ability to interact with and understand human language. They are crucial for a variety of applications, from enhancing creative writing to contributing to the understanding of human language processing, despite facing challenges in certain aspects of language comprehension. However, due to lack of transparency and the fact that LLM applications are yet to mature, it is not immediately clear how lenders are using LLM applications in the product and service lines.

³⁹ See, NFHA, et al. v. Redfin investigation summary, available at <https://nationalfairhousing.org/redfin-investigation/>.

⁴⁰ See, NFHA Press Release, *National Fair Housing Alliance and Redfin Agree to Settlement that Greatly Expands Access to Real Estate Services in Communities of Color* (April 29, 2022), <https://nationalfairhousing.org/national-fair-housing-alliance-and-redfin-agree-to-settlement-which-greatly-expands-access-to-real-estate-services-in-communities-of-color-%EF%BF%BC/>.

⁴¹ Mina Lee, Percy Liang, and Qian Yang, CoAuthor: *Designing a Human-AI Collaborative Writing Dataset for Exploring Language Model Capabilities* (2022), <https://doi.org/10.1145/3491102.3502030>.

⁴² Conor J. Houghton, N. Kazanina, and Priyanka Sukumaran, *Beyond the Limitations of Any Imaginable Mechanism: Large Language Models and Psycholinguistics* (2023), <https://doi.org/10.48550/arXiv.2303.00077>.

⁴³ See *id.*

That said, LLMs can significantly impact the housing and financial services sectors, offering innovative solutions while presenting unique challenges. Here are some key use cases and potential challenges:

1. Use Cases in Housing and Financial Services:

- a. Enhanced Customer Service: LLMs may improve customer service in housing and financial services by providing more accurate and efficient responses to customer inquiries.
- b. Automated Document Analysis: These models can analyze legal documents, loan applications, and other paperwork, speeding up processes and potentially reducing human error.
- c. Predictive Analytics: AI-driven predictive analytics can help in understanding market trends, predicting housing prices, and identifying investment opportunities in the financial sector.
- d. Personalized Financial Advice: LLMs can offer personalized financial advice based on individual customer data, improving the customer experience in financial services.
- e. Risk Assessment and Management: AI can play a crucial role in assessing risk in loan approvals, insurance underwriting, and investment decisions.

2. Potential Challenges:

- a. Data Privacy and Security: Ensuring the confidentiality and security of sensitive personal and financial data is a major challenge, especially considering the vast amount of data processed by AI systems.
- b. Bias and Fairness: There is a risk of AI systems inheriting biases from tainted training data, which can lead to unfair practices in loan approvals, housing allocations, financial advisories, and other instances.
- c. Regulatory Compliance: AI systems must comply with existing financial regulations and housing laws, which can be complex and vary across states and regions.
- d. Integration with Existing Systems: Integrating AI technologies with legacy systems in housing and financial services can be challenging due to compatibility issues.
- e. Reliability and Accuracy: Ensuring the reliability and accuracy of AI-driven decisions and outputs, especially in critical housing and financial situations, is paramount.
- f. Public Trust and Acceptance: Building public trust in AI systems, particularly in sensitive areas like financial advice and housing, is crucial for widespread adoption.
- g. Challenges in Word Understanding: Despite their advanced capabilities, LLMs sometimes struggle with understanding words based on their dictionary definitions, indicating areas for improvement in word comprehension.⁴⁴

While LLMs offer promising opportunities for innovation in housing and financial services, addressing challenges related to data security, bias, regulatory compliance, and public trust will be key to their successful implementation and acceptance.

⁴⁴ Lutfi Kerem Senel and Hinrich Schütze *Does She Wink or Does She Nod? A Challenging Benchmark for Evaluating Word Understanding of Language Models* (2021). <https://doi.org/10.18653/v1/2021.eacl-main.42>.

Using AI to Promote Fairer Outcomes

Debiasing AI Systems Using Fairness Techniques

Algorithmic fairness techniques play a crucial role in mitigating biases and ensuring equitable outcomes in AI systems. These techniques can be categorized into pre-processing, in-processing, and post-processing methods, each addressing biases at different stages of the AI model lifecycle.

1. **Pre-processing Techniques:** These techniques focus on the initial stages of AI development, where the primary goal is to rectify biases present in the training data before it is fed into the model. By identifying and modifying biased data, pre-processing methods aim to prevent the AI system from learning and perpetuating these biases.⁴⁵
2. **In-processing Techniques:** These are applied during the model training phase and involve modifying the learning algorithm to incorporate fairness. This might include adjusting the model's objective function to balance accuracy with fairness criteria. One approach is the integration of fairness constraints directly into the training process, as discussed by some researchers in their exploration of unified data and algorithm fairness through adversarial data augmentation and adaptive model fine-tuning.⁴⁶
3. **Post-processing Techniques:** These methods are applied after the model has been trained. They adjust the output of the model to ensure fairness, often through recalibration of decision thresholds for different groups. This approach is particularly useful in scenarios where modifying the training process is not feasible or when dealing with legacy systems. The work of Berkel et al. on the effect of information presentation on fairness perceptions of machine learning predictors provides insights into how post-processing interventions can alter people's perceptions of AI fairness.⁴⁶

In summary, algorithmic fairness techniques are integral to developing AI systems that are equitable and unbiased. These methods, spanning from the initial data preparation to the final model output, help in creating AI solutions that are not only effective but also fair and just.

Using AI to Detect the Risk of Discriminatory Practices and Policies

AI can significantly assist in the identification and mitigation of the risk of discriminatory practices and policies, particularly in housing and employment. This potential is explored through various stages and tools, including dataset analysis, internet or website crawling, and other enforcement-based AI tools. The following are some AI-driven approaches that support the identification and mitigation of discriminatory practices and policies, providing more equitable outcomes in housing, employment, and other critical areas.

Identifying Datasets that are Non-representative/Under-representative

AI can be used to identify datasets that are non-representative or under-representative of certain populations, which is crucial for uncovering discriminatory practices in housing and

⁴⁵ Nengfeng Zhou, Zach Zhang, Vijay Nair, Harsh Singhal, and Jie Chen, *Bias, Fairness and Accountability with Artificial Intelligence and Machine Learning Algorithms* (2022), <https://doi.org/10.1111/insr.12492>.

⁴⁶ N. V. Berkel, Jorge Gonçalves, D. Russo, S. Hosio, and M. Skov, *Effect of Information Presentation on Fairness Perceptions of Machine Learning Predictors*, (2021), <https://doi.org/10.1145/3411764.3445365>.

other areas. The work by Kamikubo et al. on the representativeness of accessibility datasets highlights the importance of ensuring that datasets are representative and inclusive of various demographic groups to mitigate bias in AI applications.⁴⁷

Internet or Website Crawling Tools

Internet or website crawling tools are instrumental in gathering online data, including housing listings and employment advertisements, which can then be analyzed for patterns of potential discrimination. Internet crawling can extract large amounts of data, enabling AI systems to analyze trends and detect potentially discriminatory language or practices. These tools can allow users to search for discriminatory phrases and comments that can impede access to fair housing and lending opportunities.

Other Enforcement-based AI Tools

AI tools, particularly those leveraging natural language processing (NLP) and large language models, can analyze textual data to identify discriminatory language or patterns in documents, policies, and communications. For example, Lobo's work on bias in hate speech and toxicity detection illustrates how AI can identify discriminatory or biased language in large datasets, which can be applied to analyze housing and employment-related texts.⁴⁸

Expanding Inclusive and Equitable Housing and Financial Services Opportunities

Using AI to Deepen Research

AI can significantly enhance the depth and effectiveness of housing and lending policy, enforcement, and legal research. This enhancement can be achieved through several applications:

1. **Dataset Analysis for Policy and Enforcement:** AI is instrumental in analyzing vast datasets to identify trends and patterns that might indicate discriminatory practices in housing and lending. The analysis can focus on various aspects, such as loan approval rates, interest rates charged, and geographical distribution of loans. These insights are critical for shaping fair housing policies and for enforcing anti-discrimination laws. The work by Wyly and Holloway on the disappearance of race in mortgage lending highlights the importance of data analysis in understanding and addressing issues in housing markets.⁴⁹
2. **Enforcement-based AI Tools Using NLP:** Natural Language Processing (NLP) and large language models can be leveraged to analyze legal documents, policy papers, and communication records for housing and lending practices. These AI tools can identify discriminatory language patterns and assist in the legal interpretation of housing policies and practices. The research by Ntoutsis et al. on bias in data-driven AI systems provides an

⁴⁷ Rie Kamikubo, Lining Wang, Crystal Marte, Amnah Mahmood and Hernisa Kacorri, *Data Representativeness in Accessibility Datasets: A Meta-Analysis*, (2022), <https://doi.org/10.1145/3517428.3544826>.

⁴⁸ Paula Reyer Lobo, *Bias in Hate Speech and Toxicity Detection* (2022), <https://doi.org/10.1145/3514094.3539519>.

⁴⁹ Elvin K. Wyly and S. Holloway, *The Disappearance of Race in Mortgage Lending*, *Economic Geography* (2002), <https://doi.org/10.1111/j.1944-8287.2002.tb00181.x>.

overview of the challenges and solutions in detecting bias, which is applicable to housing and lending legal research.⁵⁰

Leveraging AI to Design Systems with Equity

AI technologies can analyze alternative data sources, such as cash-flow underwriting data and rental housing payment histories, which are particularly beneficial for individuals with limited credit histories. AI and big data can capture weak signals in creditworthiness assessments, thereby improving financial inclusion and access to credit for traditionally underserved borrowers.⁵¹ This approach helps in expanding the scope of data considered for lending decisions, making them more inclusive and aligned with traditional ways lending decisions were assessed using the "5 Cs of credit": character, cash flow, capital, conditions, and collateral.

Using AI to Fix the Non-representative/Under-representative Data Problem

AI can be used to identify and rectify biases in non-representative or under-representative datasets, which is crucial for ensuring equitable housing and lending practices. The work of Jain and Verma highlights the importance of AI in making the credit underwriting process more accurate.⁵² By refining data to be more representative, AI can help lenders make more fair and equitable decisions.

Optimizing Privacy Protections Via AI Techniques

Privacy-enhancing techniques are essential in AI to ensure the protection of consumer data while utilizing it for housing and lending purposes.⁵³ The work of Chen et al. shares lessons learned during the development of frameworks that aid in the correct use of privacy-enhancing technologies like homomorphic encryption and secure multi-party computation.⁵⁴ These technologies are vital for developing AI systems that are both effective and respectful of consumer privacy.

⁵⁰ Eirini Ntoutsis et al., *Bias in Data-Driven AI Systems: An Introductory Survey* (Jan. 14, 2020), <https://arxiv.org/abs/2001.09762>.

⁵¹ Hicham Sadok, Fadi Sakka and Mohammed El Hadi El Maknouzi, *Artificial Intelligence and Bank Credit Analysis: A Review*, *Cogent Economics & Finance* (2022), <https://doi.org/10.1080/23322039.2021.2023262>.

⁵² Aastha Jain and Deval Verma, *Making Credit Underwriting Process More Accurate using ML* (2022), <https://doi.org/10.1109/ICACCM56405.2022.10009117>.

⁵³ Hannah Holloway, Snigdha Sharma, Samantha Gordon, and Dr. Michael Akinwumi, *Privacy, Technology, and Fair Housing – A Case for Corporate and Regulatory Action*" (Aug. 22, 2023), <https://nationalfairhousing.org/privacy-technology-and-fair-housing-a-case-for-corporate-and-regulatory-action/>.

⁵⁴ Huiii Chen, S. Hussain, Fabian Boemer, Emmanuel Stapf, A. Sadeghi, F. Koushanfar and Rosario Cammarota, *Developing Privacy-preserving AI Systems: The Lessons Learned*, 1-4 (2020), <https://doi.org/10.1109/DAC18072.2020.9218662>.

Using AI to Fix Zoning Challenges

Exclusive and restrictive zoning policies have been used over the decades to generate residential segregation, reduce affordable housing options, and thwart fair housing efforts.⁵⁵ These policies contribute to the U.S. racial wealth and homeownership gaps as well as to the affordable housing crisis.⁵⁶ It is estimated that the U.S. has a shortage of between 4 million to 7 million affordable housing units.⁵⁷ The lack of fair and affordable housing opportunities is critical since housing is the number one driver of inflation.⁵⁸

Innovators are using AI and other technologies to address zoning challenges. For example, AI can help analyze zoning regulations to identify barriers and allow developers to plan projects more efficiently. Some scholars and urban planners are examining the idea of using AI to automate the process of developing zoning codes.⁵⁹ Using AI to develop and analyze zoning codes may make it easier and quicker to identify provisions in the ordinances that present barriers to fair housing and affordable housing. This would, in turn, enable legislators to address potential discrimination risks. Using AI in this way can also help jurisdictions meet their obligation to Affirmatively Further Fair Housing.⁶⁰

Applying Fair Housing and Fair Lending Laws to AI

There are two key laws that prohibit discrimination in lending and housing: the Fair Housing Act and the Equal Credit Opportunity Act ("ECOA"). The Fair Housing Act prohibits any entity from discriminating in housing and mortgage lending on the basis of race, color, religion, national

⁵⁵ Fair Share Housing Center, *Dismantling Exclusionary Zoning: New Jersey's Blueprint for Overcoming Segregation*, (April 2023), <https://www.fairsharehousing.org/wp-content/uploads/2023/04/Dismantling-Exclusionary-Zoning-New-Jerseys-Blueprint-for-Overcoming-Segregation.pdf>

⁵⁶ Allison Hanley, *Rethinking Zoning to Increase Affordable Housing*, Journal of Housing & Community Development, National Association of Housing and Redevelopment Officials (Dec. 22, 2023), https://www.nahro.org/journal_article/rethinking-zoning-to-increase-affordable-housing/#:~:text=In%202021%2C%20home%20prices%20experienced%20%20percent%2C%20and%20rents%20surged.&text=Restrictive%20zoning%20practices%20significantly%20contribute%20restricting%20or%20banning%20apartments.

⁵⁷ Sam Khater, *One of the Most Important Challenges our Industry will Face: The Significant Shortage of Starter Homes*, Freddie Mac (April 15, 2021), <https://www.freddiemac.com/perspectives/sam-khater/20210415-single-family-shortage>; National Association of Realtors®, *Once-In-A-Generation Response Needed to Address Housing Supply Crisis* (June 16, 2021), <https://www.nar.realtor/newsroom/once-in-a-generation-response-needed-to-address-housing-supply-crisis>; National Low Income Housing Coalition, *The Gap: A Shortage of Affordable Rental Homes* (March 2023), <https://nlhc.org/gap>

⁵⁸ Katy O'Donnell, *The Main Driver of Inflation Isn't What You Think It Is*, Politico (March 18, 2022), <https://www.politico.com/news/2022/03/18/housing-costs-inflation-00015808>

⁵⁹ Norman Wright, AICP, *Using Generative AI to Draft Zoning Codes*, American Planning Association (Oct. 2023), <https://www.planning.org/publications/document/9277441/>

⁶⁰ Jurisdictions receiving federal funds for a housing or community development purpose must ensure all their laws, programs, and services are implemented in compliance with the Federal Fair Housing Act and in a manner that Affirmatively Furthers Fair Housing. This means zoning policies must not perpetuate discrimination, segregation, or other anti-fair housing principles. See National Fair Housing Alliance, *Affirmatively Furthering Fair Housing*, <https://nationalfairhousing.org/issue/affirmatively-furthering-fair-housing/>

origin, sex (including sexual orientation), disability, and familial status (also known as “protected characteristics” or “protected classes” or “prohibited bases”).⁶¹ The Fair Housing Act also requires entities receiving federal funds for a housing or community development purpose to disseminate those funds, as well as implement their programs and services in a way that Affirmatively Furthers Fair Housing. The ECOA prohibits “creditors” from discriminating in lending on the basis of race, color, religion, national origin, sex (including sexual orientation), marital status, age, and source of income.⁶²

Generally, there are two methods of proving discrimination under either the Fair Housing Act or ECOA: “disparate treatment” or “disparate impact.”⁶³ Disparate treatment occurs when an entity explicitly, overtly, or intentionally treats people differently based on prohibited characteristics, such as race, national origin, or sex. Disparate treatment can be proven through direct evidence or indirect (or “circumstantial”) evidence, for example comparator evidence, statistical evidence, or a pretextual explanation. Although disparate treatment is known as “intentional discrimination,” the law does not actually require showing prejudice, animus, or even an intent to treat someone worse because of a protected class; the differential treatment is enough to establish a violation of law.⁶⁴ An example of potential disparate treatment discrimination would be an AI model that explicitly included a protected classes (such as race) as a model variable, or that resulted in different, adverse outcomes on a prohibited basis (such as race) for similarly-situated individuals.

Disparate impact discrimination occurs when a (1) facially neutral policy or practice disproportionately harms members of a protected class, and either (2) the policy or practice does not advance a legitimate interest, or (3) a less discriminatory alternative to serve the legitimate interest exists. Disparities alone are not sufficient to impose disparate impact liability, and entities are not required to sacrifice legitimate business needs or ignore relevant business considerations. Disparate impact only requires entities to avoid considerations that disproportionately harm members of protected classes unnecessarily. An example of potential disparate impact discrimination would be an AI model that considers a mortgage applicant’s arrest record, which is not predictive of default risk.

Institutions should be aware that they need adequate Compliance Management Systems (CMS) to monitor and test AI models for potential discrimination, search for less discriminatory alternatives, and implement appropriate controls for fair lending or fair housing risk.⁶⁵ Courts and agencies have for decades recognized that disparate impact liability exists under laws like the Fair Housing Act, ECOA, and Title VII (prohibiting employment discrimination).⁶⁶ Disparate impact liability has been a part of Regulation B, which implements ECOA, since the 1970s and

⁶¹ The Fair Housing Act: 42 U.S.C. § 3601, et seq.; HUD’s implementing regulation: 24 CFR Part 100.

⁶² ECOA: 15 U.S.C. § 1619(a); CFPB’s Regulation B: 12 CFR Part 1002.

⁶³ Based on legal precedent, the federal financial regulators have also based fair lending risk assessments on these theories of discrimination. See FFIEC, *Interagency Fair Lending Examination Procedures* (2009), <https://www.ffiec.gov/pdf/fairlend.pdf>.

⁶⁴ See 12 C.F.R. Part 1002, 4(a)-1: “Disparate treatment on a prohibited basis is illegal whether or not it results from a conscious intent to discriminate.”

⁶⁵ See, e.g., Relman, Colfax, *Initial Report of the Independent Monitor, Fair Lending Monitorship of Upstart Network’s Lending Model* (April 14, 2021), https://www.relmanlaw.com/media/cases/1086_Upstart%20Initial%20Report%20-%20Final.pdf.

⁶⁶ 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).

was included in the federal agencies' 1994 Policy Statement on Discrimination in Lending.⁶⁷ It has also been sanctioned by the U.S. Supreme Court and 11 appellate courts.⁶⁸ Moreover, federal agencies have long provided guidance informing institutions of their obligations for third party oversight.⁶⁹ Institutions should be aware that using a third-party AI model does not insulate them from liability.

Policy Recommendations

The United States Must Enact Comprehensive Legislation to Advance Responsible AI

While there are significant risks of bias and discrimination in AI systems and their impact on housing and financial services, the risks are not insurmountable. The U.S. must play a leadership role in advancing Responsible AI principles and tech equity. Leading the world on these issues includes passing comprehensive legislation that forms the basis for sound policies, systems, practices, and frameworks for Responsible AI. While much of the world's technological innovations are developed in the U.S., other nations are significantly stepping up their efforts by building the infrastructure needed to spur AI innovations. This includes supporting education and training in the field and implementing rules to govern AI. The U.S. is behind the curve, and in some cases playing catch-up to other nations. The U.S. must lead the world in ensuring technologies are fair and beneficial, do not harm people and communities; and promote ideals of freedom, including ensuring robust privacy protections, equality, and equity. For these reasons, Congress should pass new legislation that mitigates the risk of algorithmic bias and ensures fairer structures by:

Ensuring Strong Civil and Human Rights Protections

Congress should develop AI legislation that reflects civil and human rights principles that are foundational to America's ideals of freedom and equality. These laws also must create equity. Civil rights, human rights, and consumer protection organizations lack the resources to ensure technologies are beneficial and not harmful, which means Congress must increase federal funding and new programs to support effective oversight and accountability. For example, a research partnership could be formed between the National Institute of Standards and Technology (NIST), civil rights organizations, consumer protection groups, non-profit trusted research agencies, and financial institutions that rely on AI to evaluate how AI or machine learning models affect fair housing and lending. Congress could also encourage the National Science Foundation to ensure that a portion of the considerable allocations for research on AI focus on the implications of using AI in housing and financial services.

⁶⁷ DOJ, HUD, FTC, Federal Housing Finance Board, Federal Reserve, FDIC, OCC, Office of Thrift Supervision, and NCUA, *Interagency Policy Statement on Fair Lending* (1994) <https://www.govinfo.gov/content/pkg/FR-1994-04-15/html/94-9214.htm>.

⁶⁸ See *Texas Dep't of Hous. & Cmty. Affairs v. Inclusive Cmty. Project*, 135 S. Ct. 2507 (2015).

⁶⁹ The agencies recently replaced each of their separate policies dating as early as 2008 for joint guidance. See Federal Reserve, FDIC, OCC, *Interagency Guidance on Third Party Relationships: Risk Management*, 88 Fed. Reg. 37920 (June 9, 2023) <https://www.govinfo.gov/content/pkg/FR-2023-06-09/pdf/2023-12340.pdf>.

Ensuring Compliance with Existing Civil Rights and Consumer Protection Laws

Congress should ensure that the federal agencies issue robust policies that remind institutions of their legal obligations to test AI models for potential disparate treatment or disparate impact discrimination.

First, all of the federal agencies with responsibility for supervision and/or enforcement of the Fair Housing Act and/or ECOA (collectively, the Agencies)⁷⁰ should emphasize that discrimination in AI models is illegal, including AI models developed or deployed by third parties. In 2023, the DOJ, FTC, CFPB, and EEOC issued a joint statement regarding enforcement efforts to protect the public from bias in AI and automated systems.⁷¹ The remaining federal agencies should immediately issue a similar announcement.

Second, consistent with the Uniform Interagency Consumer Compliance Rating System⁷² and the Model Risk Management Guidance,⁷³ the Agencies should ensure that financial institutions have appropriate Compliance Management Systems that effectively identify and control risks related to AI models, including the risk of discriminatory outcomes for consumers. Where a financial institution's use of AI indicates weaknesses in their Compliance Management System or violations of law, the Agencies should use all of the tools in their toolbelt to quickly address and prevent consumer harm, including issuing supervisory Matters Requiring Attention; entering into a non-public enforcement action, such as a Memorandum of Understanding; referring a pattern or practice of discrimination to the DOJ; or entering into a public enforcement action. The Agencies have already provided clear guidance (e.g., the Uniform Consumer Compliance Rating System) that financial institutions must appropriately identify, monitor, and address compliance risks, and the Agencies should not hesitate to act within the scope of their authority. When possible, the Agencies should explain to the public the risks that they have observed and the actions taken in order to bolster the public's trust in appropriate oversight, and provide clear examples to guide the industry.

Finally, the Agencies should clarify acceptable methods for AI testing so that institutions can align their methods accordingly. Existing civil rights laws and policies provide a framework for the Agencies to analyze fair lending risk in AI and to engage in supervisory or enforcement actions, where appropriate. That said, the Agencies can be more effective in ensuring consistent and effective compliance by setting clear regulatory expectations regarding testing for the risk of discrimination. The Agencies have been in learning mode for some time. Indeed, the

⁷⁰ The DOJ and HUD have responsibility for enforcement of the Fair Housing Act. The DOJ and FTC have enforcement authority for ECOA. The FHFA, FDIC, Federal Reserve Board, FDIC, OCC, and NCUA have supervision and enforcement authority for certain financial institutions with respect to the Fair Housing Act and ECOA. The CFPB has regulatory, supervision, and enforcement authority for ECOA.

⁷¹ CFPB, DOJ, EEOC, FTC, *Joint Statement on Enforcement Efforts against Discrimination and Bias in Automated Systems* (April 2023), https://www.ftc.gov/system/files/ftc_gov/pdf/EEOC-CRTFTC-CFPB-AI-Joint-Statement%28final%29.pdf.

⁷² Federal Financial Institutions Examination Council, *Uniform Interagency Consumer Compliance Rating System* (Nov. 14, 2016) <https://www.govinfo.gov/content/pkg/FR-2016-11-14/pdf/2016-27226.pdf>.

⁷³ See, e.g., FHFA, *Artificial Intelligence, Machine Learning Model Risk Management*, Advisory Bulletin 2022-02 (Feb. 10, 2022) (explicitly addressing fairness and equity), <https://www.fhfa.gov/SupervisionRegulation/AdvisoryBulletins/Pages/Artificial-Intelligence-Machine-Learning-Risk-Management.aspx>; Federal Reserve, OCC, *Guidance on Model Risk Management* (April 2011) <https://www.federalreserve.gov/supervisionreg/srletters/sr1107a1.pdf>.

Agencies have yet to issue guidance even after receiving a robust response to their Request for Information and Comment on Financial Institutions' Use of Artificial Intelligence in March of 2021.⁷⁴

To retain America's competitive edge in the global society, the Agencies should move quickly to issue actionable policy statements that clearly state their commitment to consumer protection and civil rights laws, including fair lending laws; insight into their supervisory expectations and methods; and useful guardrails and best practices. The time to act is now as the use of AI proliferates in every aspect of consumer financial services and housing and has the potential for far-reaching adverse impacts for consumers of color and other protected groups.

Integrating the Review of Racial Equity in the Algorithm's Lifecycle and Requiring Auditing Requirements for AI and Automated Decisions in Housing and Lending

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 model risks. Legislation must encourage Responsible AI innovation while implementing risk mitigation measures, and Congress can use tools like NIST's Risk Management Framework⁷⁵ and NFHA's AI auditing framework⁷⁶ to help protect consumers, communities, and the U.S. economy. Adequate governmental funding is essential for research, tool development, and support of civil rights, research, and consumer protection organizations to ensure equitable access to resources and address the lack of focus from philanthropies on these critical issues.

Promoting Effective Training for the Federal Workforce

Many in the federal workforce currently lack sufficient knowledge about how data-driven systems can perpetuate bias and discrimination. Many federal workers also do not have the requisite training needed to effectively test, monitor, and provide oversight over automated systems. To keep Americans safe, Congress should support comprehensive training on technology and AI bias for federal regulators and enforcement agencies and ensure the federal workforce has the equipment and resources needed to enforce U.S. laws and regulations. Training should include a focus on fair housing and racial equity. Trained professionals are

⁷⁴ CFPB, Federal Reserve, FDIC, OCC, and NCUA, *Request for Information and Comment on Financial Institutions' Use of Artificial Intelligence, including Machine Learning* (March 31, 2021) <https://www.govinfo.gov/content/pkg/FR-2021-03-31/pdf/2021-06607.pdf>. See also, NFHA et al., Response to RFI re AI (July 1, 2021), <https://nationalfairhousing.org/leading-civil-rights-consumer-and-technology-advocates-urge-the-federal-financial-regulators-to-promote-equitable-artificial-intelligence-in-financial-services/>.

⁷⁵ See NIST, *Artificial Intelligence Risk Management Framework*, U.S. Department of Commerce (Jan. 26, 2023), <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>.

⁷⁶ See Michael Akinwumi, Lisa Rice, and Snigdha Sharma, *Purpose Process, and Monitoring: A New Framework for Auditing Algorithmic Bias in Housing and Lending*, National Fair Housing Alliance (2022), https://nationalfairhousing.org/wp-content/uploads/2022/02/PPM_Framework_02_17_2022.pdf.

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, many recent innovations in developing mechanisms for debiasing tech have emerged from data scientists and engineers who were trained on issues of fairness.

Ensuring Equitable Digital Access, Public Data Access, Transparency and Explainability

America is at a critical juncture in deciding whether to develop automated systems that continue to perpetuate historical discriminatory practices or create equitable automated systems that uplift society. Currently, access to broadband and other technologies is not available on an equitable basis. For example, rural areas and communities of color disproportionately do not have access to high-speed internet. These communities also have higher instances of cell phone dead zones. Thus, government, industry, and advocacy groups must collaborate to ensure access to technological solutions are equitable and nondiscriminatory, as its far-reaching impact could either benefit or continue to severely disadvantage underserved groups and society at large.

AI legislation should mandate public availability of key data, as the lack of such data hampers efforts to develop responsible automated systems in housing and financial services. This data usage must rightly balance privacy rights with the need to protect civil and human rights.⁷⁷ Congress should encourage and support public research that analyzes the efficacy of specific uses of AI in housing and financial services and the impact of AI in financial services for consumers of color and other protected classes. For example, Congress should encourage the Consumer Financial Protection Bureau and the Federal Housing Finance Agency to release more de-personalized loan-level data from the Uniform Appraisal Dataset, National Survey of Mortgage Originations, and the National Mortgage Database so trusted researchers, advocacy groups, and the public can study potential discriminatory and inequitable outcomes in the housing and financial services, especially as they relate to the use of AI.

Finally, legislation must also ensure that AI and automated decision-making is transparent, meaning there are explanations for automated decisions and work conducted to test and promote methodologies that clarify the reasoning or design behind automated systems.

Ensuring Technologies Developed Outside of the U.S. Adhere to U.S. Rules and Regulations

Any legislation passed by Congress must ensure that foreign innovations adhere to U.S. civil rights and consumer protection standards to prevent violations through irresponsible AI.

Improve Consumers' Ability to Have Agency Over Their Data

Disclosure and notice and consent requirements are insufficient means of providing consumers agency over their data used in housing and financial services decisions. In addition to adequate privacy protections, consumers must consent to how, where, when, and under what circumstances their personal data will be utilized. The European Union and California are

⁷⁷ To more fully understand NFHA's position on balancing civil rights with privacy in housing decisions, see NFHA and Tech Equity Collaborative, *Privacy, Technology, and Fair Housing - A Case for Corporate and Regulatory Action* (Aug. 2023), <https://nationalfairhousing.org/wp-content/uploads/2023/08/NFHA-TechEquity-Paper-final.pdf>.

currently leading in this space. Congress should learn from their examples and adopt these protections and build upon them for using consumers' personal information. Legislation must clarify that consumers own their personal data. Data collected for one purpose cannot be used for other purposes without consumers' consent and assess for existing privacy risks that incorporate proper risk mitigation; and data minimization frameworks coupled with discrimination testing must be required. The legislations should also provide requirements for the development and testing of privacy enhancing technology, strengthen existing regulatory consumer complaint systems to incorporate protocols for AI and automated systems, and clarify and increase damages under existing laws for AI or automated systems generated harms.

Conclusion

It is imperative that the U.S. lead the world in establishing policies and frameworks to advance technological innovations while ensuring these systems are fair, safe, transparent, explainable, and reliable; protect consumers' privacy; and ensure that human alternatives are available when warranted. Technological innovations can provide great benefits to people and society as well as spur economic progress. Yet too many automated systems have been deployed without proper protocols, testing, and oversight. As a result, people have unfairly and inappropriately been denied housing, credit, other important opportunities and services. Researchers have found that racial inequality has cost the U.S. economy \$16 trillion over the past 20 years.⁷⁸ Congress must move with all haste to guarantee the U.S. can remain productive, strong and viable and our citizens can benefit from technological innovations.

⁷⁸ Dana Peterson and Catherine Mann, *Closing the Racial Inequality Gaps*, Citigroup, (Sept. 2020), <https://www.citigroup.com/global/insights/citiqps/closing-the-racial-inequality-gaps-20200922>.

PREPARED STATEMENT OF VANESSA PERRY

INTERIM DEAN AND PROFESSOR, THE GEORGE WASHINGTON UNIVERSITY SCHOOL OF
BUSINESS; NONRESIDENT FELLOW, HOUSING FINANCE POLICY CENTER, URBAN IN-
STITUTE

JANUARY 31, 2024

Good morning. Thank you for inviting me to address the impact of Artificial intelligence, i.e., AI, which is being employed increasingly throughout the housing and mortgage industry. For these purposes, AI refers to the use of data and algorithms in place of human decisions. This definition includes machine learning models, which are programmed to imitate the way humans learn, iteratively correcting themselves to improve their accuracy.

Compared to traditional models, AI relies on a wider range of data inputs and more complex combinations thereof. Although complex multivariate algorithms have been in place in the mortgage market for years, these models have the potential to incorporate nontraditional data sources. Due to their complexity, it is difficult, but not impossible, for anyone other than AI developers to scrutinize and monitor their inputs.

AI models are already widely applied in the mortgage market. AI digital marketing models target prospective homebuyers and communications with customers are intermediated by AI chatbots. Credit scoring companies and mortgage underwriting systems use AI to evaluate credit risk. AI models are widely used for property valuation, loan servicing, and loss mitigation. AI regulation warrants urgent attention because evidence from other domains and my research with the Housing Finance Policy Center at the Urban Institute suggests that while these models can enhance efficiency, they can have unintended impacts on fairness and equity.¹

AI models are not subject to human errors, and they enable efficient, accurate, and consistent decisions. Depending on how they are developed, their enhanced capabilities could expand access to home ownership for households currently underrepresented in the mortgage market. For example, AI can produce faster and less subjective estimates than human property appraisals and can devise credit scores for those who lack a traditional credit history.

However, because these models rely on historical data, there is the potential for these models to systematize and amplify discrimination and inequality.² For example, due to the legacy of redlining and segregation and their effects on present-day neighborhood conditions and home values, why should we expect AI models to produce estimates that are both accurate and fair? And absent guardrails, how would we know if AI models were to incorporate data elements, such as GPS location, that serve as a proxy for race, gender, or ability?

To address concerns about AI's impact on access to the housing finance system for underrepresented or marginalized communities, my coauthors and I have proposed five factors summarizing the societal, ethical, legal, and practical issues that should be considered in the development and implementation of AI.³ They form a memorable acronym, S.C.A.L.E., which stands for:

- *Societal values.* Algorithms tell us what factors the developer thinks are important, in what order, and to what degree. AI models should consider the socioeconomic and historical context (e.g., past discrimination) and should align with prevailing legal and ethical paradigms, e.g., disparate impact law, individual freedom, and racial equity.
- *Contextual integrity.* In addition to its accuracy, model inputs should be relevant to the mortgage or housing domain, and may differ substantively from those used for other or less consequential contexts.
- *Accuracy.* Models should be reliable, error-free, unbiased, and representative of all demographic and economic groups across varying macroeconomic conditions.

¹Michael Neal, Linna Zhu, Caitlin Young, Vanessa G. Perry, Matthew Pruitt (2023), "Harnessing Artificial Intelligence for Equity in Mortgage Finance", Urban Institute, November 6, <https://www.urban.org/research/publication/harnessing-artificial-intelligence-equitymortgage-finance>.

²Michael Neal, Linna Zhu, Vanessa G. Perry (2024), "To Err Is Automated: Have Technological Advances in the Mortgage Market Increased Opportunities for Black Homeownership?" *Journal of the Center for Policy Analysis and Research*, forthcoming, <https://papers.ssrn.com/abstract=4347212>.

³Perry, V., Kirsten Martin, and Ann Schnare (2023), "Algorithms for All: Can AI in the Mortgage Market Expand Access to Homeownership?", *AI*, 4(4), 888-903, <https://doi.org/10.3390/ai4040045>

- *Legality.* The model and its inputs, if used for housing or housing finance decisions, should not incorporate characteristics protected by fair lending laws or generate unjustified disparate impacts based on these characteristics.
- *Expanded opportunity.* AI models should significantly increase access to credit in addition to offering greater cost efficiency or risk assessment benefits. This criterion has perhaps the most promising impact on the economy and communities.

In terms of policy directions, the S.C.A.L.E. framework could inform new or expanded regulations, such as guidance for the use of certain types of data—such as an individual’s social media profile—for certain purposes, such as mortgage lending decisions.

While the S.C.A.L.E. criteria imply that model inputs and algorithms, due to the complexity of AI models, regulators cannot rely on traditional approaches to documentation and testing. Furthermore, concerns about potential harms related to AI are domain-specific, suggesting that regulation and enforcement efforts must be targeted specifically to housing and mortgage applications.

If designed to do so, AI models can increase access to home ownership and eradicate the effects of systemic discrimination while increasing accuracy and efficiency in the mortgage value chain. We need laws on the Federal level that turn the “S.C.A.L.E.” toward imposing these standards at every stage of the AI life cycle.

PREPARED STATEMENT OF NICHOLAS SCHMIDT

PARTNER AND ARTIFICIAL INTELLIGENCE PRACTICE LEADER, BLDS; FOUNDER AND CTO, SOLASAI

JANUARY 31, 2024

Chair Smith, Ranking Member Lummis, and Members of the Subcommittee, thank you for hosting this important hearing and for giving me the opportunity to submit this testimony.

My name is Nicholas Schmidt, and I am the founder and Chief Technology Officer (CTO) of SolasAI, as well as a Partner and the AI Practice Leader at BLDS, LLC. My passion for the responsible and fair application of artificial intelligence (AI) and machine learning (ML), including in housing, has driven my career, leading to the development of SolasAI, a software platform dedicated to mitigating bias and discrimination in algorithmic decision making.

At SolasAI, we create software that not only addresses regulatory, legal, and reputational challenges but also empowers innovation in model development. Our work centers on reducing disparate impact and eliminating disparate treatment, ensuring that AI tools used in housing are equitable, effective, and profitable.

Today, I am here to share insights on how properly implemented and regulated machine learning and artificial intelligence can transform the housing sector. I aim to provide a comprehensive understanding of the potential for AI to make high-quality housing more affordable and accessible, while also addressing the critical need for fairness and equity in these technologies.

Defining Artificial Intelligence

Before one can hope to craft effective laws or regulations around the use of artificial intelligence, we must first define it and understand its scope. While the term “AI” often conjures a vision of futuristic and sentient machines, in practical terms, AI encompasses a wide array of far less radical technologies.

Contrary to the popular focus on—and hype around—generative AI, AI’s impact on society extends through various types of machine learning (ML) and AI applications, many of which are already transforming the housing industry.¹ What can be defined as AI includes technologies ranging from predictive analytics to automated decision-making systems, all of which impact the affordability, accessibility, and equity of housing in the United States.

In practical terms, ML and AI represent a class of mathematical algorithms that learn patterns and rules from data. These learned rules may then be applied to new

¹Significant ink has been spilled over the attempt to precisely define the difference between “AI” and “ML.” Commonly, people refer to deep neural network-based techniques that are typically used for image, text, and language recognition or generation as “AI,” whereas other techniques that are used for things like credit underwriting or pricing (including, especially, tree-based ensemble models) are referred to as “ML.” A more robust and thorough discussion of the terms can be found here: <https://mitsloan.mit.edu/ideas-made-to-matter/machine-learning-explained> and <https://aima.cs.berkeley.edu/newchap00.pdf>.

data to inform real-world decisions. Thus, it is important to remember that the rules AI develops are based on mathematical (i.e., not human) insight, but that those rules are developed on historical data that encode many types of human biases. However, as I discuss below, despite the computer developing the rules, human decisions affect how the rules are developed and systems are used. Understanding this is essential for writing effective legislation.

Human Decisions Drive Algorithms

Beyond the confusion of what “AI” is, many people are unaware of how much human involvement is required to build and deploy an AI (or “algorithmic decisioning”) system. In fact, particularly coming from technology companies, there seems to be a fatalistic attitude that implies nothing can be done to improve them. This notion, dangerous in the extreme, is easily proven wrong. There are numerous places where humans interact with AI systems before and during the deployment of the algorithms that shape whether the algorithms are making reasonable, safe, and fair decisions. Understanding the extent to which the output of an algorithm is dependent on the decisions that people who build the models make is essential because—at each of these decision points—there are opportunities for humans to make better decisions that can make the algorithms more fair, accountable, and transparent.

Using a mortgage delinquency algorithm as an example, the human steps required to make such a model include, e.g.:

1. *Choosing what the model will define as a delinquency.* A data scientist might define delinquency as 60 days or 90 days of nonpayment. For people with less income security, but who are likely to be able to pay their bills over a longer period, choosing 60 days instead of 90 days may make the difference between being provided a loan or being rejected. Importantly, this decision is an entirely human-based decision.
2. *Choosing which data will be used to predict delinquency.* The computer only makes its rules based on the data it is provided. The person building the model might only include data clearly related to delinquency (e.g., the existing level of debt, past payment history, etc.), or they may include data that is not clearly causally related to delinquency (e.g., education or purchase history). The choices that the modeler makes will affect the fairness of the model, its accuracy, its reasonableness, and its reliability. While algorithms will choose how to weight different data (and possibly exclude certain data altogether), choosing which data to provide is an entirely human-based decision.
3. *Choosing the type of algorithm that will build these rules.* Model developers have many options regarding what model architecture they will use to develop the model (i.e., how the model will learn from the data). These include architectures like deep neural networks, gradient-boosted trees, or traditional linear regression. This decision has direct implications in terms of the transparency of the model and whether the model’s decisions will be explainable. Some kinds of models, such as regression models, are inherently explainable. This means it is possible to know exactly how the model arrived at a decision.² Others, such as neural networks, are not transparent. The decisions surrounding model architecture have practical and legal implications, as I explain *infra*.
4. *Choosing whether the algorithm is sufficiently accurate for the task.* No algorithm is perfect. One issue is that a model might be very effective for some people, but not others. For example, it might be highly accurate for high-credit quality individuals but be far less accurate for those with lower credit quality. A person working for the lender will ultimately make a decision whether that

²Providing a hypothetical example of how easy it is to interpret a linear regression, suppose we build an AVM using linear regression that uses three variables to predict home price: number of bedrooms, number of bathrooms, and number of square feet. With linear regression, we know the effect of each of these variables on a particular house’s estimated price because linear regression is completely transparent. Looking at the model, we could know that, for example, each bedroom adds \$50,000 to the value. We could also explicitly tell that an additional bathroom adds \$25,000 to the value. Finally, we could see that every additional 100 square feet adds \$5,000 to the value. On the other hand, a more complicated model architecture could learn from the interaction of all of these variables in highly nonlinear ways such that increasing the square footage of the house by 100 feet adds between \$0 and \$50,000, depending on the number of bedrooms and bathrooms. While, in this particular example, a more complex model would be more accurate than the regression model, it would also likely be very difficult to interpret and fully understand the model. This becomes particularly problematic when we move away from this simple three variable example to a more realistic ML model, which might include dozens, hundreds, or maybe thousands of variables.

trade-off is acceptable before putting the model into production. More generally, whether to test for different forms of inaccuracy, how to balance the varying costs of inaccuracy, and what minimum accuracy requirements are required are all human decisions.

5. *Choosing how to implement and use the algorithm.* Generally, an AI or ML model does not make a decision on its own. Frequently, only a subset of applicants will be scored by the model; which applicants are scored is a human decision. Once applicants are scored, what that score means and how it is used must be determined. For example, will a cutoff be used, or will the specific model work with another model or subjective rules to make a decision? Will the score be considered in light of other variables or factors? All of these are human decisions. In fact, when all of these decisions are examined together, it is clear that the entire model system, of which the model itself may only play a relatively modest role, is largely made up of interlocking human decisions that result in an ultimate decision.

I take pains to point out each of these decision points because, as NIST-1270 puts it, “A fallacy of objectivity can often surround these processes, and may create conditions where technology’s capacity and capabilities are oversold.”³

This fallacy of objectivity has led many to conclude that not much can be done to regulate or effectively manage these technologies. However, because so much of the use of AI is driven by choices that people make, regulators and the law do not need to “surrender” to these emerging technologies; the space is ripe for regulation of human decisions. In fact, effective regulation of these human decisions can create fairer, more equitable outcomes without stifling innovation in this space. But more than just a benefit to consumers, having defined and reasonable regulations would give companies more clarity on how they can safely use AI systems.

I next provide context for how AI is used in housing by discussing common use cases and outlining the challenges and opportunities they offer.

The Use of AI Systems in Housing

AI systems are increasingly being used across the housing industry, as companies find that many facets of housing decisions can be made far faster, cheaper, and more reliably than can be performed by humans. While, in past years, there was significant discussion about whether algorithms should be making such decisions, we are now in a world where the use of such algorithms is commonplace; companies that do not employ them are at great risk of losing out to those that do. Below, I outline several areas where I have seen algorithms used in housing and discuss the opportunities and challenges associated with their use.

Many entities have used algorithms to underwrite and price mortgages for decades. Fannie Mae’s Desktop Underwriter (DU) and Freddie Mac’s Loan Product Advisor (LPA) have histories dating back nearly 30 years, while FICO introduced its first consumer score in 1989. More recently, human appraisals are being replaced with automated valuation models (AVM), allowing fast and—hopefully—accurate assessment of the value of a home. Other ways algorithms are used include the provision and pricing of insurance, providing background screening for rentals, automating the servicing of loans, and pricing rental units.

There are two noteworthy takeaways from these applications of AI in the housing industry. First, there is a long history of the use of algorithms in housing. Correspondingly, there is a wealth of experience in building these algorithms fairly, reliably, and transparently. As such, we do not have to reinvent the wheel when it comes to effectively regulating the responsible use of algorithms. Second, the reach of algorithms in the housing industry is growing fast, which will profoundly affect people’s housing decisions. It is imperative that we learn from what we know about safely implementing algorithms in housing and apply that knowledge to these newer applications.

However, before discussing how the history and practices of regulating and monitoring the use of algorithms can be extended and made better in the world of AI, I will discuss a particular use case of algorithmic decisioning—automated valuation models—because they represent a concrete example of many of the benefits and challenges inherent in the expanded use of algorithms and AI.

³The sentence preceding this quote summarizes the issues I address: “Adding to the challenge is the reality that these systems are built and placed within organizational settings along with their accompanying—often unstated—policies and priorities, and used by subject matter experts and decision makers who have their own implicit heuristics and biases.” NIST SP-1270, “Towards a Standard for Identifying and Managing Bias in Artificial Intelligence”, p.25.

Case Study in the Use of Algorithms: Automated Valuation Models

As anyone who has purchased a house knows, obtaining an appraisal can be a drawn-out, stressful, expensive, and opaque process. Further, there is substantial evidence that human appraisals often suffer from significant discrimination, with Black or African American homeowners and people living in predominantly Black neighborhoods having their homes appraised for far less than they would have been if they were White or lived in a predominantly White neighborhood.⁴ Additionally, people living in rural areas have found it difficult to get an appraisal of their homes, and have found that appraisals are often more expensive.⁵ ⁶ Higher costs and delays in securing an appraisal may make transactions in rural areas take longer, cause sales not to occur, or prevent home buyers or refinancers from locking in lower rates.

A potential promise of automated valuation models (AVMs) is that they may solve these problems. AVMs replace the job of a human appraiser by predicting the value of a home using a wide range of data, including information such as a home's square feet, number of bathrooms, size of yard, tax history, location, sales prices of similar homes, and many other factors. While they have a long history of use by financial institutions in valuing portfolios of properties, they are now being used in a way that has a far more direct impact on consumers. Lenders are using these estimated home values for many financial decisions, including as a factor in deciding whether to offer a loan, provide favorable terms, or offer refinancing. As such, they may significantly affect a person's finances and life, influencing where they live and how much they pay for housing. Their use is instructive for understanding the benefits and harms of algorithmic decisioning.

AVMs' obvious and clear benefits are their speed, availability, and price. The National Association of Realtors (NAR) reported that, in 2022, the average cost of a home appraisal was \$500, with 9 percent of appraisals costing more than \$800.⁷ Once an AVM is built and running, the cost of calculating an appraisal is virtually zero for the model owner (though borrowers could still be charged for this service). NAR additionally reported that the median response from realtors indicated that it takes 14 days for a lender to return a completed appraisal, with 4 percent of realtors reporting that it typically takes them more than 30 days to receive a completed appraisal. In comparison, obtaining an appraisal from an AVM is virtually instantaneous; even if the lender performs subsequent reviews of the AVM-based appraisal, an AVM-based appraisal is likely to be available substantially faster than a human-based appraisal. It is likely that the use of AVMs has already ensured that home sales have closed faster, instead of languishing or being rejected.

Another benefit of an appraisal from an AVM is that it may be less discriminatory than an appraisal made by a human. A blog post from the CFPB describes a lawsuit where a Black couple's home appraised for nearly 60 percent more after they "white-washed" their home—removing evidence that they were Black.⁸ Given that the data used in an AVM would not incorporate information about the race of the homeowners, this type of discrimination is unlikely to manifest in a well-built AVM. Put another way, human appraisers are subject to cognitive biases they may not even be aware of; AVMs are not.

However, this is not to say that AVMs will be free from any discrimination or bias.⁹ AVMs are not magical—if they create accurate appraisals, those appraisals will follow the real estate maxim of "location, location, location." Of course, because of the history of housing discrimination in the United States, a home's location necessarily incorporates historical and present patterns of discrimination and bias. As a result, because the data used to build AVMs is necessarily biased, the output of the models will be as well. AVMs also cannot help when there may be insufficient data to draw conclusions.

To summarize, the discriminatory effect of AVMs is likely to be mixed. Because an AVM does not have access to or incorporate information about individual sellers or borrowers, it cannot discriminate based on those factors. However, it will still incorporate price-affecting discrimination in the data it sees. As a result, we can view

⁴ <http://tinyurl.com/3w5s2dak>

⁵ <https://www.knock.com/blog/how-long-does-an-appraisal-take>

⁶ <https://www.homelight.com/blog/buyer-how-long-does-an-appraisal-take/>

⁷ <https://www.nar.realtor/sites/default/files/documents/2022-appraisal-survey-09-27-2022.pdf>

⁸ <http://tinyurl.com/2jt4vtrs>

⁹ Throughout this report, I use the terms "bias," "discrimination," and "disparities" more or less interchangeably, according to their lay definitions. However, it should be noted that each of these has distinct definitions and interpretations in technical and legal settings. Defining these terms precisely becomes particularly important when measuring and mitigating discrimination.

AVMs as a tool to help fight discrimination in home pricing, but not one that is complete or without significant residual problems.

My work, and the software we have developed at SolasAI, measures bias and discrimination in models such as AVMs. As explained above, data may be biased due to historical patterns of discrimination; when a model is trained on such data, it results in a biased model. Our work focuses first on measuring whether a model shows evidence of unfair disparities and, if it does, attempts to identify the source of those disparities. If such disparities are found, the software searches for what is known as a “less discriminatory alternative” (LDA) model. These LDA models provide predictive power similar to or equal to the original model but have a less discriminatory effect.

In addition to these remaining concerns about discrimination, AVMs raise at least two other challenges: interpretability and accountability. Many (but not all) AI and ML algorithms are described as “black-box” processes, which means that, while we understand what data went into the algorithm, and we can see what the output of the algorithm is, it is very difficult to understand how or why the algorithm made the prediction that it made. In the case of a black-box AVM model, a lender might not be able to provide a borrower with a reliable explanation for why the home they wanted to buy received a low appraisal. While this lack of clarity may be acceptable in low-risk situations, or when the algorithm gets its prediction right, having opaque black-box models make high-risk decisions that cannot be explained may not be fair to people who receive unfavorable outcomes.

The lack of explainability also raises concerns about the quality of the models. If we cannot sufficiently understand why a model gives a particular prediction, then we should have less confidence in that model. The problem is that, while it may appear to be a high-quality model based on the data it has seen, if we cannot understand the model, then we cannot be sure that its decision-making process makes sense and is likely to continue to operate well if conditions change. As a result, for business-critical or high-human-impact decisions, model builders need to balance a desire for accuracy with the necessity that the algorithm be explainable.¹⁰

The points outlined above about AVMs apply to virtually every algorithm making housing-related decisions. With regard to discrimination, an algorithm is not inherently discriminatory, but it can discriminate if it is given discriminatory data, poorly built, or used in the wrong context. However, as we saw in the context of AVMs, even if they suffer from known biases, they may still represent a better option than human-based decisions.

Having used AVMs to illustrate the benefits and perils of using AI more generally, the next section outlines concrete steps that legislators and regulators can take to minimize likely harm while encouraging and fostering safe innovation in AI.

The Role of Regulators and Policymakers in Ensuring Responsible Innovation in AI

As we consider regulations for AI in housing, the primary goal should be to maximize the responsible use of this technology: given its potential to cause extreme harm at scale, safe and sound implementation of AI technologies is paramount. However, given its potential to represent a significant technological development that delivers real and meaningful societal benefits, we should also aim to minimize regulation’s potential to be overly burdensome, possibly stifling innovation.

Effectively accomplishing this goal is significantly more likely if we consider two key factors. First, any approach to regulation should not be overly prescriptive. Instead, we should focus on setting clear risk-based principles that encourage and enforce responsible AI development and use, where the most impactful systems (i.e., the systems with the most potential to harm or benefit people, society, or the environment) receive the most oversight. Second, it will be valuable to recognize that a significant body of existing work, regulation, and industry practice can be applied to AI systems to make them safer. Looking to these tried and validated frameworks and policies should guide our approach to making effective regulations for the use of AI.

The Benefits of a Nonprescriptive Regulatory Environment

A principles-based and less prescriptive approach to AI regulation can encourage innovation while ensuring the responsible development of AI. It recognizes the dy-

¹⁰There is a general understanding that more opaque “black-box” algorithms are better able to capture hidden patterns in data than less opaque, “white-box” algorithms. To the extent this is true, a black-box algorithm will likely be more accurate. However, there has been significant work done to show that in many contexts, less opaque and highly interpretable models perform just as well, or virtually as well, as opaque models.

dynamic nature of technology and compliance and provides flexibility for continuous improvement and adaptation.

A principles-based framework allows for innovation in both technology and compliance methods. On the compliance side, advancements like improved Less Discriminatory Alternatives (LDA) search and enhanced techniques for providing Adverse Action Notices (AANs) demonstrate how technology and compliance can complement each other and evolve together. Further, innovations in technology, such as Shapley values, explainable boosting machines, and Wells Fargo’s Python Interpretable Machine Learning (PiML) package illustrate the rapid development of new AI tools and methods that encourage responsible model building. A less rigid but strong regulatory environment encourages such advancements.

Overly prescriptive regulations, on the other hand, risk stifling innovation as they may lead to a “design-around” mentality, where the focus shifts from responsible development to merely meeting specific regulatory criteria. This can hinder genuine progress and the exploration of new AI and negate the desired helpful impact of the regulations. It also risks enforcing technical requirements that quickly drift into irrelevance as technology evolves.

Key Principles To Consider for AI Regulation

Four fundamental principles are foundational for adopting effective AI regulation: materiality, fairness, accountability, and transparency. Developing regulations with these as guideposts will help ensure that AI systems serve the public interest while advancing technological progress.

Materiality:

This principle advocates for a risk-based approach in governing AI systems. By focusing more stringent regulation on higher-risk AI applications, resources will be allocated more effectively. For example, a company should not spend as much time reviewing a marketing model as they would an underwriting model that enormously impacts both consumers and the business. Adopting such a risk-based approach ensures that systems with the most significant potential impact are carefully monitored and promotes innovation by not overburdening lower-risk initiatives with unnecessary regulatory constraints. As discussed below, SR 11–7 provides a solid foundation for guiding how materiality is assessed in AI regulation.

Fairness:

The principle of fairness is central to the responsible deployment of AI. Establishing a clear understanding and expectation of fair AI practices is crucial, particularly in applications that significantly impact individuals, such as housing. Regulators should set expectations that bias and discrimination should be identified and mitigated in AI systems. Existing frameworks relating to measuring and mitigating disparate impact, disparate treatment, and proxy discrimination should guide further regulation of AI fairness.

Accountability:

AI systems must have accountability mechanisms, especially those with high impact. This involves providing individuals affected by AI decisions a right to appeal, ensuring that there is recourse for those who may be adversely impacted. Additionally, entities that deploy AI systems irresponsibly should face appropriate consequences.

Transparency:

The principle of transparency mandates clear explanations for decisions made by AI systems. This is fundamental to building trust in AI systems. Understanding the “why” and “how” behind AI-driven decisions is crucial for public acceptance and confidence in these technologies, and is further crucial to ensure that systems are fair.

By focusing on these critical principles—materiality, fairness, accountability, and transparency—we can create a regulatory environment that encourages the development of innovative AI technologies and safeguards against potential harms and biases. This principles-based approach to AI regulation is particularly pertinent in housing, where the impact of AI can have profound implications on people’s lives and the fabric of communities. Next, I discuss a number of frameworks that can serve as guides for future regulation.

Using Existing Regulations and Frameworks To Guide Further AI Regulation

Congress and regulators will not need to devise laws and regulations from scratch to achieve the objectives that I have laid out. There are many regulations, standards, and frameworks with a proven track record of setting standards for human decisions related to AI, holding relevant actors accountable for those standards, and supporting the development and deployment of these technologies. Importantly, in

many (but of course not all) cases, the industry has welcomed these for providing clear and reasonable standards. I discuss these at a high level below.

- SR 11–7,¹¹ a supervision and regulation letter from the Board of Governors of the Federal Reserve System, constructs accountability mechanisms and organizational structures to ensure adequate and risk-based governance of credit modelers. While the document highlights the importance of technical processes such as testing and monitoring, its primary focus is on principles such as effective governance structures (e.g., independent validation teams with high stature and strong financial incentives), risk management executives with independent reporting chains, and documentation requirements.
- NIST SP 1270,¹² a special publication from the National Institute of Standards and Technology (NIST), describes technical, process, and cultural problems and solutions relating to AI bias. It highlights that many aspects of data and AI systems are strongly influenced by human behavior and decisions, and suggests that approaches from model risk management (e.g., SR 11–7) coupled with more novel approaches, such as structured feedback activities (e.g., bug bounties or red teaming), human-centered design, and information sharing are strong mitigants for managing bias in AI systems. Another prominent theme of NIST SP 1270 is that basic scientific rigor in AI needs to be improved.
- The disparate impact, disparate treatment, and proxying framework is a legal doctrine that has been developed over the course of decades under the Fair Housing Act (FHA), the Equal Credit Opportunity Act (ECOA), and Title VII of the Civil Rights Act of 1964. This framework sets forth requirements for measuring discrimination that could be used in any decision tool (including AI or ML models) and provides a conceptual framework for mitigating any meaningful disparities found through conducting an LDA search. Many AI tools that affect consumers in the housing market are likely covered by this framework via the FHA. Setting the expectation that this framework would be extended to all high-risk AI use cases throughout the housing industry would ensure that companies move towards adopting fairer models.¹³
- The NIST AI Risk Management Framework¹⁴ puts forward four risk management functions for organizations: (1) Govern, (2) Map (understanding the risk of AI systems in their operational contexts with less emphasis on their development), (3) Measure, and (4) Manage; and seven trustworthy characteristics for AI systems: (1) Safe, (2) Valid and Reliable, (3) Accountable and Transparent, (4) Explainable and Interpretable, (5) Privacy-enhanced, (6) Fair with Harmful Bias Managed, and (7) Secure and Resilient. Governance guidance is largely aligned with the risk-based principles laid out in SR 11–7 but introduces additional governance concepts from data privacy, information security, and more recent academic research. The AI Risk Management Framework has two distinct strengths: it acknowledges (1) that many AI risks arrive from real-world problems, not computer code bugs; and (2) that overlaps and connections between risks, risk controls, and governance must be recognized. It does all this while orienting governance toward traditional risk-based principles focusing on human accountability.

Conclusion

The most significant harms associated with AI are not the fantastical scenarios often depicted in science fiction, but real-world issues such as discrimination, data privacy violations, unaccountable decision making, and fraudulent activities. Effectively regulating AI systems requires recognizing these facts. As AI evolves and impacts more aspects of housing, policymakers, regulators, public advocacy groups, and industry professionals must remain vigilant and proactive. We each play a key role in ensuring that AI systems are not only technically sound and effective, but are also fair, transparent, and accountable. This will require ongoing collaboration, research, and adaptation of regulatory approaches.

Regulating AI so that it is used responsibly and safely and so that it can continue to benefit consumers and industry is a complex and crucial task. Materiality, fairness, accountability, and transparency provide a framework of principles that can

¹¹<https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm>

¹²<https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270.pdf>

¹³Schmidt, Nicholas, and B. Stephens. “An Introduction to Artificial Intelligence and Solutions to the Problem of Algorithmic Discrimination”. *Conference on Consumer Finance Law (CCFL) Quarterly Report*. Volume 73, Number 2 (October 2019). <https://arxiv.org/abs/1911.05755>

¹⁴<https://www.nist.gov/itl/ai-risk-management-framework>

serve as a touchstone for evaluating the effectiveness of proposed regulations. Further, leveraging existing regulations and frameworks such as SR 11-7, NIST SP 1270, the disparate impact framework, and the NIST AI Risk Management Framework, regulators and legislators can create an environment that fosters innovation and protects against potential harm.

The future of AI in housing presents both exciting possibilities and significant responsibilities. By embracing a principles-based regulatory approach, drawing on existing frameworks, and remaining open to continuous learning and improvement, we can build AI that reaches its potential to improve housing affordability, accessibility, and equity, while safeguarding the rights and interests of all individuals.

Thank you again for the opportunity to share my insights and perspectives on this critical issue.

**RESPONSES TO WRITTEN QUESTIONS OF CHAIR SMITH
FROM LISA RICE**

Q.1. How do we build incentives into housing and regulatory systems so that we are maximizing fairness, accountability, and transparency as we try to set guidelines both to encourage innovation but also to make sure these systems are working for the public? Can and should this be handled legislatively? Dr. Perry underscored the importance of holding “AI accountable for expanding opportunity,” and I welcome legislative and regulatory recommendations for how to best incentivize this.

A.1. Response not received in time for publication.

Q.2. You underscored the bevy of existing laws that apply to and cover most of the AI applications in housing. What are the legislative gaps you see in this space, and how do you recommend that Congress addresses them?

A.2. Response not received in time for publication.

**RESPONSES TO WRITTEN QUESTIONS OF SENATOR WARNOCK
FROM LISA RICE**

Q.1. The Fair Housing Act¹ and Equal Credit Opportunity Act² are landmark statutes to prevent discrimination in housing. Some Federal regulators have already used authorities under these statutes to engage in oversight of more advanced models that employ novel forms of data and more sophisticated statistical techniques.³

A.1. Response not received in time for publication.

Q.2. In your view, are Federal regulators using all available authorities under FHA and ECOA to engage in oversight of novel software or models that employ machine learning (ML) or artificial intelligence (AI)? Why or why not?

A.2. Response not received in time for publication.

Q.3. Are there any ways in which existing statutes, authorities, or resources fail to enable policymakers and regulators to engage in oversight of novel software or models using ML and AI? If so, what additional authorities are most helpful? Please provide specific examples.

A.3. Response not received in time for publication.

**RESPONSES TO WRITTEN QUESTIONS OF CHAIR SMITH
FROM VANESSA PERRY**

Q.1. How do we build incentives into housing and regulatory systems so that we are maximizing fairness, accountability, and transparency as we try to set guidelines both to encourage innovation but also to make sure these systems are working for the public?

A.1. Existing regulatory frameworks in the lending industry, including, aspects of the Fair Housing Act, e.g., the disparate impact

¹ See supra n. 1.

² See supra n. 2.

³ See Press Release, supra n. 3.

standard,¹ and Affirmatively Furthering Fair Housing (AFFH), can be used to establish standards for fairness and accountability in Artificial Intelligence and Machine Learning (AI/ML) systems. Existing reporting requirements (e.g., under the Community Reinvestment Act, the Home Mortgage Disclosure Act, and the Federal Housing Finance Agency’s conservatorship) collect data that can be used to evaluate outcomes, although expanded, real-time data, such as loan- and transaction-level detail, should be readily available to better inform regulators and policymakers about the impact these systems are having on the market.

The issue of transparency is controversial among scholars who study ethical AI due to concerns about the complex and dynamic nature of AI/ML models. The Consumer Financial Protection Bureau (CFPB) has recently issued guidance elaborating on the Equal Credit Opportunity Act requirement that creditors explain the specific reasons for denying loans, even when complex AI/ML models are used.² A distinguishing feature of AI models compared to algorithmic predecessors is that they often apply multiple factors simultaneously that can adapt to varying market conditions. Heretofore, these reasons have been listed in simple categories, e.g., credit history, or insufficient cash. In the case of AI/ML models, however, there could be numerous factors that may not matter by themselves but taken together result in loan rejection. To prevent consumer confusion, the CFPB should continue to provide standard categories of feedback, supported by empirical evidence on the messages that consumers are most likely to understand. At the same time, AI/ML models could be used to recommend actionable steps for denied applicants to take that would increase their likelihood of loan approval, e.g., how much debt reduction or savings would be necessary to qualify for a mortgage.

Financial regulators, e.g., the OCC, perform periodic reviews of underwriting models, and in the case of the Enterprises, model changes must be approved by FHFA in advance of adoption. To provide opportunities for innovation as well as regulatory oversight, President Biden’s recent executive order on AI relies on voluntary commitments from AI companies and “red-team-style”³ audits, but regulatory requirements and/or tax incentives could be instituted to encourage firms to participate.

Q.2. Can and should this be handled legislatively? You underscored the importance of holding “AI accountable for expanding opportunity,” and I welcome legislative and regulatory recommendations for how to best incentivize this.

A.2. To ensure that AI/ML models accomplish the goal of expanding opportunity,⁴ legislation could be established, mirroring the Affirmatively Furthering Fair Housing (AFFH) provision of the 1968

¹“Reinstatement of HUD’s Discriminatory Effects Standard”, Housing and Urban Development Department, 03/31/2023, <https://www.federalregister.gov/documents/2023/03/31/2023-05836/reinstatement-of-huds-discriminatory-effects-standard>.

²CFPB Issues Guidance on Credit Denials by Lenders Using Artificial Intelligence, Consumer Financial Protection Bureau, Sept. 19, 2023, <https://www.consumerfinance.gov/about-us/newsroom/cfpb-issues-guidance-on-credit-denials-by-lenders-using-artificial-intelligence/>.

³Red-teaming is the process of testing model effectiveness by applying an adversarial lens to your organization, <https://www.ibm.com/blog/red-teaming-101-what-is-red-teaming/>.

⁴Perry, V., Kirsten Martin, and Ann Schnare (2023), “Algorithms for All: Can AI in the Mortgage Market Expand Access to Homeownership?”, AI, 4(4), 888–903, <https://doi.org/10.3390/ai4040045>.

Fair Housing Act. AFFH requires that entities receiving Federal funding related to housing must take active steps to promote fair housing practices and combat housing discrimination, based on an analysis of impediments to housing options in their jurisdictions. A similar mandate could be incorporated into H.R. 6580, i.e., the proposed Algorithmic Accountability Act of 2022,⁵ stipulating that predeployment impact assessments identify how an AI system expands opportunity for underrepresented groups (e.g., low-income, rural area residents, residents in historically redlined neighborhoods). Organizations would subsequently be evaluated against these goals at the post-implementation evaluation stage (Algorithmic Accountability Act, Sect. 3).⁶ To strengthen these incentives, additional legislation could create Expanded Opportunity Tax Credits for organizations which can demonstrate that the adoption of a new system has significantly impacted specific goals—a model that combines approaches from Low-Income Housing Tax Credits and the affordable housing goals established for Fannie Mae and Freddie Mac. In this example, a mortgage lender would qualify for a tax credit if they could demonstrate that a new AI/ML system would result in a higher rate of mortgage approvals for low- and moderate-income homebuyers, or for residents of rural or other underserved areas.

H.R. 6580, known as the Algorithmic Accountability Act, directs the Federal Trade Commission (FTC) to require impact assessments of automated decision systems and augmented critical decision processes. Rather than place the entire burden of these assessments on the FTC, all regulators of entities in affected industries, e.g., housing, credit, and education, could implement these assessments. This would allow specialization and alignment with the agencies' mission and goals.

Q.3. We know financial technology companies and mortgage lenders are exploring new ways to incorporate AI into underwriting and property valuation. What roles do Fannie Mae and Freddie Mac, as GSEs, have to play in shaping industry implementation and responsibility for AI technology?⁷

A.3. Approximately 70 percent of mortgage loans in the U.S. are financed by Fannie Mae and Freddie Mac, and most mortgage lenders utilize their proprietary underwriting systems. As conservator of the Enterprises, the Federal Housing Finance Administration (FHFA) has issued guidance for the managing risks of AI/ML systems⁸ and has approved the use of AL and ML models for certain applications, e.g., fraud detection, property valuation, and loan servicing activities. However, according to a recent study on AI in

⁵H.R. 6580, To direct the Federal Trade Commission to require impact assessments of automated decision systems and augmented critical decision processes, and for other purpose, Feb. 3, 2022, <https://www.congress.gov/117/bills/hr6580/BILLS-117hr6580ih.pdf>.

⁶Mokander, J., Juneja, P., Watson, D.S., et al. (2022), "The U.S. Algorithmic Accountability Act of 2022 vs. The EU Artificial Intelligence Act: What Can They Learn From Each Other?" *Minds & Machines* 32, 751–758. <https://doi.org/10.1007/s11023-022-09612-y>

⁷Neal, Michael, Linna Zhu, Caitlin Young, Vanessa G. Perry, Matthew Pruitt (2023a), "Harnessing Artificial Intelligence for Equity in Mortgage Finance", Urban Institute, November 6, <https://www.urban.org/research/publication/harnessing-artificial-intelligence-equity-mortgage-finance>.

⁸Advisory Bulletin on Artificial Intelligence/Machine Learning Risk Management, Federal Housing Finance Agency, AB 2022-02, Feb. 10, 2022, <https://www.fhfa.gov/SupervisionRegulation/AdvisoryBulletins/Pages/Artificial-Intelligence-Machine-Learning-Risk-Management.aspx>.

the mortgage industry, neither the Enterprises, the CFPB, nor other regulatory entities have provided specific standards for downstream adoption and management of AI by lenders, seller/servicers of the Enterprises, or other intermediaries in the mortgage value chain. Because the underwriting standards of the Enterprises are the primary drivers of lending practices industrywide, and because of the significant legal and financial risks of noncompliance, lenders and other market participants have been reluctant to experiment with AI-based innovations.⁹ As a result of this uncertainty, these lenders may miss opportunities to leverage AI/ML to advance opportunities for currently underserved market segments, such as models to expand credit scoring or increase fairness in property valuation.

Because of the Federal Government’s close monitoring and oversight of the Enterprises, as well as their access to data, model development, and computing resources, they are in a better position than individual lenders to lead AI/ML innovation in the industry via pilot programs. For example, Freddie Mac has used AI/ML to incorporate rental and other payment data sources in their models, and partnering with Fintech companies has been one approach to expediting the development and testing of these tools.¹⁰

It is important to note that historically, changes to underwriting models that require FHFA approval have taken years to implement, failing to meet the regulatory demands posed by the rapid pace of AI/ML innovations. For example, it has taken more than 5 years for the Enterprises to update their credit score requirement despite compelling evidence that these changes would have a significant and immediate impact on homeownership access for marginalized communities.¹¹

It is also important to understand the influence of the Federal Housing Administration (FHA) on AI adoption in the mortgage industry. The FHA Technology Open to Approved Lenders (TOTAL) Scorecard has been developed using AI/ML and is used widely by mortgage lenders along with Freddie Mac and Fannie Mae’s underwriting systems. Relative to the conventional market dominated by the Enterprises, FHA insures a disproportionate share of mortgages for low- and moderate-income, first-time homebuyers and historically underrepresented racial groups. For example, in 2023, FHA’s percentage of mortgages to Black and Hispanic borrowers was 30.6 percent,¹² compared to 15.7 and 20.7 percent for Freddie Mac and Fannie Mae, respectively.¹³ FHA issues guidance for

⁹Neal, et al. (2023a).

¹⁰“Zest AI Joins Forces With Freddie Mac To Help Make Homeownership Possible for More Americans”, *PR Newswire*, Nov. 18, 2020, <https://www.prnewswire.com/news-releases/zest-ai-joins-forces-with-freddie-mac-to-help-make-homeownership-possible-for-more-americans-301176336.html>.

¹¹“Fact Sheet: Credit Score Models and Credit Report Requirements”, Federal Housing Finance Agency, Mar. 23, 2023, <https://www.fhfa.gov/Media/PublicAffairs/Pages/Fact-Sheet-FHFA-Announcement-on-Credit-Score-Models-March-2023.aspx>.

¹²“Fact Sheet: FHA’s Impact in 2023—Making Homeownership Possible for Hundreds of Thousands of Families”, U.S. Department of HUD, Nov. 15, <https://www.hud.gov/sites/dfiles/PA/documents/Fact-Sheet-MMI-Rollout.pdf>.

¹³“2023 Annual Housing Activities Report”, Freddie Mac, <https://www.freddiemac.com/about/pdf/2023-annual-housing-activities-report.pdf>; 2023 Annual Mortgage Report, Fannie Mae, <https://www.fanniemae.com/media/50736/display>.

lenders on how to implement the FHA’s AI-based systems but does not address lenders’ use of their own AI applications.¹⁴

The U.S. Department of Housing and Urban Development (HUD) also influences the adoption and oversight of AI in its role as regulator of the Enterprises’ fair lending activities along with FHFA. Significant, additional investments in technology infrastructure and expertise will be necessary for HUD and the Department of Justice to investigate violations or otherwise enforce the Fair Housing Act when AI and ML are involved. Currently, the agency relies heavily on third-party vendors for fair lending testing of these models.¹⁵

Q.4. Additionally, smaller, mission-oriented lenders such as Minority Depository Institutions and Community Development Financial Institutions tend to serve low-to-moderate income communities and minorities at higher proportions than other institutions. But we know, from one of your reports, that they have been slower to adopt AI than other financial institutions. Can you please describe this technology gap, and if it could further reinforce racial disparities in the mortgage industry and what should be done about it?

A.4. According to a recent study,¹⁶ larger lenders are more likely to adopt AI/ML systems due to greater capacity to make the necessary investments in technology and human capital, used towards in-house or third-party resources. Fintech companies, for example, can charge anywhere from several thousand to hundreds of thousands of dollars per year to develop models used for customer service, underwriting, or fraud detection, not including additional ongoing costs for maintenance and administration. Taken together, these costs present significant barriers to AI adoption for smaller and mission-oriented lenders, such as CDFIs. Anecdotal evidence also suggests that these smaller institutions rely more heavily on relationship-based approaches which are more labor-intensive, less quantifiable, and more difficult to replicate using algorithms designed for the mainstream mortgage market. However, one of the intended benefits of AI/ML is to incorporate nontraditional data sources in ways that can expand access to the mortgage market; another potential benefit is lower costs due to increased efficiency. Due to these technology gaps, smaller institutions could miss out on opportunities that AI presents to create mission-focused scale economies.

To ensure that smaller lenders have access to these technologies, regulatory agencies should partner with nonprofit housing organizations to provide technical assistance grants, similar to the Community Development Financial Institutions (CDFI) Program administered by the U.S. Department of the Treasury. This program provides grants to support investments in technology, capacity building, and new product development.

¹⁴Neal, Michael, Janneke Ratcliffe, and Matthew Pruitt (2023b), “AI Could Alter Mortgage Lending, but Government Leadership Is Needed”, *Urban Wire*, Urban Institute, November 6, <https://www.urban.org/urban-wire/ai-could-alter-mortgage-lending-government-leadership-needed>.

¹⁵Hoffberg, Adam, and Bill Reeder (2023), “The FHA TOTAL Mortgage Algorithm: Providing Machine-Learning Analysis for Two Decades”, PD&R Edge, U.S. Dept. of HUD, June 27, <https://www.huduser.gov/portal/pdredge/pdr-edge-pdrat50-062723.html>.

¹⁶Neal, et al. (2023a).

**RESPONSES TO WRITTEN QUESTIONS OF SENATOR WARNOCK
FROM VANESSA PERRY**

Q.1. The Fair Housing Act¹ and Equal Credit Opportunity Act² are landmark statutes to prevent discrimination in housing. Some Federal regulators have already used authorities under these statutes to engage in oversight of more advanced models that employ novel forms of data and more sophisticated statistical techniques.³

In your view, are Federal regulators using all available authorities under FHA and ECOA to engage in oversight of novel software or models that employ machine learning (ML) or artificial intelligence (AI)? Why or why not?

A.1. Federal regulators, such as the U.S. Department of Housing and Urban Development (HUD), the U.S. Department of Justice (DOJ), and the Consumer Financial Protection Bureau (CFPB), have leveraged their authority under the Fair Housing Act and the Equal Credit Opportunity Act (ECOA) to regulate the use of Artificial Intelligence and Machine Learning models (AI/ML) in the mortgage industry. For example, in 2019 Facebook settled a Fair Housing Act case filed by HUD. The case alleged that Facebook's AI-driven advertising platform enabled discriminatory practices, such as targeting or excluding market segments based on prohibited factors, such as race, gender, nationality, religion, familial status, and disability. This landmark case has led to heightened awareness, increased scrutiny, and regulatory oversight of advertising practices in the mortgage industry. In addition, social media and digital marketing companies have instituted stricter policies to ensure compliance with fair housing laws.

Additionally, in 2023, the CFPB, along with the DOJ, the Federal Trade Commission (FTC), and the Equal Employment Opportunity Commission (EEOC), issued a joint statement emphasizing their intentions to apply existing antidiscrimination laws and regulations to AI models, particularly when used for credit decisions and digital marketing of financial services.⁴ Later in 2023, the CFPB issued guidance reinforcing that existing legal requirements under ECOA for lenders to provide specific reasons for credit denials apply in the case of AI models.⁵ These examples underscore the potential for existing statutes to apply to emerging technological applications, such as AI/ML.

While these are important signals to the industry about best practices and legal requirements, these agencies can only take action after alleged violations have occurred and complaints have been filed. Given resource constraints including limited in-house expertise in evaluating AI/ML models, more efforts should be placed on ensuring that AI/ML models align with prevailing legal requirements and established risk management principles before

¹ See supra n. 1.

² See supra n. 2.

³ See Press Release, supra n. 3.

⁴ Khan, Linda (2023), Joint Statement on Enforcement Efforts Against Discrimination and Bias in Automated Systems, Federal Trade Commission, April 25, <https://www.ftc.gov/legal-library/browse/cases-proceedings/public-statements/joint-statement-enforcement-efforts-against-discrimination-bias-automated-systems>.

⁵ CFPB Issues Guidance on Credit Denials by Lenders Using Artificial Intelligence, CFPB, Sept. 13, <https://www.consumerfinance.gov/about-us/newsroom/cfpb-issues-guidance-on-credit-denials-by-lenders-using-artificial-intelligence/>.

implementation, such as the predeployment impact assessments proposed in the Algorithmic Accountability Act of 2022.⁶

Q.2. Are there any ways in which existing statutes, authorities, or resources fail to enable policymakers and regulators to engage in oversight of novel software or models using ML and AI? If so, what additional authorities are most helpful? Please provide specific examples.

A.2. Resource constraints, regulatory bureaucracy, and statutory restrictions explain why existing statutes, authorities, or resources limit the ability of policymakers and regulators to oversee novel software or models using ML and AI. For example, under the Fair Housing Act and ECOA, HUD, and the CFPB, respectively, have subpoena power and can make referrals to the DOJ, but the DOJ is only able to make information requests. States have subpoena power but must defer to Federal regulators for cases against national banks. Congress could increase the effectiveness of enforcement efforts by expanding the subpoena power of the DOJ and the States.

A Federal judge recently dismissed a case filed by the CFPB alleging that a mortgage company's communications would deter prospective Black mortgage applicants, ruling that the ECOA only applies to actual loan applicants, in contrast to the Fair Housing Act, which covers potential applicants.⁷ While this case did not relate to the use of AI/ML, it has implications for the ability to enforce AI/ML marketing practices. A legislative solution would be to align ECOA with the Fair Housing Act.

The Fair Housing Act's "disparate impact" rule prohibits the institution of lending policies, including the use of AI/ML models, that have an unjustified discriminatory effect based on protected characteristics. Similarly, ECOA prohibits the use of protected demographic characteristics in credit decisions. These regulations require lenders to legitimize actions that deny access to credit and housing opportunities but would be more impactful if expanded to stipulate an affirmative obligation by federally regulated entities to expand access to credit and housing opportunities.

RESPONSES TO WRITTEN QUESTIONS OF CHAIR SMITH FROM NICHOLAS SCHMIDT

Q.1. How do we build incentives into housing and regulatory systems so that we are maximizing fairness, accountability, and transparency as we try to set guidelines both to encourage innovation but also to make sure these systems are working for the public? Can and should this be handled legislatively? Dr. Perry underscored the importance of holding "AI accountable for expanding opportunity," and I welcome legislative and regulatory recommendations for how to best incentivize this.

A.1. It is essential to recognize that innovation and the principles of fairness, accountability, and transparency do not conflict. In fact,

⁶ H.R. 6580, <https://www.congress.gov/117/bills/hr6580/BILLS-117hr6580ih.pdf>.

⁷ Andreano, Jr., Richard (2023), "CFPB Suffers Significant Defeat in ECOA Lawsuit Against Townstone Mortgage", *Consumer Finance Monitor*, Feb. 6, <https://www.consumerfinancemonitor.com/2023/02/06/cfpb-suffers-significant-defeat-in-ecoa-lawsuit-against-townstone-mortgage/>.

these principles are the catalysts that drive sustainable and impactful innovation. Rather than hindering progress, these values guide the development of technologies that will ultimately lead to successful innovation, serve the public good, and ensure equitable outcomes. Fairness is an integral component of successful innovation. Innovations that exacerbate disparities based on race, gender, age, or other categories cannot be deemed successful. Implementing standards that mandate testing and validation of the fairness of models mitigates risk and promotes innovation that benefits all segments of society. The importance of accountability in fostering successful innovation becomes evident when we reflect on the causes and responses to the 2008–2009 housing crisis. The lack of accountability by individuals making investment decisions was a significant factor in the crisis. However, the resilience of the lending industry since then can be attributed, in part, to enhanced regulatory standards for model risk management introduced by the Federal Reserve and the Office of the Comptroller of the Currency in 2011. These standards have significantly improved the vetting, testing, and monitoring of algorithmic decision-making systems, and our economy’s success demonstrates that accountability is a cornerstone of robust and responsible innovation. Transparency, likewise, plays a pivotal role in promoting responsible innovation within the housing sector. AI is, without a doubt, overhyped. As a result, there has been a tendency to implement AI systems without proper oversight in many industries. By encouraging and requiring more transparency, regulators and legislators will push companies to consider whether an insufficiently tested system is worth the potential for reputational harm that might result from disclosure through transparency requirements. For example, a modeler might want to use an unnecessarily and overly complex algorithm to assess credit quality (I have seen unnecessarily complex algorithms put into production far too frequently in unregulated industries). However, suppose the company understands that it has transparency requirements and must be able to explain why the algorithm made a decision. If the algorithm is too complex, the company cannot reasonably explain the algorithm’s decisions. This may drive business and compliance leaders of the company to refuse to allow such a complex model to be used. This is beneficial because overly complex algorithms are often unfair and unsafe. Being encouraged to push back against the desire to “move fast and break things” will lead to a more considered and mature adoption of AI.

While fairness, accountability, and transparency are compatible with and necessary for successful innovation, it is crucial to acknowledge the potential for legislative efforts to inadvertently impede progress. Successful legislation should be principled rather than prescriptive, offering a framework that fosters innovation while ensuring it operates within ethical boundaries. This approach allows flexibility and adaptation, which are essential for nurturing innovation in a rapidly evolving technological landscape.

Q.2. You underscored the bevy of existing laws that apply to and cover most of the AI applications in housing. What are the legislative gaps you see in this space, and how do you recommend that Congress addresses them?

A.2. With exceptions related to the provision of transparent explanations, standards around proxy discrimination, and mandating that modelers search for less discriminatory alternative models (which I outline in my responses to Senator Warnock), large lending institutions are generally well-covered by existing laws. However, smaller companies, including third-party model providers, have generally not been subject to the rigorous standards against which large institutions are assessed in terms of model governance, model quality, fairness, and safety. While it is likely unreasonable to require that smaller companies meet the compliance standards of the top banks, they should be subject to more rigorous standards and expectations. Creating legislation that provides more funding for existing regulators to understand these issues and investigate companies that do not meet minimum standards would help ensure that these models, which have a significant effect on people's ability to rent apartments or buy homes, would be fairer and more effective.

**RESPONSES TO WRITTEN QUESTIONS OF SENATOR WARNOCK
FROM NICHOLAS SCHMIDT**

Q.1. As credit underwriting models increasingly employ newer and larger sets of data, I am concerned that historical discrimination in housing may lead to biased training data.

What is your current assessment of the quality of training data available for those building credit models?

A.1. The quality of training data used to construct credit models is often inconsistent. Data quality is generally high for individuals with sufficient historical credit history to be evaluated through traditional credit scoring models, such as FICO or Vantage. While this observation may diverge from narratives frequently encountered in the press, credit bureaus maintain robust standards for accuracy and offer mechanisms for data dispute and correction. This achievement is, in part, a testament to effective legislative frameworks that mandate transparency and consumer rights to data access.

Conversely, the emergent forms of data increasingly incorporated into credit assessments—particularly for those lacking a traditional credit score—exhibit varying degrees of quality. The issues stem from several factors, including the nascent stage of integrating such data into credit algorithms, the relative absence of mandates for data disclosure and correction, and the usage of complex algorithms that may obfuscate the impact of erroneous data on consumer credit evaluations. The latter of these issues may require significant efforts to uncover and address data inaccuracies, an effort that is often not sufficiently addressed by some model builders or users.

Further, the issue of data incompleteness is more pronounced in these newer data streams. The dependency on third-party data providers that may not consistently capture information across all consumer segments exacerbates this challenge. Notably, discrepancies in data completeness often manifest along racial lines, underscoring the need for vigilant review and monitoring to prevent systemic biases.

Despite these challenges, it is also crucial to recognize the potential benefits of incorporating alternative data sources. For individuals with limited or no traditional credit history, access to such data enables financial institutions to extend credit and housing opportunities that would otherwise be inaccessible. This can significantly enhance financial equity and provide opportunities to underserved populations. Therefore, while maintaining vigilance and concern about data quality and its associated fairness, it is essential to consider that its use often promotes financial inclusion.

Q.2. What standards are currently used by industry to assess an AI model’s quality in terms of its data inputs as well as other model characteristics (e.g., accuracy, bias and discrimination, variance, performance on sparse data, etc.)?

A.2. Before machine learning (ML) and AI became widespread, most major financial services companies had robust processes to assess the fairness and reasonableness of decision making (whether subjective or algorithmic). As algorithmic decisioning has become more common, most companies have successfully applied these processes to their machine learning and AI algorithms. In particular, the legal framework of disparate treatment, disparate impact, and proxy discrimination is the most commonly used framework to assess fairness. This framework assesses both the data used to build the algorithm and the fairness of the algorithm’s output.

This framework has a rich history. While originating in the employment context, it has been extended to housing, lending, and other industries, including insurance and technology. Its advantages include its comprehensive yet flexible nature. It also addresses how discrimination has manifested in our society, and it provides an avenue to encourage inclusion without being unnecessarily burdensome. Finally, it is well-designed to be applied to questions of fairness and inclusion in algorithmic decisioning by machine learning and AI.

AI, ML, and the data used for these algorithms present some challenges to applying this framework, particularly in assessing proxy discrimination and disparate impact. Institutions would benefit from having more clarity around what types of data they can use and what they are expected to do to minimize any disparities in their models.

Proxy discrimination involves making a lending decision based at least partially on a factor so closely related to membership in a protected class that it is essentially just a stand-in for that class. The classic example of proxy discrimination is redlining, which is the refusal to lend in certain neighborhoods that were almost always heavily minority neighborhoods. Lenders justified this by claiming that these neighborhoods were riskier to lend in than neighborhoods that were majority White. This was a racist policy, and lenders go to great lengths to avoid obvious proxies like this.

However, in machine learning and AI, modelers use far more data than traditional credit bureau attributes, which are much less likely to be proxies for a protected class. When building these models, modelers may want to incorporate information about such varied things as marketing preferences, shopping preferences, and pat-

terns of speech or writing that may cross the line and become proxies for a protected class. While it may seem clear that these variables should be avoided if they might proxy for a protected class, the question becomes more difficult for at least two reasons. First, using these variables may increase financial inclusion for people without sufficient traditional credit history—even for members of the protected class proxied by the variable. The question may be whether to deny a loan to everyone without sufficient credit history or to use these variables, accept some people without sufficient credit history, but accept relatively fewer members of a protected class than truly should be accepted if better data existed. The answer to this question is ultimately a public policy question that individual companies in the housing industry cannot answer on their own without encountering significant regulatory and legal risks. Thus, having clearer regulations around the use of potentially problematic variables would be beneficial.

The second reason why this is a difficult question is that there is not a clear dividing line between disparate impact and proxy discrimination. While we can make some distinctions on certain variables, others are less obvious. For example, a variable that indicates where a person shops for clothing is almost certainly a proxy for sex. However, a variable that indicates the percent of disposable income a person spends on clothing may not be a proxy. Suppose the second variable is still correlated with sex, but it is also a robust predictor of whatever outcome is being modeled. In that case, it might be considered a source of disparate impact, not proxy discrimination. As such, it is not necessarily illegal. If the first variable does represent a proxy risk and the second does not, then it is obvious that there is some line between the two where a variable moves from being a cause of disparate impact to being a cause of proxy discrimination. The difficulty is that there is no clear guidance on where that line is. Industry participants would benefit from regulations that set out clearer guidance on what constitutes proxy discrimination.

Machine learning and AI also present new issues related to disparate impact. While disparate treatment represents the explicit inclusion of protected class status into a decision, disparate impact is perhaps a subtler form of discrimination. In a disparate impact analysis of a model, we start with a very simple test. In the case of a yes/no outcome, we test whether a protected group receives the favorable outcome less frequently than a member of some reference group. For example, we test whether Black or African American applicants receive loan offers less frequently than non-Hispanic White applicants. In the case of a continuous outcome, such as interest rates paid for a loan, we assess whether the protected group has, on average, received worse terms than the reference group. Using interest rates as an example, we might test whether women pay higher APRs than men.

Importantly, if we find evidence of unfavorable differences, we do not necessarily assume that the algorithm is illegally discriminatory. Instead, we move to the next step of assessing whether the algorithm's use has sufficient business justification and that none of the variables in the model represent disparate treatment or proxy discrimination. For example, a well-built credit underwriting

model that has passed rigorous model governance standards and uses traditional credit bureau data would almost certainly pass this hurdle.

If the model passes this second stage, we then test whether a similar process can still meet the institution’s legitimate business needs but have a smaller negative impact on protected class members. Searching for such an alternative model is called a “search for less discriminatory alternative models” or “LDA search.”

Machine learning and AI have significantly extended the ability to effectively search for LDA models. In traditional models, the approach was to try adding or dropping a handful of variables and see whether that might decrease disparities while maintaining model quality. While sometimes effective, this approach often did not result in significant changes. In the new age of ML and AI, we have developed new techniques that have greatly expanded the ability to search for LDA models. Adopting such techniques can significantly increase equitable and fair outcomes in housing.

While we have seen significant positive results from implementing these techniques, there is a lack of regulatory clarity around whether this is required, what techniques are allowed, and what standards will meet regulatory expectations. Legislation that clearly outlines that searching for less discriminatory models is a requirement and allows regulators to set the guidelines for what is acceptable will encourage industry participants to adopt these approaches. This could drive significant and positive change by making housing outcomes fairer and more equitable.

Q.3. In your written testimony, you stated, “Many . . . AI and ML algorithms are described as ‘black-box’ processes . . . [W]hile we understand what data went into the algorithm, and we can see what output of the algorithm is, it is very difficult to understand how or why the algorithm made the prediction it made.”¹ You also note that “a lender might not be able to provide a borrower with a reliable explanation for why the home they wanted to buy received a low appraisal.”²

In your experience, how commonly are “black-box” algorithms used in housing?

A.3. These algorithms are now quite commonly used in housing. They are used throughout the industry, most notably for marketing, appraisals, credit underwriting, loan pricing, evaluating renters, and servicing loans.

Q.4. How can regulators better ensure that lenders and other institutions use models capable of providing transparent explanations for their decisions?

A.4. Regulators may be able to encourage lenders to use appropriately transparent algorithms by assessing and judging them based on the quality of their models’ explanations. If a lender knows it must give regulators evidence that its algorithms provide reasonably accurate explanations, it will be less likely to use an unnecessarily complex algorithm.

¹Testimony of Nicholas Schmidt, U.S. Senate Committee on Banking, Housing, and Urban Affairs, Subcommittee on Housing, Transportation, and Community Development (Jan. 31, 2024), <https://www.banking.senate.gov/imo/media/doc/schmidt-testimony-1-31-24.pdf> at 8.

²Id.

The difficulty is that, even with relatively simple algorithms, all explanation methodologies have shortcomings; even the most robust approach will not always provide the best explanation for why an algorithm made a decision. This lack of perfection should not be a reason to reject AI and machine learning because well-built machine learning or AI algorithms significantly benefit consumers and industry participants.

Regulators have not been clear about what constitutes a sufficiently robust explanation methodology. While this lack of clarity has been reasonable given the pace of technological change, we are now at a point where better guidance would benefit consumers and the industry. However, rather than prescribing a particular methodology, effective regulation would put forth methodologies for testing and evaluating the robustness of explanations. A company could then be evaluated on how well it tested its methodology and whether that methodology passed the tests it constructed. The benefits of this approach include that it would allow continued innovation by not prescribing an approach that may quickly find itself out of date while also providing companies with a clearer standard against which they can judge themselves.

ADDITIONAL MATERIAL SUPPLIED FOR THE RECORD



Jim Nussle
 President & CEO
 202-508-6745
 jnussle@americascrreditunions.org

99 M Street SE
 Suite 300
 Washington, DC 20003

January 31, 2024

The Honorable Tina Smith
 Chairwoman
 Committee on Banking, Housing
 and Urban Affairs
 Subcommittee on Housing, Transportation,
 and Community Development
 United States Senate
 Washington, DC 20510

The Honorable Cynthia Lummis
 Ranking Member
 Committee on Banking, Housing
 and Urban Affairs
 Subcommittee on Housing, Transportation,
 and Community Development
 United States Senate
 Washington, DC 20510

Re: Today's Hearing: "Artificial Intelligence and Housing: Exploring Promise and Peril"

Dear Chairwoman Smith and Ranking Member Lummis:

On behalf of America's Credit Unions, I am writing regarding the Subcommittee's hearing entitled, "Artificial Intelligence and Housing: Exploring Promise and Peril." America's Credit Unions is the voice of consumers' best option for financial services: credit unions. We advocate for policies that allow the industry to effectively meet the needs of their nearly 140 million members nationwide.

America's Credit Unions appreciates the Subcommittee holding this hearing on the potential transformative properties of artificial intelligence (AI) and how it impacts the housing ecosystem in this country. Credit unions are engaging in collaborative innovation by working with credit union service organizations (CUSOs) and other financial technology vendors to leverage the power of AI for many aspects of their daily operations. From member service interactions to AI chat-bots to mortgage origination – credit unions are seeing firsthand how AI is increasing staff efficiency, automating previously laborious tasks, reducing paperwork, and expediting loan decision making processes. These efficiencies have tangible benefits to the communities credit unions serve and have shown promise in helping low-and moderate-income families get access to affordable credit. It's imperative that credit unions continue to responsibly utilize AI to maintain a competitive member-focused advantage against Wall Street megabanks.

As the financial services industry continues to rapidly innovate and further integrate AI tools into their core systems, credit unions have a variety of concerns about how their regulators approach AI and the potential risks AI poses to systems they use to serve their members.

America's Credit Unions Supports Efforts to Reduce Barriers to Innovation

Financial sector regulators should tailor future actions related to AI in a way that distinguishes between the use of AI technology by regulated versus unregulated institutions. With respect to highly regulated financial institutions, such as credit unions, an appropriate regulatory

January 31, 2024
Page 2 of 3

framework should recognize the need for less prescriptive intervention and greater accommodation of innovation through pilot programs, no-action letters, waivers, and elimination of outdated rules. In an environment where non-bank fintech companies may be enjoying less rigorous supervisory oversight than traditional financial institutions, regulators should be exploring frameworks that make innovation accessible not just to the largest and most sophisticated entities, but also to smaller, community-based institutions. The need to establish a fair playing field cannot be overstated. Additionally, an overly complex or intrusive supervisory framework for assessing the innerworkings of AI algorithms would likely deter credit unions from investing in these technologies and frustrate efforts to partner with CUSOs and other technology providers.

Certain third-party providers of AI services may not grant customers access to the proprietary code that underlies machine learning algorithms or models. In these circumstances, demanding that end-users attest to the specific operational parameters of an AI product would not be feasible and would likely inhibit the use of such technology by credit unions unable to develop these products in-house. While AI explainability may be a key consideration for regulators, innovation cannot reasonably take place in an environment where every line of code is scrutinized, and financial institutions must prove a negative: that a model is incapable of error. A more reasonable approach would be to encourage regulators to understand how AI risks can be addressed through the application of existing laws and regulations, which are not limited to any particular mode of decision making or technology. As discussed below, the evaluation of AI outputs using existing compliance procedures can identify and prevent consumer harm before it occurs.

Ensuring Financial Inclusion and Fairness in Lending with AI Optimization

Credit unions were created to offer provident credit to all members of their communities and this organizing principle helps to explain the prevalence of robust relationship lending models across the industry. As cooperatives that are directly accountable to their member-owners, credit unions are focused on developing long-lasting, trusted relationships—an interest that is best served by adhering to core principles of equality and fairness.

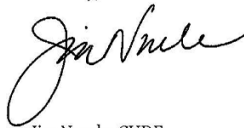
Credit unions follow existing regulations and guidance implementing the Equal Credit Opportunity Act (ECOA) and other antidiscrimination laws and have a track record of exceptional fair lending compliance. Most credit unions engage in self-tests or self-evaluations as part of their ongoing monitoring of fair-lending risks. While self-evaluations can vary in terms of their scope and sophistication based on a credit union's risk profile, they generally encompass review of denied applications; comparisons of loan files; analysis of HMDA data, and review of lending policy exceptions. Self-tests can be similarly varied and encompass a variety of analytical techniques (e.g., surveys, use of test applicants, review of credit transaction records). Both types of testing could function as post-hoc methods for evaluating the results of AI-driven lending decisions. These existing methods should be regarded as effective for detecting and addressing fair lending risks. Deconstructing the entirety of an AI algorithm to address explain-ability or overfitting risks would be costly and less productive for examination purposes.

January 31, 2024
Page 3 of 3

AI and machine learning can improve access to credit and access to housing as long as they are deployed responsibly. They can affect who is approved for credit and who sees advertisements for housing and mortgage loans. In short, they can help reduce the racial homeownership gap when the technology works as intended. Because credit unions support legislative and regulatory initiatives that promote the availability of credit to all creditworthy applicants, they recognize that responsible use of AI mandates compliance with the ECOA, Fair Housing Act, and existing applicable law. That said, credit unions and America's Credit Unions expect that legislators and regulators will weigh the regulatory burden on credit unions against the benefit to consumers when implementing new laws and regulations that govern the use of AI in housing. To do otherwise may chill technological progress and the benefits that may accrue to consumers through the responsible use of AI.

America's Credit Unions is committed to working with the appropriate financial regulatory agencies to ensure credit unions are able to innovate and use AI to help their members achieve financial well-being. On behalf of America's Credit Unions and the 140 million credit union members, thank you for holding this important hearing and considering our views on the subject.

Sincerely,



Jim Nussle, CUDE
President & CEO

cc: Members of the Subcommittee on Housing, Transportation, and Community
Development



Concept Paper

Algorithms for All: Can AI in the Mortgage Market Expand Access to Homeownership?

Vanessa G. Perry, Kirsten Martin and Ann Schnare

Special Issue
Standards and Ethics in AI

Edited by
Dr. Pablo Rivas and Dr. Gissella Bejarano



<https://doi.org/10.3390/ai4040045>



Concept Paper

Algorithms for All: Can AI in the Mortgage Market Expand Access to Homeownership?

Vanessa G. Perry ^{1,*}, Kirsten Martin ² and Ann Schnare ³

¹ Department of Strategic Management and Public Policy, School of Business, The George Washington University, Washington, DC 20052, USA

² William P. and Hazel B. White Center, Department of IT, Analytics, and Operations, Mendoza College of Business, University of Notre Dame, Notre Dame, IN 46556, USA; kmarti33@nd.edu

³ AB Schnare Associates, Washington, DC 20007, USA; annschnare@aol.com

* Correspondence: vperry@gwu.edu

Abstract: Artificial intelligence (AI) is transforming the mortgage market at every stage of the value chain. In this paper, we examine the potential for the mortgage industry to leverage AI to overcome the historical and systemic barriers to homeownership for members of Black, Brown, and lower-income communities. We begin by proposing societal, ethical, legal, and practical criteria that should be considered in the development and implementation of AI models. Based on this framework, we discuss the applications of AI that are transforming the mortgage market, including digital marketing, the inclusion of non-traditional “big data” in credit scoring algorithms, AI property valuation, and loan underwriting models. We conclude that although the current AI models may reflect the same biases that have existed historically in the mortgage market, opportunities exist for proactive, responsible AI model development designed to remove the systemic barriers to mortgage credit access.

Keywords: mortgage market; BIPOC homeownership; ethics of AI; algorithmic bias



Citation: Perry, V.G.; Martin, K.; Schnare, A. Algorithms for All: Can AI in the Mortgage Market Expand Access to Homeownership? *AI* 2023, 4, 888–903. <https://doi.org/10.3390/ai4040045>

Academic Editors: Pablo Rivas and Gissella Bejarano

Received: 8 August 2023

Revised: 18 September 2023

Accepted: 27 September 2023

Published: 11 October 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

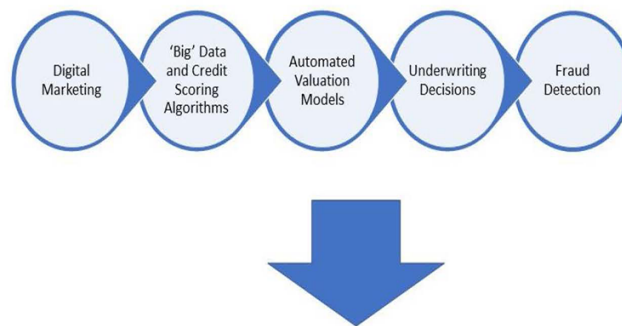
1. Introduction

In the mortgage market, digitalization, i.e., “the use of digital technologies to change a business model and provide new revenue and value-producing opportunities”, [1] is already embedded and expanding throughout the mortgage value chain. For example, mortgage lenders use sophisticated digital marketing techniques to target prospective borrowers and artificially intelligent bots to communicate with customers. Credit scoring companies are using machine learning processes to evaluate credit risk. Property valuation algorithms integrate copious amounts of data on land titles, sales, market trends, and aspects of local infrastructure to produce digital appraisals. Digitized processes are replacing manual, paper-based workflows used for loan servicing and loss mitigation. Industry participants are experimenting with blockchain implementations to manage the origination process. Despite several potential benefits of digital transformation, including increased efficiency and accuracy and lower costs, institutions and regulators are finding it difficult to keep up with the rate of technological innovation. At the same time, evidence regarding the impact of digitalization processes, namely artificial intelligence (AI), on opportunities to expand mortgage credit to underserved communities, including lower-income and minority households, is lacking and inconclusive at best. In this article, we offer a framework to examine the effectiveness of AI applications in the mortgage industry.

The purpose of this paper is to examine how the mortgage industry is leveraging AI to help overcome the historical and systemic barriers to homeownership for members of Black, Brown, and lower-income communities. As shown in Figure 1, we examine five areas where AI is transforming the mortgage market, including digital marketing; the inclusion of non-traditional “big data” in credit scoring; the use of artificial intelligence;

and machine learning in automated property valuation, underwriting, and fraud detection models. Building on prior research, we examine evidence of the potential of AI to transform market systems and outcomes. Based on this overview, we then describe how proactive, responsible technological transformation can be used to help overcome the systemic barriers to mortgage credit for historically underserved households.

AI in Housing Finance



Increased access to mortgage credit ?

Figure 1. AI is embedded and expanding throughout the mortgage value chain.

2. AI and ML in the Mortgage Industry

Artificial intelligence (AI) is a technological advancement whereby a computer or computerized actor (e.g., a robot) mimics human decision processes. Traditionally, humans program computers to perform specific computational or predictive functions, and then programmers update and improve these programs [2]. AI systems perform complex tasks in ways that mimic how humans solve problems. Machine learning (ML) is a form of AI in which the computer program optimizes its performance based on information gathered during previous tasks [3]. AI and ML are important digital transformation tools because of their ability to analyze much larger amounts of data and discover complex relationships that transcend traditional statistical assumptions and analyses. These tools have been increasingly applied in the private and public sectors [4]. Complex, multivariate algorithms have been in place for mortgage underwriting and pricing for more than two decades, and AI and ML are being used to enhance these models. AI and ML techniques have also been applied to marketing, customer relationship management, fraud detection, and loan servicing activities [4].

One supposed advantage of AI models is that they are not subject to human biases and errors; as such, they are viewed as possibly producing more accurate, consistent, and efficient decisions. Depending on how AI models are designed and developed, these enhanced capabilities could potentially expand access to credit for groups currently underserved by extant credit systems, particularly Black, Hispanic, and low-income consumers. However,

it is unclear how well these models can adapt to the changes in the market or the extent to which they would magnify the effects of past discrimination [5].

Since these models rely on historical data, critics in the academic and policy communities have raised concerns about the potential for these models to perpetuate historical discrimination and inequality [6,7]. The complexity of AI and ML tools makes it difficult but not impossible for non-developers to scrutinize and monitor their inputs [8,9].

In subsequent sections, we provide an overview of key technological and analytical advancements in the mortgage market and their potential for expanding access to credit for lower-income and minority households. We begin by establishing the criteria for evaluating the extent to which these tools and processes serve these goals. Drawing on the S.C.A.L.E. framework developed by [10], we propose five factors summarizing the societal, ethical, legal, and practical issues that should be considered in the development and implementation of AI.

3. A Framework for Evaluating the Impact of AI on Access to Homeownership for Underserved Communities

The following framework provides a parsimonious framework for evaluating the fairness and equity of AI and other technologies. While concerns about the sources of bias are not new, much of the discussion in the industry is focused on the legal risks and challenges with monitoring and oversight. These criteria are intended to summarize and simplify the key issues that responsible practitioners and policymakers should be prepared to address. In particular, this tool is forward-looking and incorporates factors that should shape the design and implementation, as well as the post-hoc evaluation of the impact of AI models in this sector.

- **Societal values.** A digitalized tool or process should be considered from the perspective of similar decisions, the larger context, and historical factors, and should align with the prevailing legal and ethical paradigms [11]. Recent political and social priorities in the U.S. have focused on racial equity and social justice, and the Biden Administration has directed regulatory agencies to increase fair access to homeownership. According to Kroll [8], credit scoring companies must “consider the context and impacts of their credit system and in particular to consider what outcomes are desired, how they might be reached, and how the deployment of a new system or changes to an existing system will alter the world.” New tools could be used to implement fair machine learning (FML) by deploying statistical algorithms to identify and correct unjust or biased outcomes [12].
- **Contextual integrity.** The appropriateness of a technological tool depends on whether it conforms to contextual norms [13,14]. Regardless of its accuracy, a particular tool must be appropriate for the mortgage lending or housing domain. Walzer [15] described “spheres of justice” to underscore the importance of context in evaluating the fairness of outcomes by arguing that someone who excels in one sphere (e.g., education) should not be granted advantages without merit in another sphere (e.g., mortgage loan access). Certain social media advertising tactics, while appropriate for less consequential product categories, may result in unfair informational asymmetries in the mortgage lending context.
- **Accuracy.** It is also important to evaluate the extent to which a tool is reliable, error-free, and widely available across all major demographic and economic groups and macroeconomic conditions. One advantage of AI is rapid, systematic, and consistent data collection and modeling. However, inaccuracies can result when certain types of data are systematically omitted or biases are built into algorithms. For example, in property valuation models, due to the varying assumptions about comparable property selection and the historical racial disparities in property values, can models produce the “accurate” measurements necessary to predict risk? What types of errors are acceptable? Accuracy also refers to the absence of bias [5].

1. *Representation bias* occurs when the sample upon which a model is based differs significantly from the characteristics of the population to which the model will be applied. For example, evidence suggests that the effect of credit scores on the likelihood of mortgage default differs for members of historically disadvantaged minority groups [16].
 2. *Historical bias* occurs when the data are accurate and correctly sampled, but also capture disparities due to past racism and discrimination. For example, Black and Hispanic borrowers pay higher rates and fees on average, and were more likely to have received subprime loans, faced foreclosure, or sustained significant equity losses during the 2008 global financial crisis.
 3. *Omitted variable bias* occurs when a model fails to include a factor that has a significant effect on the outcome. For example, it is illegal to include race or ethnicity as a factor in underwriting models. However, certain variables and combinations of variables are effective proxies for race/ethnicity. The omission of race as an explanatory variable can obfuscate the interpretation of these proxy variables and significantly underestimate the effects of racial/ethnic differences.
 4. *Selection bias* occurs when certain categories of people or transactions are systematically excluded from the data upon which a model is based. For example, households with low or missing credit scores would likely be underrepresented in the samples of prospective borrowers.
 5. *Aggregation bias* occurs when the characteristics of certain categories of people or transactions are erroneously applied to individual cases. One example is proxy methods, which rely on geographic and/or surname-based information to estimate the probability that a household belongs to a particular race and ethnicity when this information is not reported [17].
 6. *Measurement bias* occurs when a variable is systematically inaccurate, missing, or inconsistently measured. For example, a recent study found that credit scores for minority and low-income applicants were less predictive in mortgage default models due to the variations in underlying credit files [18].
- **Legality.** It is also important to assess whether adopting an AI application will have a negative and disparate impact on protected classes. The disparate impact standard prohibits any practice, including the use of a statistical algorithm, that has a negative, disparate impact on a particular racial/ethnic group when implemented. If a disparate impact occurs, the lender must provide a legitimate business justification and be able to rule out any less discriminatory alternative. The data and algorithms used for AI credit scoring, mortgage underwriting, and property valuation may run afoul of this standard.
 - **Expanded opportunity.** An AI solution should also significantly increase access to credit in addition to cost, efficiency, or risk assessment benefits. Whereas digitalization and AI have facilitated access to credit scores for previously unscorable or “credit invisible” households, it is unclear whether they increase financing opportunities for a larger group of consumers with poor credit histories [19].

It is important to note that while multiple criteria are likely to affect a particular AI application, these criteria are not mutually exclusive and are unweighted in terms of importance. Presumably, the relative importance of these factors would depend on the outcome priorities of the organization or stakeholder.

Table 1 provides examples of how the SCALE criteria apply to AI processes in mortgage lending and how these may affect access to mortgage credit to support minority homeownership.

Table 1. SCALE framework [10] and AI applications in the mortgage market.

Criteria	Digitalized Tool/Process	Impact on Minority Homeownership
Societal values	AI/ML use of GPS location	These data violate privacy norms/magnify income and wealth disparities that have resulted from historical racism and discrimination.
Contextual integrity	Targeted digital advertising that filters content based on demographic or psychographic profiles	Although these tactics work well in the context of apparel or automobiles, digital advertising may be less appropriate for mortgage lending.
Accuracy	Property valuation algorithms	On average, Black and Hispanic borrowers pay higher rates and fees and are more likely to have received high-cost subprime loans, faced foreclosure, or sustained significant equity losses during the 2008 crisis. Models based on “comparable” home values may unfairly penalize minority communities.
Legality	AI/ML mortgage underwriting algorithms	AI models may have negative, disparate impacts on certain racial/ethnic groups; due to model complexity, sources of bias may be difficult to detect.
Expanded opportunity	AI/ML using non-financial data in credit scoring algorithms	Expanded data used for credit scoring may reduce the population of unscorable households by increasing the number of households with high-risk (i.e., low) credit scores.

SCALE is not the only framework designed to evaluate the impact of AI models. For example, the AI Act, which is currently under legislative review in the European Union, designates applications with significant socioeconomic implications as high-risk, thereby requiring these AI systems to meet the standards of data quality, accuracy, robustness, and non-discrimination. This ACT also specifies the standards for documentation, transparency, and human oversight [20]. While the SCALE criteria are consistent with the EU approach, the SCALE criteria assume that AI models are far too complex, dynamic, and context-specific for substantive regulatory oversight. In addition, SCALE attempts to go beyond mere regulatory compliance with anti-discrimination laws by stipulating that these models align with societal values and expand opportunity.

The National Fair Housing Alliance (NFHA), a major civil rights and consumer advocacy organization, has put forth the Purpose, Process, and Monitoring (PPM) framework, which is focused on the stages in algorithmic model development and monitoring. The authors of this framework address the key considerations necessary for equity audits based on “fairness, accountability, transparency, explainability, and interpretability” [21]. The SCALE framework also aligns with the PPM framework and could be applied throughout the PPM process. The SCALE framework has an additional advantage in that it encompasses aspirations intended to ensure fairness while proactively advancing equity.

In the following sections, we apply the SCALE factors to understand the effects of “big” data, artificial intelligence, and machine learning on minority households in the mortgage market. We conclude with a discussion of the implications for ethical and socially responsible AI and opportunities to alleviate the existing barriers to mortgage access.

4. SCALE Criteria and AI-Driven Marketing

AI is one of several technological advancements in the mortgage industry that has been leveraged to increase access to mortgage credit for underserved consumers and to do so at lower costs and increased efficiency. However, recent research in this area has highlighted certain persistent disparities in financial market access despite the introduction of AI.

Evidence suggests that Black, Hispanic, and lower-income households have lower rates of online financial service usage. Based on the analyses of the 2019 Survey of Consumer Finances [22], Black and Hispanic households were less willing to engage in online banking

transactions and were less reliant on online information for borrowing decisions relative to White and other households. Black and Hispanic households were also more likely to have been denied credit or feared being denied credit in the past 5 years. Those who had been turned down or feared being turned down were significantly less likely to access online financial services regardless of race or ethnicity; this relationship was significantly more acute for Black and Hispanic households. According to Fannie Mae [23], the use of digital mortgage services increased significantly during the pandemic, but less so among certain groups of homebuyers. Higher-income, Asian, and Black recent homebuyers indicated a slightly higher preference for online mortgage-related activities, while lower-income and Hispanic consumers showed a stronger preference for in-person or telephone interactions [24,25]. Thus, it appears that some of the same racial and ethnic indicators of a digital divide in financial services documented before the pandemic persist today.

The term “fintech” has been defined as “technology innovations used to support or enable banking or financial services” [26,27], such as smartphone applications, wi-fi, online and mobile banking, electronic payment transactions, direct deposits, as well as transactions on peer-to-peer platform, and access to blockchain and cryptocurrencies. Friedline and Chen [28] noted that the proliferation of fintech has coincided with a decline in banking activities at brick-and-mortar institutions, and that “these trends have the potential to replicate and reinforce redlining by amplifying the existing racialized geography of financial services and exacerbating consumers’ marginalization from the financial marketplace”. They found that fintech rates among high-poverty communities are generally low, and they are even lower in areas with larger shares of Black, Latinx, and American Indian/Alaskan Native populations. Controlling for high-speed internet access, smartphone ownership, and checking account ownership, fintech usage is higher in areas with Hispanic and Asian residents; this is not the case in high-poverty areas with higher proportions of Black residents. Fintech companies, as new entrants with non-traditional business models, are more likely to introduce AI systems because they have been less subject to regulatory scrutiny.

Hauptert [29] found small yet significant racial disparities in loan approvals between similarly qualified White and non-White applicants and fewer disparities in approvals from fintech lenders versus traditional lenders. However, relative to similarly qualified White applicants, non-White applicants are more likely to receive subprime terms from both types of lenders, and the disparities in subprime lending between Black and White applicants are greater among fintech lenders than traditional lenders. The author thus recommended more careful regulation of fintech lending.

In terms of the SCALE framework, these fintech services have the potential to expand opportunities for minority homeownership, and many of these companies are innovators in AI. As it stands, however, they have had limited impact due to the significant racial and ethnic gaps in access and adoption.

Anecdotal evidence and recent legal activity suggest that AI used for targeted digital advertising practices may contribute to a less inclusive informational environment for members of traditionally underserved groups. AI marketing applications are used to harvest customer information to identify existing preferences and to recommend new product and service alternatives [30,31]. In doing so, these tactics, such as psychological targeting [30], customer prioritization based on income or profitability [32], and the targeting of vulnerable groups, can result in discrimination based on gender, age, and racial disparities. These practices have also raised concerns about the algorithms designed to optimize user acceptance in the context of social media [33]. Evans and Miller [34] argued that digital marketing techniques based on AI and machine learning (AI/ML) can increase the incidence of bias and consumer exploitation due to a lack of transparency in how they identify potential customers. Another concern related to targeted digital advertising by mortgage lenders is that advertisements may steer consumers toward particular products [34]. Enabled by AI, the customization of product and service offerings can have the unintended consequence of limiting access to information and opportunity [35], particularly in a category of profound

socio-economic significance, such as housing. There is also evidence that AI systems, because they are based on current and historical data, can reinforce prejudices, stereotypes, and historical inequities. These strategies could easily evade regulatory oversight.

These techniques also raise concerns about data privacy and contextual integrity [13,36,37]. Prior research has found that consumers can feel exploited by unauthorized uses of their data and when their data is used to categorize them in an inaccurate or biased manner [38].

Specifically, digital marketers purchase the data from third-party vendors that track users and their browsing behaviors across websites. Lenders also rely on third-party lead generators who provide lists of potential customers based on the data collected from website users who have shown interest in a particular product or category, e.g., people searching for homes or real estate agents. Additionally, lenders' digital marketing teams apply algorithms using the data extracted from various sources to estimate "e-scores" used to predict future usage behavior. Each of these techniques could exclude certain groups of borrowers from the market—particularly those who are currently underrepresented [34].

Recent cases against Facebook for Fair Housing Act (FHA) violations focused on ads for housing, but they also apply to ads for mortgages [39]. Cases filed by the National Fair Housing Alliance, other civil rights groups, and HUD found that Facebook enabled housing advertisers to screen viewers based on protected characteristics, such as race, sex, and disability, and to exclude parents, foreign-born individuals, and those seeking accessible units. In response, Facebook created a separate advertising platform that allows users to view all housing ads. The company also agreed to require advertisers to certify compliance with fair housing laws [40]. This example has prompted mortgage lenders to assess the fair lending risk in their AI marketing strategies and to carefully examine the criteria used to exclude groups based on prohibited characteristics [41].

In another example, the DOJ and CFPB settled a suit against Trustmark Bank in 2021 for using a digital marketing strategy designed for businesses in majority-White neighborhoods to generate mortgage business from majority-Black and Hispanic neighborhoods in the Memphis area [42]. The legal implications of AI targeting practices by mortgage lenders under the FHA and the Equal Credit Opportunity Act [43] raise important questions about fair access to information about mortgage loans.

That said, AI marketing tools are essential for reaching consumers in today's marketplace. However, these practices may exacerbate information gaps and steering activities, reduce competition, and further the "dual" mortgage market in which minority homebuyers pay more for mortgage credit [44].

Based on the SCALE framework criteria, targeted advertisements based on demographic categories or correlated attributes may not align with societal priorities aimed at increasing racial equity and inclusion, and it is unclear whether these practices contribute to expanded opportunity.

5. SCALE Criteria and AI Credit Scoring and Underwriting Algorithms

Policymakers and credit experts have touted the potential for the inclusion of alternative data sources to expand access to credit scores (which are necessary to access the mortgage market) for those who currently have sparse or missing credit files [45]. Current credit scoring models rely exclusively on the timeliness of past payments on consumer credit lines, i.e., credit cards, car loans, student loans, mortgages, and other consumer loans. Proponents believe that using alternative data such as rental payments, utility payments, and digital transactions in credit scoring models will expand opportunities to consumers who are currently "credit invisible" or unscorable [46,47]. One study estimated that the inclusion of telecommunications and utility payment data in traditional scoring models would increase acceptance rates by about 10 percent for the overall population, and by more than 20 percent for Black and Latinx individuals and consumers making less than \$20,000 a year [48]. Another analysis showed that rent and utility payments had a posi-

tive impact on consumers' access to credit, although the opposite was true for remittance payments [47,49].

However, critics suggest that UTR payment history data could inadvertently increase financial challenges for families who are struggling to recover from the pandemic downturn or seasonal fluctuations in energy costs. In addition, there is evidence that Black, Hispanic, and low-income households pay more not only in energy costs as a share of their incomes but also per square foot of their residences [50], and that these households are also particularly susceptible to the negative effects of extreme weather events and global warming [51,52]. Another concern is that the impact of the COVID-19 pandemic on UTR reporting, particularly "full files" of all UTR payments, would disproportionately disadvantage lower-income consumers and minority communities [53]. A recent study found that 25 to 50 percent of consumers who experienced delinquencies did so on utility or telecom tradelines, but not on credit tradelines [47]. Thus, adding these data to consumers' credit files could simply expand the population of consumers with lower credit scores.

The inclusion of rental payments to expand access to credit scores with absent or sparse credit files has garnered a great deal of recent attention. The Federal Housing Finance Administration recently approved the use of rental payments to bolster credit files used in government-sponsored enterprise (GSE) underwriting models. California, Colorado, and the District of Columbia have enacted laws to require government-subsidized landlords to report rental payments to credit bureaus, which are developing reporting standards. However, the inclusion of rental payments poses significant potential challenges. There is wide variation in the timing, consistency, and quality of rental payment and eviction data. Rental payment data are more likely to be collected from large-scale property management companies, yet Black and Latinx renters reside in only 35 percent of the units in buildings with 50 or more units and 44 percent of all units in two-to-four-unit buildings [54].

Several fintech initiatives to provide digital cash-flow data have been implemented to overcome these challenges. FinRegLab [55] analyzed the data from several non-bank financial companies that have adopted cash-flow variables in AI-driven credit decisions instead of traditional indicators and found that cash-flow variables improve predictiveness when used in tandem with traditional credit history information, and, in some cases, can predict default risk with similar effectiveness. Model developers should take note that the data and models reflect value judgments that may include certain biases [13,56]. For example, in existing scoring models, mortgage payments are weighted more heavily than other forms of credit. Blattner and Nelson [16] found that credit scores are less predictive of the default for racial and ethnic minority and low-income mortgage loan applicants, and that these errors have a significant negative impact on mortgage approvals. The authors linked these disparities to the differences in the underlying credit files rather than the biases embedded in the model specification.

Other potential predictors of credit risk include a consumer's GPS location, social media activity, health records, club memberships, educational history, academic performance, and digital footprint. Critics of these approaches have raised concerns that the alternative factors are proxies for demographic characteristics (e.g., race, ethnicity, gender, and family status) that bias credit decisions, and thus, are likely to exacerbate the effects of past marketplace discrimination [29,57]. Moreover, research suggests that the inclusion of non-financial personal data in lending decisions can pose several ethical and legal risks [10]. Models that rely on these data may do little more than automate historical discriminatory practices in mortgage markets, harkening back to the days when Federal Housing Administration guidelines explicitly advised underwriters to consider whether a borrower intends to reside "in a location inhabited by a class or race of people that may impair his interest in the property and thereby affect his motivation [to repay the loan]" [58].

Digitalization in the mortgage industry has introduced opportunities to expand the types of data used in underwriting models, thereby expanding opportunities for homeownership to historically underserved households. However, the use of UTR payment data, cash-flow (i.e., aggregated banking) data, and non-financial personal data in underwriting

raises important ethical and legal questions for those who develop and apply AI in credit scoring. While potentially predictive of repayment and default, these data raise questions based on the SCALE criteria of contextual integrity, accuracy, and perhaps even legality. Although there is overwhelming evidence that these new data sources will expand access to credit scores, it remains to be seen whether this will simply produce a larger pool of consumers with high-risk credit profiles who are more likely to be denied mortgage credit or targeted by subprime lenders [19].

6. SCALE Criteria and Automated Property Valuation Models

Digitalization in the appraisal process has improved efficiency in the loan origination process, and proponents argue that the accuracy of risk assessment has improved as well. Recent innovations include digitalized appraisal inspections whereby appraisers collect certain property data elements without in-person inspections in some cases. This information is then submitted to AI automated valuation models (AVMs) that replace traditional, more subjective procedures.

Meanwhile, appraisal bias has emerged as one of the most controversial issues in the mortgage industry, and several studies have documented systematic biases in traditional appraisals that result in lower values for Black and Hispanic homebuyers and neighborhoods [59]. One widely cited study, for example, revealed that homes owned by Black and Hispanic individuals are more likely to be appraised at a lower value than the sales price [60]. In another recent study, researchers compared traditional appraisals with those conducted via AVMs and found that homes owned by White borrowers are more likely to have an appraised value that is at least 10 percent higher than the AVM's estimated value compared to homes owned by Black borrowers; these overvaluations are also more likely to occur when White borrowers live in majority-Black neighborhoods [61]. Additional evidence suggests that AVM models are less likely to produce biased results, and as such, can be used to advance more equitable outcomes in appraisals for minority homebuyers and homeowners [62].

Concerns that plague credit scoring algorithms also apply in the case of AVMs—namely, the potential for these models to capture and amplify latent discrimination and redlining. Homes owned by Black and Hispanic families as well as homes located in minority neighborhoods have historically and consistently had lower values and rates of house price appreciation than homes owned by similarly-situated White counterparts [63]. AI models could be developed to remove the barriers to equitable outcomes and offset the effects of bias and discrimination in AVMs by assimilating a wider range of data.

In another recent analysis, researchers argued that due to the complexity and dynamic nature of AI models, it would be difficult to identify the specific cause of disparities affecting underrepresented groups or to perform standard fair lending analyses [64]. The authors suggested that existing legal, policy, and regulatory frameworks lag woefully behind in understanding these technologies or how best to oversee their application [64]. To increase transparency, some modelers develop “inherently interpretable” models, while others combine complex models with post hoc explainability methods, i.e., supplemental information. Klutz et al. [65] argued that in addition to transparency and explainability, AI models should be subjected to a higher standard of “contestability”—that is, the extent to which sufficient information is available to meaningfully challenge the model's outcomes. In contexts involving AI/ML applications, contestability would be analogous to consumer protection laws that require, for example, disclosure of the reasons for a mortgage loan denial to the applicant.

Despite the SCALE framework's concerns about accuracy, potential bias, and legality, AI/ML applications have significant potential to expand homeownership opportunities. If calibrated to do so, these models could be deployed to identify sources of bias and discrimination, as well as non-discriminatory alternatives [4]. Davis et al. [12] recently proposed an “algorithmic reparation” approach whereby AI techniques are explicitly

designed to minimize or eliminate the effects of historical disadvantages (e.g., structural racism), rather than to attempt to remove bias from existing algorithms.

7. SCALE Criteria Applied to AI Fraud Detection Models

An advantage of the SCALE framework is that it can be used to evaluate cases where there might or might not be bias. An example is the use of AI models for mortgage fraud detection. Fraud in the form of misrepresentation of borrower identity, reported income or assets, and deed or property-related information poses significant risks to mortgage lenders and housing market intermediaries. Companies increasingly rely on AI and ML techniques to predict the likelihood of misrepresentation based on analyzing patterns in previous fraudulent cases, third-party validation of employment data, account activity, and consumer behavior data from website analytics, digital marketing, and social media activity. High-risk cases, once identified, are referred for further investigation [66].

The SCALE framework could be used to evaluate the impact of AI fraud detection models on racial equity. Examining image patterns and comparing applicant data with verified fraudulent cases aligns with the Societal Values for integrity, and since legitimate applicants grant their approval for data such as credit score, employment, and banking information to be accessed from third-party sources, the Societal Norms for data privacy are also met. Social media activity and other personal behavioral patterns may pose questions about Contextual Integrity, but otherwise, the factors used to predict fraudulent cases are likely appropriate because they mirror those used in underwriting decisions. The primary goal of fraud detection is to increase Accuracy, which depends on the ability of these models to reduce the rate of prediction errors and the likelihood and severity of fraud incidences. A key question about Accuracy is how reliable the data used for validation and/or prediction are and to what extent are these data consistently available across demographic groups. The Legality question pertains to whether AI fraud detection models have a negative, disparate impact on under-represented groups, i.e., if certain applicant profiles are correlated with race or ethnicity, will members of protected classes of prospective borrowers be subjected to more scrutiny than others? This is especially relevant in the case of fraud detection, which is largely unregulated relative to the activities fundamental to the underwriting decisions. Lastly, the Expanded Opportunity factor implies that AI fraud detection models would ultimately either enable previously underserved mortgage borrowers to enter the market or lower the costs of mortgage credit for the members of these groups.

8. Discussion and Conclusions

The emerging use of AI in the origination and servicing of mortgages promises to transform the housing industry and potentially open doors to segments of the population that have previously been underserved. At a minimum, enhanced technologies and the harvesting of non-traditional data should make the process more efficient, consumer-friendly, and less costly. However, unless properly designed, AI could also serve to solidify the historic inequities that currently characterize the housing market and conflict with the stated public policy goal of increasing the homeownership rates of Black, Brown, and lower-income families.

Perhaps the most fundamental issue relates to the types of data that can or should be used in the creation of the underlying models; for example, the assessment of credit risk. As noted earlier, the fact that a given variable is predictive of future loan performance—for example, an individual's educational background or expenditure patterns—is not enough to justify its use. The SCALE typology presented above offers a framework that the industry, its regulators, and Congress could use to assess the pros and cons of deploying various kinds of data in the origination or servicing of mortgages. In the end, such analysis could serve as a basis for drafting new regulations, expanding fair lending laws, and/or enacting more general privacy legislation that explicitly prohibits the use of certain types of data—for example, an individual's social media profile—for certain purposes, including but not necessarily limited to the granting of mortgage credit.

At a minimum, AI models have raised concerns about socio-political priorities to advance racial equity. The use of such tools in other applications has been shown to embed or exacerbate some of the biases that plague human decision-makers. For example, Microsoft's AI chatbot "learned" to respond using racist language gathered from social media users [67], and reports claim that a Twitter algorithm automatically edited out images of Black faces [68]. Racial bias also has been found in popular facial recognition programs and tenant screening algorithms adopted by landlords [69,70]. According to a recent Brookings paper, the use of AI in the mortgage context can embed "biased feedback loops" whereby consumers who previously encountered barriers to traditional forms of credit and obtained financing via higher-risk and more expensive subprime loans have lower credit scores, thereby capturing these circumstances in models for future credit decisions and pricing [59]. Based on the SCALE framework, these approaches also raise concerns in terms of accuracy due to the potential for bias in the representation and selection of samples upon which models are based, in addition to omitted variables and historical factors which could also contribute to systematic errors.

Mortgage lenders, policymakers, and other industry stakeholders should also consider the elements of the SCALE framework when designing, adapting, evaluating, and monitoring digitalized tools. These perspectives could help inform recently proposed legislation (e.g., the H.R.6580 Algorithmic Accountability Bill proposed by Senators Wyden and Booker, and Representative Yvette Clarke) intended to expand the FTC enforcement of AI in housing, financial services, and other industries [71] (<https://www.wyden.senate.gov/news/press-releases/wyden-booker-and-clarke-introduce-bill-to-regulate-use-of-artificial-intelligence-to-make-critical-decisions-like-housing-employment-and-education>, accessed on 26 September 2023). Trade publications and the blogosphere are replete with examples of digitalized solutions claiming to increase efficiency in marketing, operations, risk assessment, regulatory compliance, and servicing. However, there are far fewer frameworks for proactive responsible digital transformation that could provide solutions to the systemic barriers to mortgage credit in the current market structures.

The proposed Algorithmic Accountability Act would direct the FTC to require impact assessments of AI systems and "augmented critical decision processes". Although this proposal acknowledges the issues and would increase resources for the FTC and other agencies to evaluate AI models, it is unclear how audits would account for the iterative, dynamic, and rapidly changing nature of model development. How often should models be evaluated and at what stage of development or implementation? In addition to the concerns about transparency and interpretability described above, as well as the pervasiveness of AI modeling in these industries, it is hard to imagine that the government would ever have the resources to meaningfully evaluate these practices. In the case of AI, industry self-regulation might be a more viable, less costly alternative to traditional regulatory oversight. In partnership with FinRegLab, Blattner, and Spiess, they compared ML credit decision models in terms of explainability and fairness and found that these models vary in the extent to which they can identify characteristics that have a negative and disparate impact on protected classes. These authors propose an approach for "evaluating the quality and usability of information produced about machine learning models' behavior" which could be adopted by lenders and regulators who are seeking transparency in the context of fair lending" [64].

Another policy recommendation is to revisit HUD's 2020 "disparate impact rule" that requires "a robust causal link between the challenged policy or practice and the adverse effect on members of a protected class". The "robust causal link" standard has been difficult to prove or enforce and harkens back to a time when manual underwriting decisions based on a few discrete factors were the norm. Due to the substantial number of factors and combinations thereof in AI/ML models, causal links, including those that unduly harm disadvantaged groups, are difficult to uncover. New language and interpretation of this standard would foster more effective enforcement of the disparate impact legal standard.

As described above, digitalization and AI in the mortgage market can help advance social and political goals of eradicating racism and discrimination, as captured in Davis et al.'s [12] notion of digital reparation. Digitalization strategies could improve accuracy and remove rather than introduce bias; however, such strategies require thoughtful design, development, and implementation. Theoretical and empirical research on the effects of AI and ML, for example, suggest that if designed to do so at the outset, these tools have the potential to identify and eradicate the effects of systemic discrimination while simultaneously increasing predictive accuracy and efficiency in the mortgage value chain. It is important, however, to ensure that tools and approaches adapted from other contexts are appropriate for mortgage lending. The developers of these tools should address potential legal and regulatory issues, such as the potential for discrimination in the form of disparate impact. Lastly, in addition to increased efficiency, lower transaction costs, and/or improved predictiveness, digitalization strategies should be designed to expand opportunities by reducing the barriers associated with manual, more subjective, and biased processes used by many traditional brick-and-mortar institutions.

The Black-White homeownership gap persists due to economic and social disadvantages that have accumulated over generations. The effects of "color-blind" regulations are a subject of heated debate among scholars and policymakers, and some argue that to account for racial effects, race must be explicitly included in models that predict outcomes such as loan defaults. As captured by Samuel Myers's [72] "Minnesota paradox", a reliance on race-neutral metrics of homeownership and other economic outcomes can obfuscate segregation, poverty, and other conditions that exist for Black communities. Ifeoma Ajunwa [73] more broadly described the paradox of automation where more automated decision-making is positioned as an anti-bias intervention, yet "has served to replicate and amplify bias". Recent research on ethics in AI and ML suggests that the models need to include race at the design stage, rather than simply as a test for bias on the back end. Existing, well-intentioned public policies prohibit the inclusion of race as a factor in credit or valuation models. This paradigm fails to acknowledge that race is an endogenous and recursive measure of systematic and institutional discrimination. To address the societal goals of advancing equity and expanding homeownership opportunities, these same models could be used to measure and potentially offset the effects of race in the estimates of credit costs and risks.

Several important unanswered questions remain. For example, what are the appropriate goals for adopting AI/ML approaches, particularly those used to inform lending decisions? Replicating human decisions is one such goal. Should the outcomes of these models (i.e., the ability to assess and price risk) necessarily be superior to bias and other errors often associated with human decisions? We assume firms should use established criteria to assess whether digitalization projects and any innovative programs or projects in mortgage lending expand opportunities for homeownership in underserved communities. More work should be conducted to translate the established criteria of success so they can be applied to the outcomes of these tools. As Thomas and Uminsky [74] noted, defining the outcomes or metrics of success (accuracy, effectiveness, etc.) narrowly, and without regard to the context of the decision, exacerbates underlying problems. Metrics should be broad, multi-faceted, and informed by an understanding of those stakeholders most impacted by the program—in this case, mortgage applicants from historically disadvantaged groups.

Another key issue is how AI/ML can be used to verify fair and equitable treatment of individuals for each of these types of decisions, as prescribed by Davis et al.'s [12] notion of algorithmic reparation. Rather than simply making the decisions, these tools could be used to support or validate decisions being made by humans and/or AI. This suggests an additional, responsible use of AI in the mortgage market: using novel data analytic techniques to monitor, assess, and verify the fair and equitable treatment of mortgage applicants.

Author Contributions: Conceptualization, V.G.P., K.M. and A.S.; investigation, V.G.P., K.M. and A.S.; writing—original draft preparation, V.G.P., K.M. and A.S.; writing—review and editing, V.G.P., K.M. and A.S.; project administration, V.G.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created for this research.

Conflicts of Interest: The authors declare no conflict of interest.

References

- Gartner (n.d.). Digitalization. Available online: <https://www.gartner.com/en/information-technology/glossary/digitalization> (accessed on 7 February 2023).
- IBM Cloud Education (n.d.). What Is Artificial Intelligence (AI)? Available online: <https://www.ibm.com/cloud/learn/what-is-artificial-intelligence?lnk=ile> (accessed on 7 February 2023).
- Brown, S. Machine Learning, Explained. 2021. Available online: <https://mitsloan.mit.edu/ideas-made-to-matter/machine-learning-explained#:~:text=What%20is%20machine%20learning%3F,to%20how%20humans%20solve%20problems> (accessed on 7 February 2023).
- Akinwumi, M.; Merrill, J.; Rice, L.; Saleh, K.; Yap, M. An AI Fair Lending Policy Agenda for the Federal Financial Regulators. Policy Brief. 2021. Brookings Center on Regulation and Markets, Brookings Institution. Available online: <https://www-brookings-edu.cdn.ampproject.org/c/s/www.brookings.edu/research/an-ai-fair-lending-policy-agenda-for-the-federal-financial-regulators/?amp> (accessed on 7 February 2023).
- FinRegLab. The Use of Machine Learning for Credit Underwriting. The Milken Institute. 2021. Available online: https://finreglab.org/wp-content/uploads/2021/09/the-Use-of-ML-for-Credit-Underwriting-Market-and-Data-Science-Context_09-16-2021.pdf (accessed on 7 February 2023).
- Arnold, D.; Dobbie, W.; Hull, P. Measuring Racial Discrimination in Algorithms. *AEA Pap. Proc.* **2021**, *111*, 49–54. [CrossRef]
- Martin, K. Designing Ethical Algorithms. *MISQ Exec.* **2019**, *18*, 129–142. [CrossRef]
- Kroll, J.A. The Fallacy of inscrutability. *Philos. Trans. R. Soc. A* **2018**, *376*, 2018008420180084. [CrossRef]
- Johnson, K.N.; Pasquale, F.A.; Chapman, J.E. Artificial intelligence, Machine Learning, and Bias in Finance: Toward Responsible innovation. *Fordham Law Rev.* **2019**, *88*, 499.
- Perry, V.G.; Schnare, A.B. Tipping the SCALE: Will Alternative Data in Credit Scoring Promote Or Impede Fair Lending Goals? Presentation at the National Association of Realtors Public Policy Forum. 2021. Available online: <https://www.nar.realtor/events/public-policy-forum/tipping-the-scale-will-alternative-data-in-credit-scoring-promote-or-impede-fair-lending-goals> (accessed on 7 February 2023).
- Martin, K. Ethical Implications and Accountability of Algorithms. *J. Bus. Ethics* **2019**, *160*, 835–850. [CrossRef]
- Davis, J.L.; Williams, A.; Yang, M.W. Algorithmic Reparation. *Big Data Soc.* **2021**, *8*. [CrossRef]
- Martin, K.; Nissenbaum, H. What Is It About Location? *Berkeley Technol. Law J.* **2020**, *35*, 251–326. [CrossRef]
- Nissenbaum, H. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*; Stanford Law Books; Stanford University Press: Stanford, CA, USA, 2009.
- Walzer, M. *Spheres of Justice: A Defense of Pluralism*; Basic Books: New York, NY, USA, 1984.
- Blattner, L.; Nelson, S. How Costly Is Noise? Data and Disparities in Consumer Credit [Working Paper]. 2021. Available online: <https://arxiv.org/abs/2105.07554v1> (accessed on 7 February 2023).
- Consumer Financial Protection Bureau (CFPB). Using Publicly Available Information to Proxy for Unidentified Race and Ethnicity. Consumer Financial Protection Bureau. 2014. Available online: https://files.consumerfinance.gov/f/201409_cfpb_report_proxy-methodology.pdf (accessed on 7 February 2023).
- Heaven, W.D. Bias Isn't the Only Problem With Credit Scores—And No, AI Can't Help. *MIT Technology Review*, 17 June 2021. Available online: <https://www.technologyreview.com/2021/06/17/1026519/racial-bias-noisy-data-credit-scores-mortgage-loans-fairness-machine-learning/> (accessed on 7 February 2023).
- Schnare, A.B. Alternative Credit Scores and the Mortgage Market: Opportunities and Limitations. Progressive Policy Institute. 2017. Available online: http://www.progressivepolicy.org/wp-content/uploads/2017/12/UpdatedCreditScoring_2017.pdf (accessed on 7 February 2023).
- Engler, A. The E.U. and U.S. Diverge on AI Regulation: A Transatlantic Comparison and Steps to Alignment. Brookings Institution. 2023. Available online: <https://www.brookings.edu/articles/the-eu-and-us-diverge-on-ai-regulation-a-transatlantic-comparison-and-steps-to-alignment/> (accessed on 26 August 2023).

21. Akinwumi, M.; Rice, L.; Sharma, S. Purpose, Process, and Monitoring: A New Framework for Auditing Algorithmic Bias in Housing and Lending. National Fair Housing Alliance. 2022. Available online: https://nationalfairhousing.org/wp-content/uploads/2022/02/PPM_Framework_02_17_2022.pdf (accessed on 26 August 2023).
22. Federal Reserve. The 2019 Survey of Consumer Finances. 2020. Available online: <https://www.federalreserve.gov/econres/scfindex.htm> (accessed on 28 August 2023).
23. Fannie Mae. COVID-19, Mortgage Digitization, and Borrower Satisfaction. 2021. Available online: <https://www.fanniema.com/media/40491/display> (accessed on 7 February 2023).
24. Atske, S.; Perrin, A. Home Broadband Adoption, Computer Ownership Vary by Race, Ethnicity in the U.S. Pew Research Center. 2021. Available online: <https://www.pewresearch.org/fact-tank/2021/07/16/home-broadband-adoption-computer-ownership-vary-by-race-ethnicity-in-the-u-s/> (accessed on 7 February 2023).
25. Vogels, E.A. Digital Divide Persists Even as Americans with Lower Incomes Make Gains in Tech Adoption. Pew Research Center. 2021. Available online: <https://www.pewresearch.org/fact-tank/2021/06/22/digital-divide-persists-even-as-americans-with-lower-incomes-make-gains-in-tech-adoption/> (accessed on 7 February 2023).
26. Ehrentraud, J.; Garcia Ocampo, D.; Quevedo Vega, C. Regulating Fintech Financing: Digital Banks and Fintech Platforms. FSI Insights on Policy Implementation, 2020, No. 27, Financial Stability Institute, BIS. Available online: <https://www.bis.org/fsi/publ/insights27.pdf> (accessed on 7 February 2023).
27. nLIFT. Fulfilling the Promise of Fintech: The Case for A Nonprofit Vision and Leadership. The Aspen institute. 2018. Available online: https://www.aspeninstitute.org/wp-content/uploads/2018/09/nLIFT-Manifesto-FINAL-1.pdf?_ga=2.176913637.1579357171.1541431738-807926508.1541431738 (accessed on 7 February 2023).
28. Friedline, T.; Chen, Z. Digital Redlining and the Fintech Marketplace: Evidence From US Zip Codes. *J. Consum. Aff.* **2021**, *55*, 366–388. [CrossRef]
29. Hauptert, T. The Racial Landscape of Fintech Mortgage Lending. *Hous. Policy Debate* **2022**, *32*, 337–368. [CrossRef]
30. Matz, S.C.; Menges, J.I.; Stillwell, D.J.; Schwartz, H.A. Predicting individual-Level income From Facebook Profiles. *PLoS ONE* **2019**, *14*, e0214369. [CrossRef]
31. Matz, S.C.; Netzer, O. Using Big Data as a Window into Consumers' Psychology. *Curr. Opin. Behav. Sci.* **2017**, *18*, 7–12. [CrossRef]
32. Libai, B.; Bart, Y.; Gensler, S.; Hofacker, C.; Kaplan, A.; Kötterheinrich, K.; Kroll, E.B. Brave New World? On AI and the Management of Customer Relationships. *J. Interact. Mark.* **2020**, *51*, 44–56. [CrossRef]
33. Ali, M.; Sapiezynski, P.; Bogen, M.; Korolova, A.; Mislove, A.; Rieke, A. Discrimination Through Optimization: How Facebook's Ad Delivery Can Lead to Biased Outcomes. *Proc. ACM Hum.-Comput. Interact.* **2019**, *3*, 199. [CrossRef]
34. Evans, C.; Miller, W. From Catalogs To Clicks: The Fair Lending Implications of Targeted internet Marketing. *Consum. Compliance Outlook* **2019**, *3*. Available online: <https://consumercomplianceoutlook.org/2019/third-issue/from-catalogs-to-clicks-the-fair-lending-implications-of-targeted-internet-marketing/> (accessed on 7 February 2023).
35. Hermann, E. Leveraging Artificial intelligence in Marketing for Social Good—An Ethical Perspective. *J. Bus. Ethics* **2022**, *179*, 43–46. [CrossRef]
36. Martin, K.D.; Murphy, P.E. The Role of Data Privacy in Marketing. *J. Acad. Mark. Sci.* **2017**, *45*, 135–155. [CrossRef]
37. Martin, K.D.; Palmatier, R.W. Data Privacy in Retail: Navigating Tensions and Directing Future Research. *J. Retail.* **2020**, *96*, 449–457. [CrossRef]
38. Stefano, P.; Reczek, R.M.; Giesler, M.; Botti, S. Consumer Experiences With Marketing Technology: Solving the Tensions Between Benefits and Costs. *NIM Mark. Intell. Rev.* **2022**, *14*, 25–29.
39. National Fair Housing Alliance. Fair Housing Groups Settle Lawsuit with Facebook: Transforms Facebook's Ad Platform Impacting Millions of Users. 2019. Available online: <https://nationalfairhousing.org/national-fair-housing-alliance-settles-lawsuit-with-facebook-transforms-facebooks-ad-platform-impacting-millions-of-users/> (accessed on 7 February 2023).
40. Jan, T.; Dwoskin, E. Facebook Agrees to Overhaul Targeted Advertising System for Job, Housing and Loan Ads After Discrimination Complaints. *The Washington Post*, 19 March 2019. Available online: https://www.washingtonpost.com/business/economy/facebook-agrees-to-dismantle-targeted-advertising-system-for-job-housing-and-loan-ads-after-discrimination-complaints/2019/03/19/7dc9b5fa-4983-11e9-b79a-961983b7e0cd_story.html (accessed on 7 February 2023).
41. Brown, A.K. Fair Lending—Digital Marketing and HMDA 2018. Compliance Session at the Marquis User's Conference, Skadden Foundation. 2019. Available online: <https://gomarquis.com/wp-content/uploads/2019/10/Austin-Brown-Fair-Lending-Digital-Marketing-and-HMDA-2018.pdf> (accessed on 7 February 2023).
42. Ballard Spahr. DOJ/CFPB/OCC Settle Redlining Lawsuit Against Mississippi-Based National Bank. *Consumer Finance Monitor*. 2021. Available online: <https://www.jdsupra.com/legalnews/doj-announces-major-new-initiative-2017442/> (accessed on 7 February 2023).
43. Humber, N.J.; Matthews, J. *Fair Housing Enforcement in the Age of Digital Advertising: A Closer Look at Facebook's Marketing*; Working Paper; Roger Williams University: Bristol, RI, USA, 2020. Available online: https://docs.rwu.edu/cgi/viewcontent.cgi?article=1308&context=law_fac_fs (accessed on 7 February 2023).
44. Aronowitz, M.; Golding, E.; Choi, J. The Unequal Costs of Homeownership. MIT Golub Center for Finance and Policy. 2020. Available online: <https://gfcf.mit.edu/wp-content/uploads/2020/10/Mortgage-Cost-for-Black-Homeowners-10.1.pdf> (accessed on 7 February 2023).

45. Ramirez, E.; Brill, J.; Ohlhausen, M.K.; McSweeney, T. Big Data: A Tool for Inclusion or Exclusion? Understanding the Issues, FTC Report, Federal Trade Commission; 2016. Available online: <https://www.ftc.gov/system/files/documents/reports/big-data-tool-inclusion-or-exclusion-understanding-issues/160106big-data-rpt.pdf> (accessed on 7 February 2023).
46. Kreiswirth, B.; Schoenrock, P.; Singh, P. Using Alternative Data to Evaluate Creditworthiness. Consumer Financial Protection Bureau; 2017. Available online: <https://www.consumerfinance.gov/about-us/blog/using-alternative-data-evaluate-creditworthiness/> (accessed on 7 February 2023).
47. Cochran, K.T.; Stegman, M.; Foos, C. Utility, Telecommunications, and Rental Data in Underwriting Credit. Urban Institute Research Report. 2021. Available online: https://www.urban.org/research/publication/utility-telecommunications-and-rental-data-underwriting-credit/view/full_report. (accessed on 7 February 2023).
48. Lee, A.S.; Schnare, A.; Turner, M.A.; Walker, P.D.; Varghese, R. Give Credit Where Credit Is Due: Increasing Access To Affordable Mainstream Credit Using Alternative Data. Urban Markets Initiative Report. Policy and Economic Research Council and the Brookings Institution. 2006. Available online: <https://www.brookings.edu/research/give-credit-where-credit-is-due-increasing-access-to-affordable-mainstream-credit-using-alternative-data/> (accessed on 7 February 2023).
49. CFPB. Report on the Use of Remittance Histories in Credit Scoring. Consumer Financial Protection Bureau; 2014. Available online: <https://www.consumerfinance.gov/data-research/research-reports/report-on-the-use-of-remittance-histories-in-credit-scoring/> (accessed on 7 February 2023).
50. Drehobl, A.; Ross, L. Lifting the High Energy Burden in America's Largest Cities: How Energy Efficiency Can Improve Low Income and Underserved Communities. ACEEE. 2016. Available online: <https://www.aceee.org/sites/default/files/publications/researchreports/u1602.pdf> (accessed on 7 February 2023).
51. Byrne, J.; Portanger, C. Climate Change, Energy Policy, and Justice: A Systematic Review. *Anal. Krit.* **2014**, *36*, 315–343. [CrossRef]
52. Carley, S.; Konisky, D.M. The Justice and Equity Implications of the Clean Energy Transition. *Nat. Energy* **2020**, *5*, 569–577. [CrossRef]
53. National Consumer Law Center. The Credit Score Pandemic Paradox. 2022. Available online: https://www.nclc.org/images/pdf/special_projects/covid-19/IB_Pandemic_Paradox_Credit_invisibility.pdf (accessed on 7 February 2023).
54. Choi, J.H.; Young, C. Owners and Renters of 6.2 Million Units in Small Buildings Are Particularly Vulnerable During the Pandemic. 2020. Available online: <https://www.urban.org/urban-wire/owners-and-renters-62-million-units-small-buildings-are-particularly-vulnerable-during-pandemic> (accessed on 7 February 2023).
55. FinRegLab. The Use of Cash-Flow Data in Underwriting Credit: Empirical Research Findings. The Milken Institute. 2019. Available online: https://finreglab.org/wp-content/uploads/2019/07/FRL_Research-Report_Final.pdf (accessed on 7 February 2023).
56. Friedman, B.; Nissenbaum, H. Bias in Computer Systems. *ACM Trans. Inf. Syst.* **1996**, *14*, 330–347. [CrossRef]
57. Odinet, C.K. Predatory Fintech and the Politics of Banking. *Iowa Law Rev.* **2021**, *10*, 1739–1800.
58. Federal Housing Administration. *Underwriting Manual: Underwriting and Valuation Procedure Under Title II of the National Housing Act*; Federal Housing Administration: Washington, DC, USA, 1938. Available online: <https://www.huduser.gov/portal/sites/default/files/pdf/Federal-Housing-Administration-Underwriting-Manual.pdf> (accessed on 7 February 2023).
59. Rothwell, J.; Perry, A.M. Biased Appraisals and the Devaluation of Housing in Black Neighborhoods. Brookings Institute. 2021. Available online: <https://www.brookings.edu/research/biased-appraisals-and-the-devaluation-of-housing-in-black-neighborhoods/> (accessed on 7 February 2023).
60. Folk, J.; Chen, K. Avoiding Overvaluation Risk and Appraisal Bias in Today's Uniquely Challenging Market Session. Clear Capital. 2021. Available online: <https://www.clearcapital.com/avoiding-overvaluation-risk-and-appraisal-bias-in-todays-uniquely-challenging-market-session> (accessed on 7 February 2023).
61. Williamson, J.; Palm, M. Appraising the Appraisal: A Closer Look at Divergent Appraisal Values for Black and White Borrowers Financing Their Home. Fannie Mae. 2022. Available online: <https://www.fanniemae.com/media/42541/display> (accessed on 7 February 2023).
62. House Canary. Reducing Racial Bias in Home Appraisals Using Automated Valuation Technology. 2021. Available online: <https://www.housecanary.com/wp-content/uploads/2021/12/Reducing-Racial-Bias-in-Home-Appraisals-Using-Automated-Valuation-Technology-December-2021.pdf> (accessed on 7 February 2023).
63. Perry, V.G.; Aronowitz, M.; Choi, J.H.; Golding, E.; Green, M.; Green, R.K.; Jourdain-Earl, M.; Aiko Nelson, A.; Rhue, L.; Rice, L. 2020 State of Homeownership in Black America. 2020. Available online: <https://www.shuba2020.com/> (accessed on 7 February 2023).
64. FinRegLab; Blattner, L.; Spiess, J. Machine Learning Explainability Fairness: Insights From Consumer Lending. Working Paper. 2022. Available online: <https://finreglab.org/wp-content/uploads/2022/04/FinRegLab-Stanford-ML-Explainability-and-Fairness-insights-from-Consumer-Lending-April-2022.pdf> (accessed on 7 February 2023).
65. Klutznick, D.N.; Kohli, N.; Mulligan, D.K. Shaping Our Tools: Contestability as a Means To Promote Responsible Algorithmic Decision Making in the Professions. In *After the Digital Tornado: Networks, Algorithms, Humanity*; Werbach, K., Ed.; Cambridge University Press: Cambridge, UK, 2020; pp. 137–152.
66. Huard, F. How Machine Learning Is Making It Easier to Spot Fraud and Mitigate Risk for Underwriters. Core Logic. 2023. Available online: <https://www.corelogic.com/culture-stories/how-machine-learning-is-making-it-easier-to-spot-fraud-and-mitigate-risk-for-underwriters/> (accessed on 26 August 2023).

67. Schwartz, O. In 2016, Microsoft's Racist Chatbot Revealed the Dangers of Online Conversation. *IEEE Spectr.* 2019. Available online: <https://spectrum.ieee.org/in-2016-microsofts-racist-chatbot-revealed-the-dangers-of-online-conversation> (accessed on 7 February 2023).
68. Collier, K. Twitter's Racist Algorithm Is Also Ageist, Ableist and Islamophobic, Researchers Find. *NBC News* 2021. Available online: <https://www.nbcnews.com/tech/tech-news/twitters-racist-algorithm-also-ageist-ableist-islamophobic-researchers-rcna1632> (accessed on 7 February 2023).
69. Guynn, J. Google Photos Labeled Black People "Gorillas". *USA Today*. 2015. Available online: <https://www.usatoday.com/story/tech/2015/07/01/google-apologizes-after-photos-identify-black-people-as-gorillas/29567465/> (accessed on 7 February 2023).
70. Rosen, E.; Garboden, P.M.E.; Cossyleon, J.E. Racial Discrimination in Housing: How Landlords Use Algorithms and Home Visits to Screen Tenants. *Am. Sociol. Rev.* 2021, *86*, 787–822. [CrossRef]
71. Kaye, K. This Senate Bill Would Force Companies to Audit AI Used for Housing and Loans. *Protocol*. 8 February 2022. Available online: <https://www.protocol.com/enterprise/revised-algorithmic-accountability-bill-ai> (accessed on 7 February 2023).
72. Myers, S.; Samuel, M., Jr. On the Surprise of the Minnesota Paradox and Large Racial Disparities. Video. Humphrey School of Public Affairs, University of Minnesota. 2020. Available online: <https://www.youtube.com/watch?v=6TSLsv1KCWM> (accessed on 7 February 2023).
73. Ajunwa, I. The Paradox of Automation as Anti-Bias intervention. *Cardozo Law Rev.* 2019, *41*, 1671–1742.
74. Thomas, R.; Uminsky, D. The Problem with Metrics is a Fundamental Problem for AI. *arXiv* 2020, arXiv:2002.08512.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.



January 31, 2024

Chair Tina Smith
 Ranking Member Cynthia Lummis
 U.S. Senate Committee on Banking, Housing and Urban Affairs
 Subcommittee on Housing, Transportation, and Community Development
 226 Dirksen Senate Office Building, Washington D.C. 20515

Re: CAIDP Statement for the Record: "Artificial Intelligence and Housing: Exploring Promise and Peril"

Dear Chair Smith, Ranking Member Lummis, and Members of the Subcommittee,

We write to you regarding the hearing on "Artificial Intelligence and Housing: Exploring Promise and Peril."¹ The Center for AI and Digital Policy (CAIDP) appreciates your leadership on addressing the risks and benefits of AI systems, and your work towards establishing standards of AI governance.

The CAIDP is an independent research and education non-profit based in Washington, DC.² Our global network of AI policy experts and advocates advises national governments, international organizations, and congressional committees regarding artificial intelligence and digital policy. Our President, Merve Hickok testified at the first congressional hearing on AI last year—"Advances in AI: Are We Ready For a Tech Revolution?"³ CAIDP routinely provides advice to Congressional Committees on matters involving AI policy. We previously advised the Senate Judiciary Committee on AI in Criminal Prosecutions⁴, AI and Human Rights⁵, Senate HELP Committee on AI and Healthcare,⁶ the

¹ U.S. Senate Committee on Banking, Housing and Urban Affairs, Subcommittee on Housing, Transportation, and Community Development, *Artificial Intelligence and Housing: Exploring Promise and Peril*, (Jan. 31, 2024), <https://www.banking.senate.gov/hearings/artificial-intelligence-and-housing-exploring-promise-and-peril>.

² CAIDP, About, <https://www.caidp.org/about-2/>.

³ Testimony and statement for the record of CAIDP President Merve Hickok, *Advances in AI: Are We Ready For a Tech Revolution?*, House Committee on Oversight and Accountability, Subcommittee on Cybersecurity, Information Technology, and Government Innovation (March 8, 2023), https://oversight.house.gov/wp-content/uploads/2023/03/Merve-Hickok_testimony_March-8th-2023.pdf.

⁴ CAIDP, *Statements*, <https://www.caidp.org/statements/>

⁵ CAIDP, *Statement to Senate Judiciary Committee on "AI and Human Rights"*, (June 13, 2023), <https://www.caidp.org/app/download/8462575863/CAIDP-SJC-06132023.pdf>

⁶ CAIDP, *Statement to Senate HELP Committee on "Avoiding a Cautionary Tale: Policy Considerations for Artificial Intelligence in Healthcare"*, (November 8, 2023), <https://www.caidp.org/app/download/8487454163/CAIDP-Senate%20HELP-AI-Healthcare-11082023.pdf>



Senate Rules Committee on AI and Elections⁷. We also publish the annual *Artificial Intelligence and Democratic Values Report*,⁸ providing a comprehensive review of AI policies and practices in 75 countries.

In brief, our recommendations to this Committee are:

- 1) to exercise oversight on federal government tasked with strengthening AI and civil rights in the broader economy under the Biden AI Executive Order.⁹
- 2) to move forward with AI legislation, such as the the Algorithmic Accountability Act of 2023

Recommendation 1: Exercise oversight on federal agencies tasked with strengthening AI and civil rights in the broader economy under the AI Executive Order

President Biden’s Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence¹⁰ (“AI EO”) is a sweeping document on AI guardrails and seeks to position the federal government as a model of accountable AI development and use.

The principles and policies section of the AI EO states, “*Artificial Intelligence policies must be consistent with my Administration’s dedication to advancing equity and civil rights. My Administration cannot—and will not—tolerate the use of AI to disadvantage those who are already too often denied equal opportunity and justice. From hiring to housing to healthcare, we have seen what happens when AI use deepens discrimination and bias, rather than improving quality of life. Artificial Intelligence systems deployed irresponsibly have reproduced and intensified existing inequities, caused new types of harmful discrimination, and exacerbated online and physical harms.*”¹¹

⁷ CAIDP, *Statement to Senate Rules Committee on “AI and Elections”*, (September 27, 2023), <https://www.caidp.org/app/download/8478562663/CAIDP-SRC-AI-ELECTIONS-09272023.pdf>

⁸ CAIDP, *Artificial Intelligence and Democratic Values (2023)*, <https://www.caidp.org/reports/aidv-2022/>.

⁹ Executive Order 14110 of October 30, 2023, *Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence*, Federal Register Vol. 88, No. 210, pg. 75191-75226, <https://www.govinfo.gov/content/pkg/FR-2023-11-01/pdf/2023-24283.pdf>

¹⁰ Executive Order 14110 of October 30, 2023, *Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence*, Federal Register Vol. 88, No. 210, pg. 75191-75226, <https://www.govinfo.gov/content/pkg/FR-2023-11-01/pdf/2023-24283.pdf>

¹¹ *Id.*, Section 2(d), pg. 75192



Section 7.3 of the AI EO addresses the obligations of the federal government for “*Strengthening AI and Civil Rights in the Broader Economy*.”¹² Notably, the AI EO mandates “[t]o address discrimination and biases against protected groups in housing markets and consumer financial markets”¹³ and encourages the Director of the Federal Housing Finance Agency and the Director of the Consumer Financial Protection Bureau to consider using their authorities, as they deem appropriate, to require their respective regulated entities, where possible, to use appropriate methodologies including AI tools to ensure compliance with Federal law.¹⁴ To achieve such purpose, the AI EO requires the relevant agencies to evaluate their underwriting models for bias or disparities affecting protected groups and evaluate automated collateral-valuation and appraisal processes in ways that minimize bias.

The AI EO directs the Secretary of Housing and Urban Development (HUD) shall, and the Director of the Consumer Financial Protection Bureau (CFPB) to issue additional guidance: “... to combat unlawful discrimination enabled by automated or algorithmic tools used to make decisions about access to housing and in other real estate-related transactions.”¹⁵ and specifically requires the guidance to address:

- “the potential violation of the Fair Housing Act (Public Law 90–284), the Fair Credit Reporting Act (Public Law 91–508), or other relevant federal laws by tenant screening systems. This includes examining how the use of data, such as criminal records, eviction records, and credit information, can result in discriminatory outcomes that violate Federal law, and
- how the Fair Housing Act, the Consumer Financial Protection Act of 2010 (Title X of Public Law 111–203), or the Equal Credit Opportunity Act (Public Law 93–495) apply to the advertising of housing, credit, and other real estate-related transactions through digital platforms, including those that use algorithms to facilitate advertising delivery. Also, providing best practices to avoid violations of Federal law.”¹⁶

We urge this Committee to exercise oversight on the actions to be completed by HUD and CFPB under the AI EO, specifically the issuance of additional guidance to address the risks of bias, discrimination, exclusion, and predatory practices in housing, credit related decisions arising from the use of automated decision-making systems.

¹² *Id.*, pg. 75213

¹³ *Id.*

¹⁴ *Id.*

¹⁵ *Id.*, Section 7.3(c), pg. 75213.

¹⁶ *Id.*



Recommendation 2: Move forward with comprehensive AI legislation

We commend the agency actions to guide the appropriate use of AI— including the April 2023 joint statement by the FTC, DOJ, EEOC, and CFPB on bias in automated systems,¹⁷ and the White House Blueprint for an AI Bill of Rights.¹⁸

However, these measures do not address the full spectrum of governance required for the development, deployment, and use of AI systems. We support the Blumenthal-Hawley Bi-Partisan Framework for U.S. AI Act¹⁹ and the Algorithmic Accountability Act²⁰. We believe the implementation of either of these initiatives would go a long way in filling the legislative vacuum in which high-risk AI systems operate.

AI-powered screening, rental programs can greatly exacerbate the current problems with access to housing and credit, for already underserved populations and vulnerable groups.²¹ Generative AI systems can produce convincing yet fake datasets,²² enabling the spread of housing market misinformation through fake reviews, testimonials, articles, or emails. There are also examples of algorithmic systems facilitating collusion among landlords to raise rental prices to the detriment of consumers in the market. Just last year, the Department of Justice joined an antitrust lawsuit filed by dozens of tenants against Texas-based software provider Real Page for its price-setting algorithms which inflated rents.²³

¹⁷ EEOC, CRT, FTC, and CFPB, *Joint Statement On Enforcement Efforts Against Discrimination And Bias In Automated Systems*, (Apr. 25, 2023), https://www.ftc.gov/system/files/ftc_gov/pdf/EEOC-CRT-FTC-CFPB-AI-Joint-Statement%28final%29.pdf.

¹⁸ White House Office of Science and Technology Policy, *Blueprint for an AI Bill of Rights*, (October 2022), <https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf>.

¹⁹ Senator Richard Blumenthal & Senator Josh Hawley, *Bipartisan Framework for U.S. AI Act*, <https://www.blumenthal.senate.gov/imo/media/doc/09072023bipartisanaiframework.pdf>

²⁰ Office of Senator Cory Booker, *Booker, Wyden, Clarke Introduce Bicameral Bill to Regulate Use of Artificial Intelligence to Make Critical Decisions like Housing, Employment and Education*, Press Release, (Sept. 21, 2023), <https://www.booker.senate.gov/news/press/booker-wyden-clarke-introduce-bicameral-bill-to-regulate-use-of-artificial-intelligence-to-make-critical-decisions-like-housing-employment-and-education>

²¹ Consumer Reports, *How Tenant Screening Reports Make It Hard for People to Bounce Back From Tough Times*, (Mar. 11, 2021), <https://www.consumerreports.org/electronics/algorithmic-bias/tenant-screening-reports-make-it-hard-to-bounce-back-from-tough-times-a2331058426/>

²² Miryam Naddaf, *ChatGPT Generates Fake Data Set to Support Scientific Hypothesis*, 623 *Nature* 895 (2023).

²³ ProPublica, *DOJ Backs Tenants in Case Alleging Price-Fixing by Big Landlords and a Real Estate Tech Company*, (November 16, 2023), <https://www.propublica.org/article/doj-backs-tenants-price-fixing-case-big-landlords-real-estate-tech>



Automated systems used to determine who gets housing rely on faulty algorithms and systems which are opaque, predictive, and frequently based on inaccurate data.²⁴ The discriminatory consequences are not just theoretical; there is growing body of litigation because of disparate impact on minorities.²⁵ More concerning is that the widespread use of AI driven decision making relating to housing, renting, or credit is constricting avenues for meaningful redress to contest adverse outcomes in court as petitioners struggle to establish “standing.”²⁶ The models making these decisions are not interpretable²⁷, and as such creates barriers to establishing “traceability” for the purpose of “standing” when attempting to enforce private rights of action under applicable statutes.

Specific to housing and related credit or real-estate transactions, we urge the Committee to consider the following recommendations in developing legislation:

- i. The consumer must be provided with clear, transparent, and complete information about AI-driven decision-making processes, including the subjects and elements involved. Simple processes should be in place for appealing any such decision.
- ii. Affected parties should have a right to contest adverse decisions made by AI systems.
- iii. Decisions about housing allocation should not rely solely on algorithmic outcomes, and companies utilizing AI systems be able to provide explanations for their decisions.
- iv. Implement ex-ante impact assessments and ex-post evaluation or audit mechanisms for any AI system that implicates civil rights or public safety.

²⁴ TechEquity Collaborative, *The Promises and Perils of Residential Proptech - Year 1 Summary Report*, (Apr. 2023), <https://techequitycollaborative.org/2023/04/25/the-promises-and-perils-of-residential-proptech-year-1-summary-report/>

²⁵ *Louis v. Saferent Sols., LLC*, No. 22-CV-10800-AK, 2023 WL 4766192 (D. Mass. July 26, 2023); *Connecticut Fair Hous. Ctr. v. Corelogio Rental Prop. Sols., LLC*, 369 F. Supp. 3d 362 (D. Conn. 2019); *In re: TRANSUNION RENTAL SCREENING SOLUTIONS, INC. FCRA Litigation.*, 2021 WL 4932957; *Perez v. Transunion Rental Screening Solutions, Inc.*, 3:18CV02509

²⁶ U.S. Const. art. III, § 2, cl. 1. This situation is mainly due to the fact that “[t]o meet the constitutional minimum standing requirements, all plaintiffs must establish they have (1) suffered an injury in fact, (2) that is fairly traceable to the challenged conduct of the defendant, and (3) that is likely to be redressed by a favorable judicial decision.”

²⁷ Pauline T. Kim, *Data-Driven Discrimination at Work*, 58 Wm. & Mary L. Rev. 857 (2017), <https://scholarship.law.wm.edu/vmlr/vol58/iss3/4>



Given the serious challenges to housing access, affordability, and equity that are worsened by the widespread use of AI systems, last year, a coalition of 15 Attorney Generals submitted a comment letter to the FTC and CFPB recommending many of the measures set forth above.²⁸ AI Legislation is now urgent and we impress upon this Committee to support the Algorithmic Accountability Act 2023 in moving forward to mark-up.

Thank you for your consideration of our views. We ask that this statement be included in the hearing record. We would be pleased to provide you and your staff with additional information.

Sincerely yours,

Marc Rotenberg
CAIDP Executive Director

Merve Hickok
CAIDP President

Christabel Randolph
Law Fellow

Natalia Alarcón Rueda
Research Assistant

Md Abdul Malek
Research Assistant

²⁸ State of California, Department of Justice, *Attorney General Bonta Submits Comment Letter Recommending Reforms to the Tenant Screening Process*, Press Release, (May 31, 2023), <https://oag.ca.gov/news/press-releases/attorney-general-bonta-submits-comment-letter-recommending-reforms-tenant>



**FAMILY HOUSING
FUND**

Senate Committee on Banking, Housing, and Urban Affairs
Subcommittee on Housing, Transportation, and Community Development
534 Dirksen Senate Office Building
Washington, D.C. 20510

February 7, 2024

Dear Chairwoman Smith, Ranking Member Lummis, and Committee members,

Thank you for the opportunity to submit a statement for the record for the hearing **Exploring the Promises and Perils of Artificial Intelligence in America's Housing Sector**.

Family Housing Fund is a nonprofit intermediary and supporting organization of government dedicated to building a strong, equitable housing system that ensures everyone has a home. We explore emerging issues, pilot innovations, and build pathways for systems change in the Twin Cities region, adding data and nuance to issues to inform potential solutions.

We have illuminated a thread across two distinct bodies of work related to the role and impact of Artificial Intelligence (AI) and automation that poses a challenge to ensuring fair and equitable access to rental housing and homeownership opportunities.

The Tenant Screening Industry

In our work to ensure fair access to rental housing and create a just tenant selection system, we have been exploring research related to the impact and utility of tenant screening reports and examining emerging research around the Tenant Screening industry more broadly. In November 2022, the Consumer Financial Protection Bureau (CFPB) issued two reports on the tenant background check industry.¹ The CFPB's analysis of more than 24,000 complaints highlighted how errors and false information in tenant background checks raise costs and barriers to quality rental housing. Among key findings in their analysis, the CFPB specifically highlighted the following:

- **As corporate landlords have increased their rental holdings, the demand for digital, algorithmic scoring of prospective tenants has increased:** The automated property management systems with centralized databases relied on by corporate landlords and private equity firms substitute a single algorithmic score for the more nuanced and holistic evaluation of prospective tenants done historically by smaller landlords and property managers.
- **Market dysfunctions result in companies selling erroneous data to landlords:** Tenant screening companies appear inclined to include negative information on a report even if that information might be inaccurate. The tenant scores produced for landlords make decision-making easy, but the social scores can hide data errors and magnify the negative impact of erroneous and outdated information.

¹ [CFPB Reports Highlight Problems with Tenant Background Checks | Consumer Financial Protection Bureau \(consumerfinance.gov\)](https://www.consumerfinance.gov)

Additionally, In August 2022, researcher Wonyoung So studying with the Department of Urban Studies and Planning at MIT published "*Which Information Matters? Measuring Landlord Assessment of Tenant Screening Reports*"² designed to understand how landlords assess the content and/or presentation of tenant screening reports and use these reports in their rental decisions. In this study, So found that when risk scores or recommendations from the screening agency were present, the landlord defaulted to the score in their assessment regardless of any detail present with the record – referred to as *Automation Bias*.

The growing reliance on these scoring tools suggests that increasingly, tenant screening algorithms – not housing providers – are determining which renters can access housing opportunities. More exploration is needed to understand to what extent explicit racial bias is present in these algorithms, yet the evidence we have to date suggests a concerning fair housing issue.

Single-Family Rentals and Corporate Investors

In the Twin Cities and other major U.S. markets, large investors purchased a substantial portion of single-family homes following the foreclosure crisis and Great Recession, and these investors continue to grow their share of the housing supply. Research shows these investors disproportionately purchase distressed homes in neighborhoods with high BIPOC populations, and digital technology is a central component of private equity real estate operations, from acquisition to leasing to management.

Increasing utilization and reliance on automation in tenant screening and property management supports greater capacity to expand holdings and may contribute to unequal access to single-family rental homes. Emerging research has shown how these companies' rental management strategies deliberately minimize physical maintenance and evade reasonable requests for repairs and government oversight, while also pursuing increased rents, fees, and other revenue.

At the same time, Private Equity Investors are outbidding first-time homebuyers – often with cash offers made sight-unseen. For the acquisition process, private equity's algorithmic systems consider "neighborhood desirability, proximity to employment centers, transportation corridors, community amenities, construction type, and required ongoing capital needs."³ This allows firms to quickly acquire huge amounts of property, assess value, and underwrite mortgages without the need for major human labor inputs. In one example, researchers found that Invitation Homes assessed about a million homes to purchase 50,000.⁴ In another example, within 15 minutes of a home coming on the market, Progress Residential could assess whether it should be considered for acquisition and could make an offer within two hours.⁵

As corporate investor ownership increases in the Twin Cities, these dynamics frustrate the efforts of local government and the affordable housing community to support first-time home buyers, to close

² [Which Information Matters? Measuring Landlord Assessment of Tenant Screening Reports \(tandfonline.com\)](https://tandfonline.com)

³ [Automated landlord: Digital technologies and post-crisis financial accumulation - Desiree Fields, 2022 \(sagepub.com\)](https://sagepub.com)

⁴ Desiree Fields. *Automated landlord: Digital technologies and post-crisis financial accumulation*. Economy and Space, 2019. p. 169.

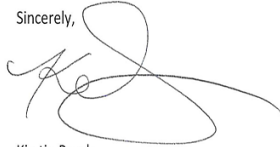
⁵ Peter Whoriskey, Spencer Woodman, and Margot Gibbs, "This Block Used to Be for First-Time Homebuyers. Then Global Investors Bought In," Washington Post, December 15, 2021, <https://www.seattletimes.com/business/real-estate/this-block-used-to-be-for-first-time-homebuyers-then-global-investors-bought-in/>

racial gaps in homeownership and wealth, and to ensure a safe, healthy, and equitable leasing and living experience for Twin Cities renters who choose single-family homes.

Overall, the increased utilization of automation and Artificial Intelligence in an effort to streamline processes and service delivery may be posing a greater threat to the Fair Housing rights of our communities.

Thank you for this opportunity to comment on this important issue.

Sincerely,

A handwritten signature in black ink, appearing to read 'K. Burch', with a large, stylized flourish extending from the end of the signature.

Kirstin Burch
Program Director
Family Housing Fund
kirstin@fhfund.org
612-274-7691



February 2, 2024

The Honorable Senator Tina Smith
 Chair, Senate Subcommittee on Housing,
 Transportation, and Community Development
 720 Hart Senate Office Building
 Washington, DC 20510

The Honorable Senator Cynthia Lummis
 Ranking Member, Senate Subcommittee on Housing,
 Transportation, and Community Development
 Russell Senate Office Building, Suite SR-1274
 Washington, DC 20510

Re: Subcommittee Hearing on Artificial Intelligence and Housing: Exploring Promise and Peril, January 31, 2024

Dear Chairwoman Smith and Ranking Member Lummis:

Thank you for holding this important hearing on the role of artificial intelligence (AI) in the mortgage and rental industries. The Urban Institute is a nonpartisan, nonprofit research organization that provides data and evidence to advance upward mobility and equity. We appreciate the opportunity to share recent findings from our analysis on how housing stakeholders have begun adopting AI to promote efficiency and equity in mortgage finance. As an organization, Urban does not take positions on issues; rather, it empowers its researchers to follow the evidence and make recommendations based on their findings. As such, the views expressed in this letter are those of the authors and should not be attributed to the Urban Institute, its trustees, or its funders.

Based on extensive qualitative research, we believe AI is poised to transform the mortgage industry. The potential efficiency gains for the sector are significant, as is the promise of AI to overcome human biases and the challenges of evaluating nontraditional financial profiles, more common for mortgage applicants of color. But reports across various economic sectors illustrate AI's potential to perpetuate racial inequities.¹ In addition, without challenging or correcting for the underlying causes of bias in the data they use, AI models can simply perpetuate and embed racial inequality on a larger scale.²

1. AI can bring relief for an industry facing high production costs, but greater efficiency must be matched with expanded access. Loan production costs are high and have risen in recent years. Between the first quarter of 2009 and the first quarter 2019, before the recent refinance wave, mortgage production costs rose 250 percent, from \$3,700 to \$9,300 per loan. Over this same period, inflation, measured as the change in the Consumer Price Index, rose only 19 percent. As of the fourth quarter of 2022, the average cost was \$12,450.³

¹ Crystal Grant, "Algorithms Are Making Decisions about Health Care, Which May Only Worsen Medical Racism," ACLU, October 3, 2022, <https://www.aclu.org/news/privacy-technology/algorithms-in-health-care-may-worsen-medical-racism>; and Associated Press, "U.S. Warns of Discrimination in Using Artificial Intelligence to Screen Job Candidates," NPR, May 12, 2022, <https://www.npr.org/2022/05/12/1098601458/artificial-intelligence-job-discrimination-disabilities>.

² Nico Grant and Kashmir Hill, "Google's Photo App Still Can't Find Gorillas. And Neither Can Apple's," *New York Times*, May 22, 2023, <https://www.nytimes.com/2023/05/22/technology/ai-photo-labels-google-apple.html>. Several accounts of biased AI business practices have been reported in popular media. For example, Microsoft AI's chatbot "learned" to respond using racist language gathered by social media users; racially biased facial recognition programs, such as Google's facial recognition program, labeled photos of Black people as gorillas; and a Twitter algorithm was discovered to automatically edit out images of Black faces.

³ Connie Kim, "What Opportunities Will the ICE-Black Knight Deal Present for Lenders?" *HousingWire*, April 4, 2023, <https://www.housingwire.com/articles/what-opportunities-will-the-ice-black-knight-deal-present-for-lenders/>.

500 L'Enfant Plaza SW
 Washington, DC 20024

urban.org

By eliminating human bias in decisionmaking, AI also holds great potential to address the racial homeownership gap. But AI must first overcome the biases and inequities embedded into the data it analyzes. And policymakers and the mortgage industry must reckon with historical and present-day barriers that lock would-be homebuyers of color out of the market altogether.

2. **The use of AI, when aided by alternative data sources that illustrate a mortgage applicant's ability to pay a mortgage, can help close racial differences in mortgage originations.** AI can help extend mortgage credit to borrowers whose credit scores or histories would have disqualified them under traditional underwriting standards. And by incorporating alternative data sources, the benefits for applicants of color could be further improved. As a result, AI can facilitate an expanded use of data that can help a lender improve underwriting outcomes for people of color.

By incorporating more data that capture people of color's ability to qualify for a mortgage, underwriting can become more equitable. Applying AI and machine learning to typical credit bureau tradeline data to pull more variables allows for more data that provide a clearer picture of thin-file borrowers. Separately, more information can be obtained using alternative data, which can also be useful to score thin-file borrowers. Combining alternative measures of credit, such as rental payment histories, into AI-based models can optimize the ability of AI to close racial mortgage denial rate gaps.⁴

3. **Several risks emerge with AI and the data it uses that could disproportionately affect low-income communities and communities of color.** Research suggests that machine learning models, which also use alternative data, can result in larger biases in mortgage underwriting by uncovering sophisticated ethnicity and race proxies invisible to statistical models.⁵ Additional research confirms this conclusion while illustrating that adopting fairness techniques may result in worse outcomes for both Black and white mortgage applicants.⁶

AI models also often lack transparency and explainability and are sometimes seen as "black boxes." The increasing complexity of AI models may mean that humans have a hard time understanding what the model is doing or the relationship between the variables that go into it and the results it generates. AI also has the potential to undermine privacy.⁷ AI is optimized with a lot of data as inputs, including personal information. In addition, as AI expands its footprint, its use could cover more processes that use personal information. And AI can create false information about people. Unaccounted for, these risks could undermine the safety and soundness of the mortgage application process.

Based on these three results, we find that without intentional design, AI could further existing disparities in the mortgage market and even create new ones. Designers of AI models should determine whether the inputs and the underlying data that models learn from reflect the housing market's disparate racial impacts. Prioritizing equity must happen from the start, not as an afterthought. Predicting potential harms ahead of development or implementation—rather than playing catch-up after the product or program is already in

⁴ Authors' interview, from Michael Neal, Linna Zhu, Caitlin Young, Vanessa G. Perry, and Matthew Pruitt, *Harnessing Artificial Intelligence for Equity in Mortgage Finance* (Washington, DC: Urban Institute, 2023).

⁵ Xueru Zhang, Mohammad Mahdi Khalili, and Mingyan Liu, "Long-Term Impacts of Fair Machine Learning," *Ergonomics in Design* 28, no. 3 (2019): 7, <https://doi.org/10.1177/1064804619884160>.

⁶ Leying Zou, and Warut Khern-am-nau, "AI and Housing Discrimination: The Case of Mortgage Applications," *AI and Ethics* 3 (2023): 1271, <https://doi.org/10.1007/s43681-022-00234-9>.

⁷ Tom Wheeler, "The Three Challenges of AI Regulation," Brookings Institution, June 15, 2023, <https://www.brookings.edu/articles/the-three-challenges-of-ai-regulation/>.

motion⁸—and documenting inputs, processes, and results can bring transparency to the innovative process. Intentional design should thus focus on equitable AI. This approach combines the AI tool in question with the power of human insight and thoughtfulness about potential unintended benefits or harms.

We also recommend that improved federal guidance on best practices for adoption can deepen the use of AI throughout the mortgage industry in a safe and sustainable way. Federal actors have two vital roles that will fundamentally set the course for AI in mortgage finance: (1) regulatory guidance and standard setting and (2) advancing adoption of AI-based systems. Federal agencies play a significant role in the mortgage industry. Their presence can help expand and standardize processes powered by AI. But lack of clear regulatory standards governing the use of AI is a major concern across the mortgage industry. Regulation in the AI space should clarify standards for lenders partnering with third-party vendors (e.g., financial technology companies) offering AI tools. Federal government actors such as the government-sponsored enterprises, Fannie Mae and Freddie Mac, and the Federal Home Loan Bank System should use pilot programs to test how AI models affect industry and consumer outcomes, particularly around data use.

For questions or to meet with any of the Urban Institute's research experts, please contact Justine Davenport, senior federal affairs adviser, at jdavenport@urban.org.


Sincerely,

Michael Neal
Senior Fellow, Housing Finance Policy Center
Urban Institute

Linna Zhu
Senior Research Associate, Housing Finance Policy Center
Urban Institute

Matthew Pruitt
Research Assistant, Housing Finance Policy Center
Urban Institute

⁸ Casey Fiesler, "AI Has Social Consequences, but Who Pays the Price? Tech Companies' Problem with 'Ethical Debt,'" *The Conversation*, April 18, 2023, <https://theconversation.com/ai-has-social-consequences-but-whopays-the-price-tech-companies-problem-with-ethical-debt-203375>.



RESEARCH REPORT

Harnessing Artificial Intelligence for Equity in Mortgage Finance

Michael Neal

Linna Zhu

Caitlin Young

Vanessa Gail Perry

Matthew Pruitt

November 2023





ABOUT THE URBAN INSTITUTE

The Urban Institute is a nonprofit research organization that provides data and evidence to help advance upward mobility and equity. We are a trusted source for changemakers who seek to strengthen decisionmaking, create inclusive economic growth, and improve the well-being of families and communities. For more than 50 years, Urban has delivered facts that inspire solutions—and this remains our charge today.

FHLBank San Francisco

ABOUT THE FEDERAL HOME LOAN BANK OF SAN FRANCISCO

The Federal Home Loan Bank of San Francisco is a member-driven cooperative helping local lenders in Arizona, California, and Nevada build strong communities, create opportunity, and change lives for the better. The tools and resources we provide to our member financial institutions—commercial banks, credit unions, industrial loan companies, savings institutions, insurance companies, and community development financial institutions—foster homeownership, expand access to quality housing, seed or sustain small businesses, and revitalize whole neighborhoods. Together with our members and other partners, we are making the communities we serve more vibrant and resilient.

Contents

Acknowledgments	iv
Executive Summary	v
Harnessing Artificial Intelligence	1
The Potential of AI in Mortgage Finance	8
AI throughout the Mortgage Life Cycle	13
How Can AI Embed Discrimination into the Mortgage System?	23
The AI Ecosystem	31
Recommendations to Advance Racial Equity through AI	38
Notes	45
References	49
About the Authors	51
Statement of Independence	53

Acknowledgments

This report was funded by the Federal Home Loan Bank of San Francisco. We are grateful to them and to all our funders, who make it possible for Urban to advance its mission.

The views expressed are those of the authors and should not be attributed to the Urban Institute, its trustees, or its funders. Funders do not determine research findings or the insights and recommendations of Urban experts. Further information on the Urban Institute's funding principles is available at urban.org/fundingprinciples.

The authors are grateful for the time and insights of the organizations and individuals with whom we spoke in developing this report. In addition to those profiled in the report, we acknowledge participants CoreLogic, the Federal Housing Finance Agency, the Mortgage Industry Standards Maintenance Organization, Steve O'Conner and the Mortgage Bankers Association, Upstart, F. Dan Siciliano, Peter Hull, David Arnold, Michael Akinwumi, Will Dobbie, Peter Bergmen, Ed Sivak, and Kareem Saleh for their contributions to this piece. We also thank David Hinson for his thoughtful and careful copyedits.

Executive Summary

Artificial intelligence (AI) is poised to transform the mortgage industry. The potential efficiency gains are significant, as is the promise of AI to overcome human biases and the challenges of evaluating nontraditional financial profiles, more common for mortgage applicants of color. But reports across various economic sectors illustrate AI's potential to perpetuate racial inequities.¹ Without challenging or correcting for the underlying causes of bias in the data they use, AI models can simply perpetuate and embed racial inequality on a larger scale.²

In this report, we highlight current use of AI and its potential to achieve both efficiency and racial equity at several key steps of the mortgage process. In addition, based on interviews with industry stakeholders, we describe the mortgage ecosystem and the role various stakeholder groups play in influencing AI adoption and penetration.

The analysis in this report is based on a qualitative research approach. The data collection included nearly 50 interviews with key stakeholders in the mortgage ecosystem. This includes staff members in the federal government, financial technology (fintech) companies, mortgage lenders, consumer advocates, and research organizations.

Over these interviews, several key trends consistently arose:

- There is no clear agreement on the definition of "artificial intelligence."
- Firms largely use AI to pursue efficiency, often through machine learning.
- AI could amplify existing racial disparities or create new ones.
- The ability of AI to improve racial equity can be undermined by the data used to train the algorithm, not just by the algorithm itself.
- Some of the most promising AI-based underwriting models are also the most controversial, such as explicitly incorporating race up front in underwriting models.
- Within the mortgage process, AI is being used in marketing, underwriting, property valuations, and fraud detection and is beginning to be incorporated into servicing.
- In terms of adopters, AI has been used by the government-sponsored enterprises (GSEs), large mortgage lenders, and fintech firms. Adoption rates appear lower among smaller and mission-oriented lenders, such as minority depository institutions (MDIs) and community development financial institutions (CDFIs).

Based on our interviews and our understanding of AI use in the mortgage finance ecosystem, we recommend actions in the following areas:

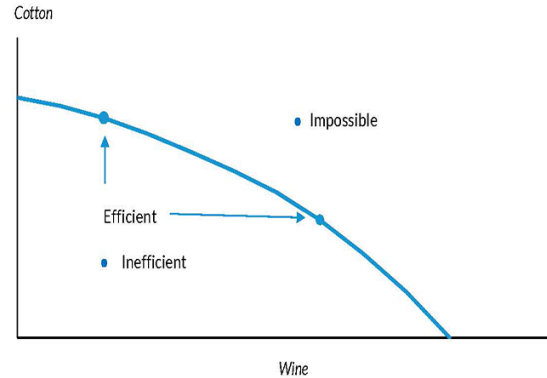
- **Intentional design.** Designers of AI models should determine whether the inputs and the underlying data that models learn from reflect the housing market's disparate racial impacts. Prioritizing equity must happen from the start, not as an afterthought. MDIs and CDFIs provide a blueprint for detecting the underlying bias in a model's data, and policy should move toward a requirement that firms should debias AI models to correct these disparities.
- **Pilot programs and innovation.** The GSEs and Ginnie Mae should use pilot programs to test how AI models affect industry and consumer outcomes, particularly around data use and mortgage servicing.
- **Increased regulatory guidance.** Federal regulators must improve trust and equity in AI in the mortgage finance industry through interagency coordination and by setting clear guidelines for its application. Federal regulators should also clarify the rights of consumers to access private data sources, especially those from protected classes.

Harnessing Artificial Intelligence

Framing the Efficiency and Racial Equity Dimensions of Artificial Intelligence

One of the benefits of AI is its ability to integrate and analyze a broad range of data to maximize predictive accuracy and other outcomes, such as equity. In theory, AI should help close racial inequities in mortgage outcomes. To motivate this intuition, we start with an idea in economics known as the production possibilities curve.³ In its simplest form, a production possibilities curve illustrates the trade-offs an economy must make to achieve production of a desired mix of products or services. Allocating more resources to the production of one item means there are fewer remaining resources to produce the other item.

FIGURE 1
Production Possibilities Curve



URBAN INSTITUTE

Source: The Balance Money.

Equity Accelerator for Homeownership

This report is part of a series, commissioned by the Federal Home Loan Bank of San Francisco, that examines how innovations within the mortgage finance system can help narrow Black-white homeownership gaps. Homeownership is the primary way many US households build wealth. But because of historical racism and its ongoing legacies, the path to homeownership for Black households is rife with structural barriers, and, even once obtained, homeownership's benefits are not equitably distributed. To address some of the persistent racial disparities in homeownership and wealth, the Federal Home Loan Bank of San Francisco has partnered with the Urban Institute to launch a research and product development initiative called the Racial Equity Accelerator for Homeownership. The accelerator hosts several research workstreams investigating methods for facilitating and sustaining Black homeownership:

- incorporating alternative data into mortgage underwriting
- mitigating the impact of student loan debt on Black homeownership
- innovating loss mitigation strategies to help households sustain homeownership during times of stress
- using artificial intelligence and advancing technologies that can overcome mortgage lending biases (the focus of this report)

Although the accelerator focuses on Black homeownership, many of the barriers Black households face apply to other households of color, and the solutions to reduce the Black-white homeownership gap can help other households who struggle to become homeowners and build wealth.

Addressing the Black-white homeownership gap is essential to achieving racial equity and ensuring all households have access to homeownership. Historically, the mortgage finance system purposefully excluded Black households from homeownership through racist practices such as redlining, and the legacies of these practices persist. To undo the effects of explicit historical racism in housing, an equally explicit commitment must be made to address racial homeownership disparities. Without such a commitment, homeownership and wealth gaps will widen.

Rooting out systemic racism is a complicated process that will require sustained collaboration between many actors in the housing finance system, and the policy and practice changes proposed in this report may improve Black homeownership only at the margins. But our research in this area is promising, and if the housing finance system can rally the necessary political will, we may be able to make tangible improvements for hundreds of thousands of Black households.

In a simple and traditional example, an economy chooses between two items, cotton and wine. The inputs needed to produce cotton cannot be used to produce wine, and conversely, if the inputs are used to produce wine, they cannot be used to produce cotton. In extreme situations, if all resources are used

to produce cotton, no wine can be produced, and similarly, if all resources are used to produce wine, no cotton can be produced.

Figure 1 illustrates the production choices in this hypothetical economy, with each choice represented by a point on the chart. Any point along the production possibilities curve is considered “productively efficient,” which means the economy is using all its productive resources as efficiently as possible. The combination of goods the economy produces can be influenced by preferences or market demand.

How Can a Production Possibilities Curve Relate to AI?

In the previous example of a simple economy, the maximum output of cotton and wine and the continuum of trade-offs associated with their production were easily quantified. Though less readily quantifiable, policy analysis has demonstrated a similar trade-off between efficiency and equality (Okun 2015).

Efficiency refers to the allocation of resources that maximizes overall productivity. Equality exists when intersections of social identities, residence in marginalized communities, or experiences with oppressive systems do not determine opportunities, access to resources, and outcomes in life. Achieving equality requires acknowledging, addressing, and dismantling systemic biases in mind-sets, practices, and policies (Venkateswaran et al. 2023).

Actions that boost efficiency often come at the expense of greater equality, while steps taken to improve equality often limit efficiency. Policies that promote efficiency often prioritize market mechanisms and individual incentives, which can lead to unequal outcomes. Moreover, efficiency-driven policies may overlook externalities and social costs that can negatively affect vulnerable populations. If market preferences or market demand desired only equality, the industry would be less efficient. Conversely, if the industry chooses to focus solely on efficiency, the trade-off would be less equitable outcomes than are otherwise possible.


Policymakers may put their finger on the scale in favor of societal priorities. Finding the right balance between efficiency and equality involves complex considerations.

But production theory posits that innovative techniques could move the trade-off curve outward. In other words, a more innovative process can boost the production of cotton and wine simultaneously

and therefore raise production overall. For example, the development of machines to churn wine may free up human resources that can be applied to cotton production.

The same is true for the trade-off between efficiency and equality in mortgage finance (figure 2). For example, use of an automated valuation model (AVM) to determine home values—in place of a human appraiser—can improve both equality and efficiency. AVMs may improve efficiency by reducing the amount of time required to estimate a property's value. And eliminating the human appraiser may remove instances of conscious or unconscious racial bias.

FIGURE 2
The Equality-Efficiency Trade-Off and the Effect on Innovation
The production possibilities curve moves out because of innovation



In this report, we approach AI as a trade-off curve-shifting innovation. On one hand, AI can boost efficiency and accuracy while reducing labor costs. In addition, AI provides greater modeling flexibility that can account for the experiences for people of color that traditional models may omit or misunderstand. As a result, one may hypothesize that AI use, as an innovative technique, should boost both efficiency *and* equality. On the other hand, experiences in other sectors suggest that AI carries the risk of codifying and amplifying the inequities embedded in the system today.

To apply this theoretical model to the realities of the mortgage system, we conducted interviews with stakeholders and reviewed the latest reports on AI. The first step in each interview was to determine how key stakeholders defined AI.

Defining AI

In interviews with nearly 50 housing finance stakeholders, we heard many definitions of AI. Notably, there was a lack of universal agreement about what AI actually is (box 1). A central step in assessing AI's impact is to first define it.

For example, AI as a "replication of human decisionmaking" was a less prevalent definition used by representatives from larger financial institution leaders we spoke to. Other stakeholders offered less specific definitions of AI that were indistinguishable from automation more broadly.

BOX 1

Definitions of AI, According to Our Interviewees

Our research included interviews with housing policy stakeholders, including mortgage lenders, consumer advocates, credit union executives, academics, tech equity experts, and fintech leaders. The most common definition we heard focused on AI's ability to mimic human behavior. Another definition focused on leveraging AI models to improve organizational efficiency. As interviewees noted, these definitions may be too broad to be useful:

1. Interviewees focused on AI's ability to mimic human decisionmaking:

"AI is some automated decisionmaking rule or any database decision rule." —Academics studying AI issues

"Any system that tries to replicate similar decisionmaking that humans do—that is, AI and automated statistical models and machine learning models—are instances of AI. For example, computer vision is AI because it tries to visualize similar objects that humans see. First, there is a type of AI that learns from experience, and the second does not need experience to learn; these are called machine learning and zero-shot learning, respectively." —Tech equity expert

"[AI is] the ability to create computer programs that mimic human decisionmaking and algorithms that will allow you to be more precise in understanding the variables and their relationships, akin to machine learning." —Financial services trade association executive

"At the risk of generalization, AI is math that attempts to replicate human intelligence. For example, statistical models attempt to classify risks like humans do. Generative AI attempts to generate creative outputs like text, images, and audio, like humans do. Computer vision attempts to see things like humans do, etc." —Tech equity experts

2. Interviewees working in the industry emphasized AI's ability to improve individual or organizational efficiency:

"AI is technology that helps do something more at scale." —Financial services trade association executive

"[AI is] a tool that is improving efficiency by looking at the data and trying to assist with process and decisionmaking." —Credit union executive

"[AI is] leveraging of technology to create efficiencies within one's operations." —Credit union executive

3. Interviewees who noted how common definitions of AI could be too broad:

"There is no standard definition... [it] can include statistical, mathematical models or automation. A lot of the processes used by creditors that are called AI may not truly fit into that box. We are trying to understand that distinction." —Consumer rights advocates

"Several pieces of state legislation have mentioned 'automated decision systems,' but that definition could be too broad—e.g., a stoplight would fall under that definition. So, 'replicates human decisionmaking' may be too broad." —Mortgage lender

For this report, we develop and offer the following working definitions. We build off Stanford University's John McCarthy, who defines AI as "the science and engineering of making intelligent machines, especially intelligent computer programs."⁴ As IBM states, "It is related to the...task of using computers to understand human intelligence."⁵

In its simplest form, AI combines computer science and robust datasets to enable problem solving. It also encompasses subfields of machine learning, deep learning, and natural language processing that are frequently mentioned in conjunction with AI (Black Knight, n.d.). These disciplines are composed of AI algorithms that seek to create expert systems to make predictions or classifications based on input data. Deep learning uses artificial neural networks that loosely simulate the human brain to identify patterns in data or to predict outcomes. Machine learning focuses on the development of computer algorithms that can learn and perform tasks without specific additional programming. Machine learning, increasingly being employed in algorithms throughout the mortgage industry, is a branch of AI and computer science that focuses on the use of data and algorithms to imitate the way humans learn, gradually improving its accuracy.

It is important to distinguish between automation and AI, which are often conflated. Automation can significantly reduce the human effort needed to complete certain routine tasks, such as employment verification, property appraisals, and certain types of marketing. But automation without AI is limited to the preprogrammed rules implemented by the developer of the automated process and relies on their "thinking" for its effectiveness.

Literature Review

Emerging research has helped shed light on the implications of AI for the mortgage industry. Recent reports show how machine learning models can predict delays in the mortgage origination process (Brahma et al. 2021), improve forecasting of mortgage delinquency (Azhar Ali et al. 2021), increase the efficiency of fraud detection (van Zetten, Ramackers, and Hoos 2022), and improve the efficiency and accuracy of AVMs (Steurer, Hill, and Pfeifer 2021).

At the same time, AI may pose risks to communities of color. Research suggests that machine learning models, which also use alternative data, can result in larger biases in mortgage underwriting by uncovering sophisticated ethnicity and race proxies invisible to statistical models (Zhang, Khalili, and Liu 2019). Additional research confirms this conclusion while illustrating that adopting fairness techniques may result in worse outcomes for both Black and white mortgage applicants (Zou and Khern-am-nuai 2022). In addition, analyses of AVMs suggest that they may produce greater relative error in majority-Black neighborhoods compared with majority-white ones (Neal et al. 2020).

A growing body of research has begun to develop a template for assessing uses of AI in mortgage lending. Through its Tech Equity Initiative, the National Fair Housing Alliance (NFHA) has developed a template that can conduct a critical analysis of an algorithmic system to identify its assumptions and limitations and can produce appropriate recommendations to mitigate consumer fairness and privacy risks (Akinwumi, Rice, and Sharma 2022). Recent research suggests that fair lending analyses of machine learning models must focus on outcomes as opposed to inputs (Gillis 2021). FinRegLab's work also informs the discussion of AI in the credit industry more broadly, with implications for policymakers (Bailey et al. 2023; Cochran et al. 2021).

Theory and research suggest that AI, often using machine learning, can improve efficiency within the mortgage process. But racial bias may undermine its use in practical settings. And federal regulators have acknowledged the potential downside risk.⁶ Some federal agencies have begun to outline expectations of AI and its implications for fairness in mortgage lending.⁷

The rest of this report delves deeper into the applications of AI. The next section describes the dynamics boosting the adoption of AI. We then identify key areas of the mortgage application process and describe the implications of AI for efficiency and racial equity at each stage of production, and we offer a case study of how these principles apply to valuation. Next, we examine how key actors in the mortgage ecosystem use AI. We end with policy recommendations to strengthen racial equity in the mortgage market without sacrificing efficiency.

The Potential of AI in Mortgage Finance

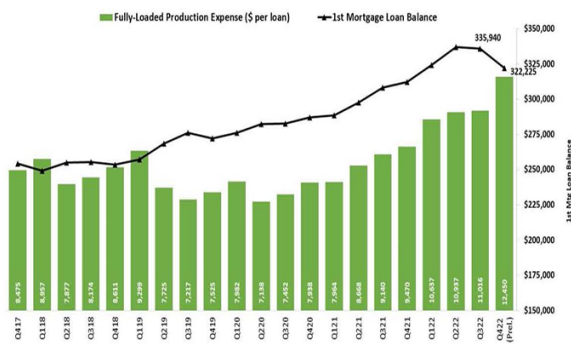
In this section, we discuss AI’s potential to drive greater efficiency and the promises and risks related to equity.

AI and Efficiency

Potential efficiency gains are a major factor driving AI deployment in the mortgage system. Given the vast amounts of information required throughout the mortgage life cycle, and the industry’s cost structure, this is not surprising.

Loan production costs are high and have risen in recent years. Before the recent refinance wave, between the first quarter of 2009 and the first quarter 2019, mortgage production costs rose 250 percent, from \$3,700 to \$9,300 per loan. Over this same period, inflation, measured as the change in the Consumer Price Index, rose only 19 percent. More recently, as of the fourth quarter of 2022, the average cost was \$12,450.⁸

FIGURE 3
Mortgage Loan Production Expenses, Fourth Quarter of 2017 through Fourth Quarter of 2022



Source: Mortgage Bankers Association. Used with permission.

The high fixed costs of mortgage production, excess capacity during periods of low loan production volume, and inefficient loan manufacturing processes all contribute to rising production costs. The increased costs of regulatory compliance with rules such as the 2010 Dodd-Frank Act and the False Claims Act also play a role (Freddie Mac Single-Family 2021; McCargo 2017).

Macroeconomic fluctuations magnify these structural challenges. In response to the broader economic cycle between 2019 and 2021, the Federal Reserve cut the federal funds rate and, with the onset of the COVID-19 pandemic in 2020, further eased policy. In response, mortgage demand soared in 2020 and 2021, driven largely by refinance loans. Refinances often cost less to produce than purchase loans.

But tightening monetary policy over the past two years, largely in response to elevated inflation rates and solid labor market conditions, has boosted mortgage interest rates. Mortgage applications have responded predictably, with applications for purchase loans and refinances both declining, especially refinance applications. These trends have dragged down volumes and profit margins, as lenders take time to shed excess capacity. The mortgage industry's characteristic cycles of ramping up and then ratcheting back based on economic conditions means efficiency is a key focus for most lenders. With lenders looking for ways to establish leaner operating processes to mitigate this risk, increased automation through AI can be one long-term solution.

AI and Racial Inequities in the Housing and Mortgage Markets

By eliminating human bias in decisionmaking, AI also holds great potential to address the racial homeownership gap. It must first, however, overcome the biases and inequities already embedded into the data it analyzes. And policymakers and the mortgage industry must reckon with historical and present-day barriers that lock would-be homebuyers of color out of the market altogether.

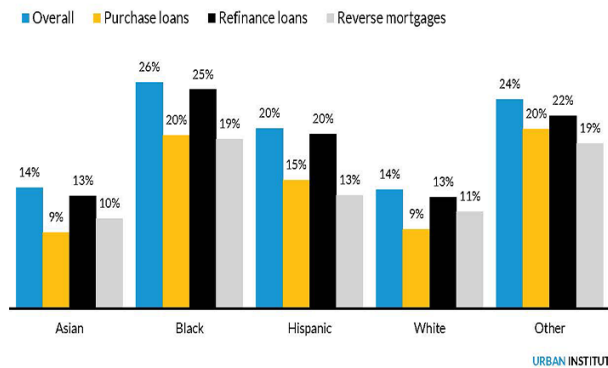
Several factors limit success for people of color throughout the entire mortgage cycle. These disparities are driven by factors that appear to be race neutral but in fact reflect and proxy for the results of a long history of racial discrimination.⁹ This discrimination, rooted in structural racism, has manifested in less access to homeownership and its intergenerational wealth-building effects, in the lack of access to mainstream banking and credit-building mechanisms, in community-level disinvestment, and in broader disparities in education, employment, and income.

Historical discrimination such as redlining limited access to mortgages for people of color, and restrictive racial covenants systematically excluded people of color from owning properties in specific, usually higher-opportunity, areas (Gerken et al. 2023; Santucci 2019). Predatory lending often targeted households of color with mortgages they could not afford, setting them up for foreclosure, wiping away their equity, and ruining their credit profiles in the 2008 foreclosure crisis (Rugh and Massey 2010). The impact of these racist actions, combined with instances of outright destruction of Black property and wealth, have led people of color to mistrust financial institutions.¹⁰

Racial discrimination also exists in other economic sectors and affects a household's ability to qualify for a mortgage. For example, Black workers are often the first fired during an economic recession. In addition, historical experiences, combined with fewer resources channeled to students of color, have limited access to education for people of color, lowering their incomes and making it more difficult to achieve homeownership.¹¹ Black and Hispanic workers typically receive less income, even for the same level of education.

Higher denial rates, particularly for Black and Hispanic applicants, have persisted (figure 4).¹² A larger proportion of mortgage applicants of color have low, or even missing, credit scores. In addition, renters of color are more likely to be cost burdened, increasing their reliance on debt. This results in less savings needed to purchase a home and higher debt-to-income ratios. In addition, inadequacy of property (i.e., collateral) may hinder applicants of color from qualifying for a mortgage, especially when vying for lower-price, more affordable homes or homes in neighborhoods where race-based valuation bias leads to undervaluation.

FIGURE 4
Denial Rates, by Race, Ethnicity, and Application Type



Source: 2021 Home Mortgage Disclosure Act data.
Note: For owner-occupied, primary-residence applications only.

It has also set applicants of color up for subtle forms of discrimination in the mortgage application stage through process and documentation. Increased reliance on easily documentable and easy-to-assess factors, such as a 3-digit FICO score and a consistent salary documented via a W-2, has sidelined applicants who, while fully creditworthy, have a profile that cannot easily feed into an automated, rules-based system.

When they are approved for a home loan, Black and Hispanic homebuyers are charged higher rates and fees for mortgage loans and higher mortgage insurance premiums because of risk-based pricing based on down payment amounts (a marker of personal and family wealth) and credit scores.

Research also points to the potential for racial bias in home appraisals research (Narragon et al. 2021; PAVE 2022). Property valuation has generated significant equity concerns. Evidence indicates that borrowers of color are more likely to have their property undervalued by appraisers relative to white borrowers. The potential for appraisal bias may be at least partly rooted in racial prejudices appraisers hold. Moreover, structural racism affects the values of the properties owned by households of color or in neighborhoods of color. Appraisal error and bias have systematically resulted in either under- or overvaluations (Neal et al. 2020). In addition, industrial sites were often placed near neighborhoods of color, further depressing home values. And a lack of neighborhood investment restrained home values in communities of color. Because of local property tax assessment policies,

homeowners of color also pay more for property taxes as a proportion of their home values (Avenancio-León and Howard 2022). Taken together, homeowners of color accumulate less equity than white homeowners.

These racial inequities are also reflected in racial disparities in refinancing. Homeowners of color are less likely than white homeowners to refinance.⁴³ Low collateral values, high debt positions, and the timing of a home purchase are key factors that reduce refinancing for mortgaged homeowners of color. In turn, this leaves homeowners of color with housing costs higher than those for homeowners who can refinance, costs that reduce households' ability to tap any wealth gained.

In these ways and more, racial disparities in homeownership, once the result of mechanisms that were legal, have been embedded and compounded. Today, the Black-white homeownership gap is worse than it was in 1968 (Goodman and Zhu 2021). In the pursuit of efficiency, will AI embed and reinforce this cycle, or can stakeholders use AI to counteract the racial homeownership gap?

AI and Its Implications for Efficiency and Equity

The tension between equity and efficiency competing for the same resources sets up a trade-off curve that describes the limits of that system across various combinations of efficiency and equity; you can move along the production possibilities curve and adjust its performance to be more efficient or more equitable, but you cannot increase both at the same time.

In the next section, we examine what we learned from interviewees and experts about how AI is being deployed across the mortgage process and the implications for efficiency and equity.

AI throughout the Mortgage Life Cycle

Here, we outline some of the ways AI is increasingly being used in the mortgage process, from marketing all the way through to closing and servicing.

Marketing and Loan Production Selection (Sorting)

AI tools allow lenders to launch hyperindividualized marketing campaigns. This corresponds with recent technology trends in the mortgage process—such as those that track the saving and spending habits of bank account holders—that can be useful for identifying potential clients and tailoring offers to them. AI tools can also be useful for identifying patterns in the data to determine whether existing borrowers are looking to refinance their loan.

Marketers can also use third-party data to target people at certain life stages, such as newly married people or new parents, who may be more likely to apply for a mortgage. AI tools also analyze customer demographics, behaviors, and sentiments about lenders and the broader market and craft more targeted and appealing content. Additionally, AI helps organizations develop content to meet their search engine optimization goals by determining what keywords can attract more customers.

AI can also be used to facilitate interaction with prospective applicants, such as chatbots.¹⁴ As a result, an AI-based chatbot functions as a first-response tool that greets, engages with, and serves customers in a familiar way.

Although AI can allow for more targeted outreach and steering of potential borrowers, lenders using AI need to be careful to ensure their practices do not create disparate treatment or impact. One of the marketing experts we spoke to provide an example of an organization that tried to target “sports enthusiasts.” This ostensibly race-neutral marketing category could disproportionately exclude people of color depending on the sport targeted (e.g., white people make up a larger share of fans of golf, for instance, compared with other sports, such as football or basketball). In such ways, a bias that was historically based on factors such as cultural affinity or, for example, living in neighborhoods without bank branches, can become baked into an algorithm.

Because lending is so highly monitored, lenders need to ensure that their marketing campaigns do not disparately isolate or exclude certain demographics. Moreover, lenders have previously marketed lower-quality products to people of color compared with white counterparts, such as when financial institutions targeted subprime loans to Black and Hispanic households. Lenders should learn from previous actions to minimize the financial risks for communities of color.

In terms of sorting borrowers and matching them to loan products, AI could also play a role in suggesting or steering applicants toward certain loan products, such as an adjustable-rate loan versus a fixed-rate loan, or toward a certain channel, such as Federal Housing Administration (FHA), Veterans Administration (VA), or Fannie Mae- or Freddie Mac-backed loans; bank-held portfolio loans; or private investors.

There is a long and well-documented history of racialized sorting, where borrowers of color have been steered to costlier options, such as loans from finance companies, FHA-backed loans, and subprime loans (Bayer, Ferreira, and Ross 2018). AI could replicate such steering by presenting borrowers options that maximize lender profits, or it could result in borrowers being matched with the product and channel that is optimal for them.

Underwriting and Pricing

Underwriting is a multistep process by which a mortgage lender verifies a potential borrower's income, assets, credit history, debt, and property details to issue final approval on a loan application. The underwriting process usually takes three to six weeks.¹⁵ This risk assessment process often also determines the price charged for the loan.

Underwriting rules can depend on intended mortgage execution (e.g., FHA versus VA versus GSE). Mortgage underwriting tools can use AI models in their process. They can automate distinct tasks, such as document review, and meet other underwriting parameters.¹⁶ And it can be employed to meet several underwriting parameters.

In terms of equity, AI has been cited as a mechanism for potentially improving access to mortgage credit for people of color. Over the past several decades, lenders have relied on traditional credit measures for assessing a potential borrower's ability to qualify. But these traditional measures tend to disproportionately affect people of color and often fail to include predictors of mortgage performance, such as on-time rental payments (Choi et al. 2022). Interestingly, many of these predictors were once relied upon when underwriting was more of a manual, human judgment-centered process.

In reckoning with this problem, lenders and regulators have slowly begun to adopt “alternative” data—such as on-time rental, telecommunications, and utility payments—to extend access to mortgage credit to those who have been historically excluded. Research shows that these data points are both strong predictors of mortgage payments and could reduce racial disparities in mortgage lending¹⁷ with certain precautions (Wu 2022). And this is certainly an area where AI can be of use. First, AI will expand the use of data, which can help a lender improve underwriting outcomes for people of color. And by incorporating alternative data sources, the benefits for applicants of color could be further improved. But second, AI can also better analyze traditional tradeline data. For example, Yolanda McGill of Zest AI notes their fintech company has employed AI and machine learning models to identify potential borrowers’ financial situations and ability to access credit.¹⁸ In this way, AI can help extend mortgage credit to borrowers whose credit scores or histories would have disqualified them under traditional underwriting standards.

By incorporating more data that capture people of color’s ability to qualify for a mortgage, mortgage underwriting can become more equitable. Applying AI and machine learning to typical credit bureau tradeline data to pull more variables from it gives more data that provide a clearer picture of thin-file borrowers. Separately, more data can be obtained using alternative data, which can also be useful to score thin-file borrowers. Obtaining more data points on a borrower even using only normal tradeline data is possible with machine learning and can help lenders reach creditworthy borrowers they would have overlooked. But including traditional forms of information for assessing the ability to repay a mortgage into an AI-based model has proven inferior to combining alternative measures of credit, such as rental payments, into AI-based models.¹⁹

Some of the most promising AI-based underwriting models are also the most controversial because they incorporate race or ethnicity into the underwriting model. Including data on race allows these models to produce more equitable outcomes for people of color, reducing their likelihood of denial relative to white applicants (Gillis 2021). But by including explicit data on protected classes, they could conflict with current federal law unless they are part of a special purpose credit program under the Equal Credit Opportunity Act (ECOA).

But there may be methods for incorporating protected class status. An interviewee working at a fintech company noted that protected class information ought to be used during model training, such that the algorithm can learn to set the relative weights on variables in ways that make its decisions more sensitive to the credit characteristics of protected class applicants. That is, consciousness of protected status can be used so that models are built to be fair from the beginning and thus achieve more equitable outcomes without the need for protected status information to be used as an input at

decision time.²⁰ But the federal finance regulators have never issued guidance on whether this use of protected class must be reconciled with the ECOA. Many fintech companies, often the early-stage creators or adopters in an innovation system, argue that mitigating racial disparities requires revisiting prohibitions on the use of protected status, which may have made sense at one point but may have outlived their usefulness. At the same time, these fintech companies continue to look for ways to build credit strategies that are sensitive to protected classes while keeping their products aligned with federal regulations and increasing the likelihood that borrowers and lenders will adopt them. But the impact may not diverge from models seeking to avoid regulatory violations but still produce disparate outcomes. This challenge raises the need for other race-conscious approaches that can reduce racial inequities in underwriting.

Employment, Income, and Asset Verification

Critical to the mortgage process is income and asset verification. Generally, the borrower submits W-2 forms, bank statements, tax returns, and pay stubs to verify income and employment. Some lenders ask for contracts to verify the borrower's employment status. Similarly, lenders use borrowers' bank and investment account records to verify their assets.²¹ These verifications help lenders evaluate whether the borrower can pay their monthly mortgage reliably.

Income and asset documentation may consist of physical papers that an underwriter needs to scan, but AI, specifically machine learning processes, can efficiently extract and verify all the information in these documents. These applications can scan huge volumes of case files, extract relevant information, and pass it on to the underwriter in summary form.

Use of AI in employment and income verification is more straightforward and presents fewer equity concerns than other parts of mortgage underwriting. There also exists the potential to advance equity, as improvements in these verification technologies may make it easier for lenders to assess the ability to qualify among people with multiple or inconsistent income streams, such as gig workers, who are more likely to be people of color and whose income has traditionally been difficult to assess.

Appraisals

In the underwriting process, the determination of a home's value is critical. First, it gives the underwriter a sense of whether the price agreed upon by the buyer and seller is an accurate

representation of the home's true value. Second, it provides the underwriter a sense of the risk being taken if the loan is originated.

A home appraisal is a process through which a real estate appraiser determines a home's fair market value. Separately, collateral issues—such as an appraisal that comes in below the sales price or mortgage requested—could undermine the mortgage application and result in the potential borrower being denied a mortgage. Appraisals are also often used by local governments to determine property taxes.

Traditionally, a human appraiser estimated property values. A traditional appraisal also involves a visit to inspect and measure the home. From this in-person visit, an appraiser can assign a value by comparing the home's features, age, and condition with similar homes in the area and what they sold for.

A more efficient and faster process of valuing homes has emerged: automated valuation models. AVMs incorporate large amounts of historical data on home sales to estimate the subject property's value. In an increasing share of loan origination transactions, AVMs are also accepted instead of appraisals. For GSE mortgage originations, AVM use is largely confined to refinance loans with loan-to-value (LTV) ratios below 80 percent. But they are used for some purchase originations where the LTV ratio is 90 percent or higher.

AVMs can determine property values in seconds. This facilitates the mortgage industry by increasing refinances and helping consumers discover the price of a home they may wish to purchase.²² And during the pandemic, appraisal waivers helped sustain the industry when in-person activities were discouraged.

AVMs incorporate AI to quicken estimates and to increase accuracy. And with machine learning, the AVM can use large amounts of data more efficiently. Before the use of AI, creating a highly accurate AVM involved developing a base model. But that statistical model would not have the same level of accuracy in all markets. And creating multiple models for accurate use in different geographies unwinds the benefits of automation.

Incorporating machine learning helps AVMs perform complex calculations that more accurately capture how humans think through neural networking and cloud computing power. As a result, one model can be written that can adapt to all the nuances of different markets.

AVMs are increasingly being used to remove the human potential for bias. AI and machine learning can improve the accuracy of AVMs through more flexible modeling and can better absorb and interpret data. In addition, AI can also be used to test an AVM for its potential for bias. The results of a machine

learning-based testing model should provide for more precise estimates of the racial inequities in property valuation.

But in addition to the algorithms in place, an AVM's performance also relies on the data it is fed. And input data that are racially biased will produce racially biased results. This is a huge issue for AVM developers and users because of the long-standing history of racism and racial disparities in appraisal values. The sales comparison approach used by human appraisers—in which appraisal values are based in part on those of comparable properties—has baked in racial inequity by pulling forward historic undervaluation over time. As a result, an AVM relying on these data will likely experience the same problem and can perpetuate racial discrimination.

Fintech companies are innovating in the AVM space. They are now able to design a model, based on AI, that can be used to estimate a home's value more quickly and using more data drawn from a broader range of sources. But the AVM is still producing an estimate of the property focused on what the housing market will bear. One interviewee pointed out that an AVM does not try to incorporate the impacts of systemic racism on a home's value; it just knows what others have paid for similar properties.²³ Comparing purchases on similar properties and factoring in the broader context that undervalues homes in neighborhoods of color are two different approaches.

Charu Singh of Just Value Inc. leads a property technology and financial solutions firm focused on undervaluations. Singh points out that "AI is not the first nor most impactful tool to solve" systemic appraisal bias, noting that the housing industry and tech experts "have not yet built the foundation to allow AI models to reflect anything back to us other than the bias already present in the data and market."²⁴ Singh also points out that society's actions have driven systemic undervaluation, that AI can amplify those actions, and that communities of color must "lead the way" to identify equitable and transparent AI models.²⁵

Although AI use in AVMs can produce greater accuracy, it will be important to further define what the industry means by accuracy. If accuracy corresponds with a home's ultimate sale price, a comparable-sales method improved with AI-based AVMs will reduce racial inequities.

Fraud Detection

There are several types of fraud, and it may be committed by borrowers, lenders, appraisers, or people involved in the real estate transaction. Income fraud risk, where an applicant falsifies their income, is a critical concern and can be perpetrated via doctored pay stubs or W-2s to longer-term falsification

schemes involving using false information with "seasoned" bank account information. Meanwhile, asset rental fraud occurs when loan applicants borrow or rent other people's assets to make themselves appear more qualified for mortgage financing. Equity skimming occurs when investors use straw buyers, or someone who purchases property on behalf of another person.

Homeowners may also fall victim to foreclosure relief scams. In this type of mortgage fraud, homeowners who are at risk of defaulting on their loans or whose homes are in foreclosure are misled into believing they can save their homes by putting the property in the name of a third-party investor. Scammers often use false or stolen identities to commit mortgage fraud. This takes place when the scammer obtains financing by using an unknowing victim's financial information, including Social Security numbers, stolen pay stubs, and falsified employment verification forms, thereby obtaining a fraudulent mortgage on a property they do not own or occupy.

False appraisals are another common way scammers commit mortgage fraud. Appraisers may commit the appraisal fraud on their own or with the help of other professionals, including a builder or a mortgage banker.

There are several ways to avoid mortgage fraud. They include using an attorney to thoroughly review all legal paperwork; checking the references and referrals of all participating parties, including real estate brokers and loan officers; researching and verifying the property's title history; reviewing all final loan documents to ensure all information is accurate; and researching and reviewing property tax assessments to verify the actual assessed value.

Automated computer systems using statistical modeling and analysis to fight online fraud are mature. The software would review previous cases and then look for statistical relationships among many data variables and then use these patterns to flag suspect cases for review by internal antifraud teams.

Increasingly, AI is being implemented in these automated systems for greater speed and accuracy. For example, Resistant.AI trains the application, and in turn, the AI system processes information, receives ongoing corrections, and learns from them.²⁶ By training the application, an AI process uses historical data to help uncover important characteristics and features of a particular dataset. As decision logic has evolved, complex pattern recognition has become a requirement, and employing machine learning algorithms can accelerate the use of data. Through machine learning techniques, a lender can better understand the probability that a mortgage application will have fraudulent information and the part of the application that is likely to be misrepresented.

For example, machine learning–based fraud detection software would be able to improve risk predictions by relying on large volumes of diverse, granular, and high-quality data. In addition, machine learning can quantify any potential fraud risk, capturing actual and nonfraud evidence. A consistent and ongoing feedback loop allows the technology to continuously learn, creating stronger predictions while identifying new and emerging fraud behavior trends.

Machine learning can improve on the inefficiency associated with an underwriter manually reviewing applications. Relying on machine learning technology can help underwriters decide whether a loan needs further investigation for potential fraud.

The potential for racial bias may exist in fraud detection software as well. For example, if the training data used to develop the AI-based algorithm are biased, the model could produce biased outputs, thinking that a person of color is more likely to commit fraud. In turn, applicants of color may be more likely to be turned down for loans.

To date, fraud detection in mortgage lending is an area for potential research. There is little empirical analysis describing fraud detection models and their potential for racial bias. This research would bring greater transparency to an opaque part of the mortgage process.

Servicing

In loan servicing, the lender collects the principal, interest, and escrow payments from the borrower over the life of the loan and passes those payments to the loan funders (investors). There are two distinct aspects of mortgage servicing: first, managing performing loans is a high-scale, high-efficiency process that is fairly automated; second, managing nonperforming loans, from their first delinquency through reperformance, restructuring, foreclosure, or some other resolution, is more complex, labor intensive, and costly. Mortgage lenders can streamline loan servicing workflows using AI-powered support automation platforms.²⁷ These tools alleviate repetitive manual processes and provide solutions to lenders, underwriters, and borrowers. By speeding up loan servicing processes, lenders can deliver better, more tailored customer service to borrowers.

But nonperforming loan servicing has been less automated than earlier stages of the mortgage life cycle, such as underwriting. And the lack of robust data suggests that AI may be less useful in mortgage servicing because there is less information to train the model on. This lack of data often extends to race and ethnicity as well. Datasets such as those collected under the Home Mortgage Disclosure Act

provide this information at origination, but there is no requirement that race or ethnicity Home Mortgage Disclosure Act data be maintained as the loan moves to servicing.

Federal Regulations Promoting Equity in the Mortgage Process

Federal fair housing laws seek to eliminate discrimination in the mortgage industry, and these all have implications for AI. Most notably, the ECOA, enacted in 1974, prohibits discrimination in any aspect of a credit transaction on several factors, known as protected classes (Federal Reserve, n.d.).

Under the ECOA, it is illegal for creditors to discriminate against applicants based on race, sex, age, national origin, or marital status or because an applicant receives public assistance. Creditors are also prohibited from asking about marital status (with some exceptions) and from asking an applicant whether they plan to have children or additional children, though they can ask about the number, ages, and financial obligations relating to all existing children (Wu 2021).

Within the Civil Rights Act, passed in 1968, Titles VIII and IX are commonly known as the Fair Housing Act because they established new antidiscrimination laws regarding housing. The 1968 Fair Housing Act expanded on previous laws by prohibiting discrimination concerning the sale, rental, and financing of housing based on race or color, religion, national origin, sex, family status, or disability (Federal Reserve, n.d.).

The prohibitions codified by the Fair Housing Act cover discrimination in all aspects of residential real estate-related transactions, including purchasing real estate loans or appraising residential real estate (Federal Reserve, n.d.).

Because both the Fair Housing Act and the ECOA apply to mortgage lending, lenders may not discriminate in mortgage lending on the basis of any of the prohibited factors listed. Under both laws, the following activities are illegal if performed on the basis of a protected class:

- failing to provide information or services or providing different information or services relating to any aspect of the lending process, including credit availability, application procedures, and lending standards
- discouraging or selectively encouraging applicants, with respect to inquiries about or applications for credit

- refusing to extend credit or using different standards in determining whether to extend credit
- varying the terms of credit offered, including the amount, interest rate, duration, and type of loan
- using different standards to evaluate collateral
- treating a borrower differently in servicing a loan or invoking default remedies
- using different standards for pooling or packaging a loan in the secondary market (Federal Reserve, n.d.)

Regulation around AI is still in its early stages, and much of the discussion in the mortgage space has centered more heavily on AI's use in underwriting and appraisals. Additional oversight is necessary to ensure that AI tools are not driving racial disparities in access to mortgage credit throughout the process.

How Can AI Embed Discrimination into the Mortgage System?

In the previous section, we covered some of the equity issues raised by AI use at different stages of the mortgage process. But more broadly speaking, our research highlighted several ways AI can embed discrimination or systemic racism in the mortgage system across many different processes: through biased data, biased models, and biased standards. This section uses an example of AVM error and equity to illustrate how bias affects racial disparities in the mortgage system.

Biased Data

Across our interviews, many industry stakeholders underscored the fact that any AI or machine learning tool is only as good as the data that go into it. Any data resulting from racist practices will inform inequitable outcomes under newer AI or machine learning models.

This issue is not exclusive to the mortgage industry; equity-minded people are grappling with it across AI's different uses. One researcher highlighted an issue with AI-based algorithms that set bail in the criminal justice system. People of color, especially Black people, are overrepresented in the criminal justice system and are more likely to have been arrested, despite committing crimes at similar rates to white people. This overrepresentation in the data leads to higher bail being set for Black people, despite the algorithm being ostensibly race neutral.

Similarly, the history of racism and racial disparities in the credit and housing finance systems creates the same problem for AI users relying on data from these industries. Systemic and individual appraiser racism have led to undervalued homes in many communities of color. So an AI-based AVM may reduce racial bias to some degree by removing the human appraiser, but it can still perpetuate systemic racism by, for example, using undervalued homes as comparisons when assigning a value to a property in a neighborhood of color.

Another issue with AVMs relying on biased data is that their efficiency means they have the power not only to reinforce systemic racism but to amplify it as well. AI and machine learning models have the capacity to do the work of a human but faster, meaning that lenders can use them to significantly increase loan volume. And that pace will only increase as AI technology improves. But the cost of this

type of efficiency is that if the industry fails to understand or grapple with the racism baked into much of the data, they will simply further entrench that racism and worsen existing disparities.

Moreover, the ability of AI to reduce racial bias could improve if full demographic information was available for every loan, because this information would allow lenders to test, after the fact, for racial bias in lending-related systems. Right now, lenders do not always have full demographic information for every loan because it is an optional part of the loan application.

In addition, there are privacy concerns for consumers whose data are used in AI modeling. For instance, advancing the use of cash flow data in underwriting requires increased use of consumer-permissioned data, particularly from bank accounts. But if the use of this type of data is not regulated and monitored, there are significant data privacy concerns for consumers. And these concerns may be necessarily heightened for consumers of color, who have an understandably greater distrust of the financial system because of past and ongoing harms.

Biased Models

Another racial equity issue that can arise with AI use is whether the models themselves are biased. Given the restrictions around fair lending, it is unlikely AI developers would build models that disproportionately harm people of color. But the ways some AI models learn and evolve on their own mean that if developers and users are not careful about the parameters they set, AI models can still create inequitable outcomes by race.

Fair lending standards are understood to prevent lenders from using race (or other protected statuses) as inputs in their lending algorithms.²⁸ But many other variables can serve as proxies for race, especially as data quality and granularity improve.

Perhaps a lender is hoping to market to potential borrowers, and the lender trains an AI-based marketing tool based on past borrowers' characteristics. Perhaps these past borrowers skew disproportionately white compared with the population, as white households are more likely to achieve homeownership. Even if the data on which the AI is trained do not include information on race, the model, as it learns, may identify proxies for race that result in a preference for potential white borrowers over borrowers of color, thereby amplifying the impact of biased data.

This is one reason some stakeholders advocate for including protected statuses as model inputs. To correct for historical discrimination, lenders could use race in their models as a means of debiasing their

outputs. But for this solution to avoid being harmful will require strong regulatory supervision of lenders' use of race variables.

Another risk associated with AI models is a lack of transparency and explainability, sometimes referred to as a "black box." The increasing complexity of AI models may mean that humans have a hard time understanding what the model is doing or the relationship between the variables that go into it and the results it generates. This is especially true for AI-based systems because they can generate different decisions even for identical circumstances. Under the ECOA, lenders are required to be able to explain the reasons for denial. But this may become increasingly complicated as the models become increasingly complicated. A lack of transparency could harm consumers and reduce their ability to dispute what they perceive as an unfair or arbitrary decision. This could disproportionately affect borrowers of color, who are more likely to be denied a mortgage.

Finally, because some AI models are designed to learn, there can be an inherent risk of model drift²⁹ that needs to be closely monitored. A model that begins by offering racially equitable evaluations and outcomes can change to become inequitable or discriminatory. A lack of careful monitoring poses a risk to consumers of color.

Discriminatory and Inequitable Standards

Beyond more specific AI-related concerns, developers and users also need to consider the end goal of their AI use. Throughout the mortgage process, certain standards applied to applicants and borrowers have been shown to have racially inequitable outcomes. Given the efficiency gains and amplification power of AI, mortgage industry stakeholders should be thoughtful regarding where they plug in AI tools.

A good example of the issue of discriminatory or inequitable standards are traditional debt-to-income ratio and credit score requirements in the underwriting process.

Recently, the mortgage industry has begun to reckon with its history of systemic racism and how that racism affects what may seem, at first glance, to be reasonable markers of one's ability to qualify for a mortgage. Using automation or AI to simply speed up this process or extract income and debt information from application documents would be an example of using AI to increase efficiency but would perpetuate an inequitable standard.

The inequities baked into debt-to-income ratios and credit scores are one reason the use of alternative data in underwriting has gained steam, because AI-driven models can process more complex and nonstandardized information than rules-based systems currently in use. If it accelerates that trend, AI will not be perpetuating a racially biased status quo but will be used to extend access to credit to people who are disadvantaged by traditional metrics. Even so, underlying even this improvement is the fact that even though people of color have lower incomes and greater job instability, they may still be less likely to make the on-time rent payments that can sometimes substitute for credit scores. In this regard, AI can be used both to make progress on advancing racial equity in the mortgage space and to encourage the industry to think deeper about underlying disparities and innovative ways to remedy them.

In the next section, we extend the discussion of AVMs with a case study to illustrate how these biases can be baked into AI and how steps can be taken using AI to improve both accuracy and equity. We discuss how AI balances the accuracy-equity trade-off in AVMs and how models that focus on being more precise might skew to exclude or undermine communities of color.

Case Study: Integrating Equality into the Design of Machine Learning Algorithms by Balancing Accuracy and Equity in AVMs

To demonstrate how an AI-based innovation can increase or mitigate bias, we will delve into AVMs. Using an AVM instead of a traditional appraisal can yield significant efficiencies and removes bias risk arising from human judgment. Nevertheless, prior research has found AVMs to have less accuracy, as a percentage of value, in neighborhoods of color (Neal et al. 2020). To avoid embedding valuation bias in AVMs, the federal government has recommended rulemaking to address potential bias in AVMs (PAVE 2022). As building AI into AVMs can improve their performance and efficiency and expand their use, this case study sheds light on how regulators can think about setting performance standards for equity in this product.³⁰

Measuring Accuracy

An AVM produces estimates that can differ from the property's actual market value because of limitations of the algorithms, omitted variables, and other factors. To measure the accuracy of an AVM's estimates, the forecast standard deviation (FSD) is used to represent the probability that a particular

AVM estimate falls within a statistical range of the subject property's actual market value. For example, if an AVM produces an estimated value of \$200,000 with an FSD of 5 percent, there is a 68 percent probability (1 standard deviation) that the subject property's actual market value falls between \$190,000 and \$210,000. The lower the FSD, the more accurate an AVM is.

Measuring Equity

The simplest measure of fairness in machine learning is statistical parity. Because we are concerned about algorithmic discrimination against people of color in the form of less accuracy, we want a model that results in the same average FSD for both people of color and white people.

But the goal of statistical parity in outcomes overlooks disparities in inputs, factors that affect the accuracy an AVM's performance. Previous findings have shown that properties in neighborhoods with more homogeneous properties, fewer distressed sales, and higher median household incomes tend to have greater AVM accuracy—that is, lower FSDs. Properties owned by white households are more likely to be in neighborhoods with factors contributing to lower FSDs (Neal et al. 2020). This means that if we managed to find a perfect AVM—that is, one that takes the data input of any property, regardless the neighborhood's housing characteristics, and always minimizes the FSD—statistical parity forces us to make difficult trade-offs. This is because the actual prediction difficulty of the two groups differs, but fairness requires us to achieve equal FSDs.

Here again, we face a trade-off curve.

Fairness versus Accuracy

To illustrate the trade-off between fairness and accuracy in the context of AVMs, we can use a simple example. Assume two models: model 1 has a higher accuracy overall (a lower FSD score), but it has a lower FSD score in majority-white neighborhoods than in majority-Black neighborhoods. Model 2 has a higher FSD score but still below the satisfactory threshold, and the accuracy of its outcomes is less correlated with neighborhood racial composition.

If our goal is pure accuracy, we would select the model with the lowest FSD, which is model 1, but then we violate the fairness requirement. Moving from model 1 to model 2 indicates that improving fairness will degrade accuracy, and vice versa.

But why not build separate models for the two types of neighborhoods? This more complex hybrid model would increase both fairness and accuracy. This hybrid model, though, now explicitly uses race as an input, as the first step is to pick which model to use based on race. The example illustrates several points raised in the preceding section. First, hidden bias in the data is a critical source of AVM error. The training data used in the AVM algorithms (i.e., home prices of comparable sales) can contain all kinds of hidden bias, and building complex models from such data can both amplify these biases and introduce new ones. Second, the algorithms employed by the AVM may contribute to AVM error. Without specifying a nondiscriminatory fairness factor in the algorithms, the optimal AVMs would maximize the predictive accuracy. When maximizing accuracy across multiple racial groups, an algorithm will naturally maximize for the majority population, at the expense of others, as the majority contributes more to the model's accuracy.

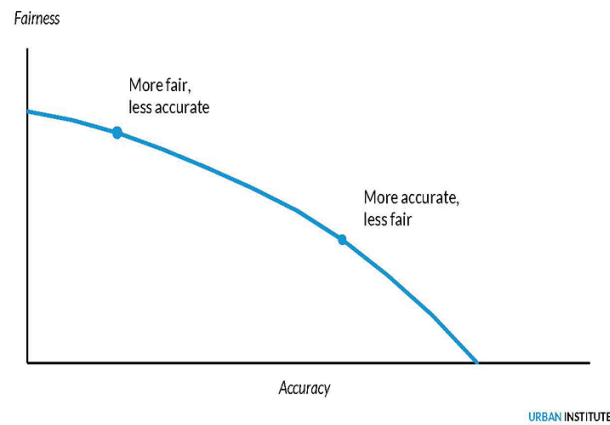
Thus, building the optimal model for predictive accuracy without specifying a fairness criterion often results in demonstrable unfairness to some applicants and borrowers. The sensible response to this tension between accuracy and fairness is to acknowledge it and then directly quantify and manage it so that we can figure out a solution to reduce it. This means including racial and ethnic information in the AVMs' inputs can lead to more equitable outcomes.

How Can We Optimize This Accuracy-Fairness Trade-Off?

FIRST, ESTABLISH THE TRADE-OFF FRONTIER OF ACCURACY AND FAIRNESS ACROSS VARIOUS MODELS

For this first step, we measure the FSD and the unfairness scores of each model's output. Each point in figure 5 corresponds to a different AVM. Here, we have drawn a curve connecting a set of models that outperform the other models. We should consider all models on the curve to be reasonable candidates for optimal performance. Moving from model 1 to model 2, we are gaining a greater degree of fairness at the cost of reducing accuracy, as we are getting a greater error score. This yields the frontier of "best" performance between error and unfairness.

FIGURE 5
The Trade-Off Frontier of Accuracy and Fairness



URBAN INSTITUTE

NEXT, ENCODE FAIRNESS INTO ETHICAL MACHINE LEARNING ALGORITHMS

The trade-off frontier is silent about which point we should choose along the frontier, as that is a matter of judgment about the relative importance of accuracy and fairness. We could assign a single numerical objective that takes a weighted combination of the error and unfairness scores. The weights ascribe a "penalty" to losing accuracy or fairness.

Interestingly, this optimization decisionmaking process relies on human judgement, including how to define the notion of fairness and equity in the first place, which groups of individuals to protect, and how to assign penalty weights in consideration of the relative importance of accuracy and fairness.

The accuracy-fairness trade-offs illustrate the complexities and the nuanced nature of algorithmic fairness. Therefore, the design of ethical algorithms becomes an exercise in understanding which fairness criteria are most relevant in a context and then balancing these considerations.

Without intentionally addressing these equity concerns, when such models become the basis for widely deployed services such as appraisal tools, the bias can be further propagated and even amplified by their reach and scale. It may be practically impossible to solve the bias in data input as the data, such as income and home prices, already bake in the impact of historic racism. But it is possible to think about

how to integrate fairness and racial equity into the design of ethical algorithms and in the testing of algorithms for bias.

We believe federal regulators are positioned to set these governing principles because of their authority and because of their ability to set consistent standards on a large scale. In other words, to think about an algorithm for building machine learning models that can avoid amplifying or that can reduce racial bias requires dedicated governmental resources focused on transparently assessing trade-offs that can emerge in principled rulemaking.


The AI Ecosystem

In the previous section, we assessed AI use and its implications for both efficiency and racial equity through the mortgage application process. In this section, we look at AI adoption from the point of view of various entities involved throughout the process. Many key actors drive the pace and forms of adoption, including fintech companies, lenders, researchers, community advocates, the federal government, and the GSEs.

Technology Adoption Life Cycle

Broadly speaking, our qualitative evidence suggests the diffusion of AI parallels the technology adoption life cycle. According to this framework, a small proportion of AI users are innovators and early adopters. In our research, fintech companies and the largest lenders appeared to belong to these two groups, respectively.

FIGURE 6
Diffusion of Innovation Curve



URBAN INSTITUTE

Source: Adapted from Matt S. Smith, "Models for Predicting the Future: Geogrey Moore's 'Crossing the Chasm,'" SmithHouse blog, February 21, 2018, <https://smithhousedesign.com/models-predicting-future-geoffrey-moores-crossing-chasm/>.

But there appeared to be a gulf between early adopters and other adopters, the distribution to the right of the chasm in figure 6. This may be attributable to a lack of clarity on issues documenting an AI process or regulations governing the appropriate AI-based decisionmaking. But it may also reflect the

high up-front costs associated with adopting AI. Below, we look at the state of adoption within each category.

Fintech Companies

Fintech companies are a key source of innovation. They often develop financial technology to improve and automate the delivery and use of financial services. That said, the mortgage industry is heavily regulated. And at times, the diffusion of innovation can be restrained by the regulatory environment.

Fintech companies are pushing the bounds of the evolution of AI as it pertains to underwriting. One interviewee who works with fintech companies noted that lenders have begun “moving from conventional credit scoring techniques to more sophisticated ones powered by alternative data and AI.”³¹ Additionally, Zest AI, a mission-oriented fintech company, works with banks, fintech companies, and credit unions to adapt ethical AI practices into their underwriting, pricing, and other decisionmaking systems.³²

Mortgage Lenders

The high cost of adopting AI partly reflects the need for new technological infrastructure that can analyze AI-based algorithms. In addition, labor costs are elevated because of a need to hire data scientists and other technology specialists who can build and execute AI-based models. Further, regulatory costs may also be high, reflecting a degree of uncertainty on sanctioned uses of AI.

Given the financial and regulatory costs of AI, smaller lenders appear further behind in AI adoption. But the lack of AI penetration among these lenders may also reflect their relationship banking style of operating that is geared toward achieving their mission. Moreover, some MDIs expressed skepticism of AI's benefits for communities of color as well as the history of structural racism.

One interviewee with experience at a CDFI noted that larger financial institutions probably incorporated AI more often than smaller, community-based institutions. Additionally, the interviewee noted that few such institutions look beyond traditional models and prefer incorporating a hands-on approach. “These lenders are not going to implement anything that walks them backwards on their mission, even if it improves their efficiency.”³³ Such financial institutions may be more comfortable adopting AI use for income verification, particularly institutions that make easy approval decisions. This would free up their resources for more challenging cases that require a hands-on approach.³⁴

Moreover, mission-oriented financial institutions can serve lower-income communities and communities of color through high-touch activities as a way to separate themselves from larger lenders using AI to drive efficiency and grow markets.³⁵ Additionally, Nicole Elam of the National Bankers Association highlights the importance of MDIs improving their tech efficiency while maintaining relationship banking models.³⁶ Elam notes that MDIs have not found a scalable relationship-driven way to incorporate new technologies into their initiatives to help community members become mortgage ready.

And there may be a high degree of skepticism, rooted in the history of structural racism that undermined communities of color and MDIs. Some MDIs' skepticism may reflect the view that AI mimics human behavior, and human behaviors are affected by structural racism. So MDI lenders hope to see more transparency in how AI is developed, where the data come from, how the data are being refined, and what is the racial equity lens being built in.³⁷

As AI's use expands across the financial landscape, some mission-oriented lenders have begun exploring AI options. Gary Perez, president and CEO of USC Credit Union, which is deploying AI in small-dollar, nonmortgage consumer lending, notes that incorporating AI-assisted credit underwriting services will enable the credit union to better support the majority-Black and majority-Hispanic neighborhoods surrounding the University of Southern California. Perez notes that the small organization's community charter expansion, as well as its CDFI and MDI designations, are new, thus limiting their ability to do empirically based relationship banking. As a result, USC Credit Union hopes to use AI to responsibly expand access to community members who would otherwise be denied.³⁸

Civil and Consumer Rights Groups

Civil and consumer rights groups generally expressed skepticism about whether AI will be a force for racial equity. Keisha Deonarine of the NAACP pointed out, "The credit system was not created with [them] in mind. And this raises the importance of humans assessing the information provided to see whether there are any potential risks."³⁹

Chi Chi Wu and Odette Williamson of the National Consumer Law Center note, "AI has systematically undermined people of color" in other areas of the economy, such that it breeds a degree of skepticism and fear. These feelings "could also be heightened by skepticism people of color have with the financial system in the past," and Wu and Williamson note that companies need to develop transparent, nondiscriminatory systems to build trust with communities of color.⁴⁰ They are cautious of

new models, indicating that nothing prevents a newer model or approach from being more discriminatory. But they also feel AI could improve the current system.⁴¹

NFHA is playing a leading role in tech equity. Critically, NFHA has developed an approach for auditing algorithms (Akinwumi, Rice, and Sharma 2022). NFHA posits that many issues AI is creating are not generated by AI itself. Rather, AI mirrors issues in the underlying data. NFHA's Tech Equity Initiative is grounded in the hope that intentional AI design could help people of color gain access to housing and economic opportunities, and if businesses can see how AI can help them achieve equity while still achieving profit targets, they are more likely to be on board.⁴²

Through this work, NFHA has identified five key goals⁴³ to address technological bias:

1. developing solutions for removing bias from the technologies that shape our lives
2. increasing transparency and explainability for AI tools
3. advancing research to help reduce bias in tech
4. developing policies that promote more effective oversight for AI tools
5. supporting efforts to increase diversity, equity, and inclusion in the tech field

Incorporating these goals has led to important leadership on topics intersecting with racial equity and racial bias. Importantly, NFHA has helped deepen the industry's understanding of machine learning in credit underwriting models, algorithmic bias in housing-related transactions, and a fair lending policy agenda for federal financial regulators (Akinwumi et al. 2021).⁴⁴ NFHA is also helping guide federal policy related to key areas of mortgage lending, including AVMs, which increasingly rely on AI.

Researchers

Systematic analysis of AI-based systems is critical for establishing its potential and its limits. Given their independence, researchers can provide added transparency, particularly for a black-box technology in the private sector.

Research has found mixed results on whether automated underwriting discriminates. Some research has suggested that fintech algorithms discriminate in mortgage underwriting, but "40% less than face-to-face lenders" (Bartlett et al. 2020, 4). Updates to this research find the rate differences between nonfintech and fintech lending to be negligible (Bartlett et al. 2022). AI can be used to audit the findings from a standard statistical model (Zhu, Neal, and Young 2022). In this case, when using an

AI-based method, the extent of error in AVMs was modestly lower than the results produced by an ordinary least squares regression (Zhu, Neal, and Young 2022).

FinRegLab has also produced reports targeting policymaking in AI in financial services broadly. The organization started with a market context report to try to understand what kinds of lenders are using machine learning, what kinds of challenges they need to solve, and what policy questions require more clarity around regulatory guidance (Cochran et al. 2021). Bailey and coauthors (2023) provided important information to the industry on explainability and fairness in machine learning for credit underwriting. Additionally, FinRegLab evaluated tools provided by seven companies that provide platforms and services to help lenders develop and manage machine learning models responsibly, including for bias-related risk. They hope this work sheds light on the changes needed to support the fair and responsible use of machine learning models.⁴⁵

The Secondary Market and Guarantors

One major force is the secondary market. Many lenders sell a significant portion of their loans to Fannie Mae and Freddie Mac to boost lender capacity, increase liquidity, and offload some risks. Loans securitized by the GSEs are a sizeable portion of the one-to-four-family residential market (Goodman et al. 2023). But to sell to the GSEs, a lender must meet Fannie Mae and Freddie Mac's underwriting standards regarding the types of loans they will buy.

Given the role the GSEs play in the secondary market, their use of AI could significantly affect AI's adoption throughout the rest of the industry. Per a Federal Housing Finance Agency (FHFA) report, both GSEs are using a "cautious approach" to adopting artificial intelligence and machine learning and remain in the early stages of adoption. And although this report suggests that neither GSE uses AI directly in underwriting, Freddie Mac plans to incorporate an AI credit scoring solution in its automated underwriting system (OIG 2022).

In addition to the FHFA and the GSEs, the US Department of Housing and Urban Development (HUD) plays an important role in the mortgage process. And given the disproportionate concentration of borrowers of color in FHA mortgages, HUD's influence on AI could have a stronger short-term impact on outcomes for applicants of color.

The FHA has implemented its Technology Open to Approved Lenders Mortgage Scorecard algorithm for its automated underwriting. But HUD may still have steps to take to modernize its systems.⁴⁶ Use of AI will also require internal updates to their systems. Updating with AI-based systems,

however, will affect lenders because FHA lenders would have to ensure their systems are up to date to communicate with theirs. And HUD plans to assess whether loans it supports through all of its programs violate the Fair Housing Act.⁴⁷

In sum, the adoption of AI by secondary market agencies is moving slower than the development of AI tools and technology. Caution is not a bad thing. But this slow-moving process also means lenders are wary about using AI tools that could advance equity—such as underwriting tools that go beyond traditional credit scoring—because the loans still have to be approved by and managed by rules these agencies set. And lenders may wait until these rules are developed or sufficiently clarified before expanding AI adoption. Because more than 80 percent of mortgages are backed by the federal government or the federally chartered GSEs, they set the rules of engagement. The FHA and the GSEs could encourage more lenders to adopt equity-advancing AI and ensure that the necessary oversight infrastructure exists to prevent the entrenching or exacerbating of racial inequality through these tools.

Federal Regulators

Through the Fair Housing Act and the ECOA, the federal government has issued rules that protect historically marginalized communities, including people of color, and uses its authority to enforce these rules. In addition, the federal government, through both public commentary and the use of its own systems, directly influences the role of AI in the mortgage application process. The Biden administration and financial regulators have taken steps to promote transparent, cautious use of AI that promotes both efficiency and equality in the housing industry.

Partly rooted in the administration's growing interest in AI, agencies such as the Consumer Finance Protection Bureau (CFPB) are identifying the tools needed to support sustainable and equitable AI use.⁴⁸ To further its understanding of the risks associated with AI and to craft better policies, the CFPB has called for pilots that illustrate the success of automation and where AI guidance will be needed.⁴⁹

The federal government can also provide guidance to the industry by outlining its thinking on AI more broadly. The most notable releases are the AI Bill of Rights and President Biden's executive order on October 30, 2023, which identified key principles for automated systems. The principles include safe and effective systems, algorithmic discrimination protections, data privacy, notice and explanation, and human alternatives, consideration, and fallback.⁵⁰ President Biden's executive order directs federal agencies to develop protocols for safe use of AI and requires developers to share their models with the National Institute of Standards and Technology (NIST) for approval. The executive order also provides

resources that help developers and users avoid algorithmic bias, and the US Department of Justice will advise landlords and recipients of federal funding how to avoid fair housing violations when using AI technology and other algorithmic models.⁵¹ Recently, the Federal Reserve also indicated that the ECOA must evolve with the realities of AI.⁵²

Agencies created to oversee the mortgage industry have also published guidance on AI. Before the release of the report mentioned in the previous section, the FHFA issued an advisory bulletin to the GSEs in February 2022 that offered risk management guidance regarding the use of AI and machine learning.⁵³ The bulletin outlined four areas of heightened risk: (1) model risks, including lack of transparency (or “black box risk”) and model performance degradation; (2) data risks, including relying on unrepresentative or unsuitable data sources; (3) other operational risks, particularly insufficient information technology infrastructure, lack of information security, reliance on third-party providers, and potential disruptions to business continuity; and (4) regulatory and compliance risks, including lack of compliance with consumer protection, fair lending, privacy, and employment discrimination laws and regulations.

Recommendations to Advance Racial Equity through AI

AI has the potential to advance racial equity, but it also has the potential to harm communities of color if not used properly. In other sectors, AI has been shown to bias outcomes against people and communities of color. This is a moment to make sure the benefits of AI are equitable in the housing finance system, which underpins homeownership opportunity.

Based on our conversations with mortgage industry stakeholders, we have developed several recommendations for policymakers, regulators, developers, and users to ensure that AI equitably expands access to mortgage services. We recommend three areas of focus: intentional design for equity, carefully studied pilot programs, and regulatory guidance.

Design with Intention

The key to intentional design is making equity a top priority rather than an afterthought. Experts in this area advocate for speculation on potential harms ahead of development or implementation, rather than playing catch-up after the product or program is already in motion.⁵⁴ And documenting inputs, processes, and results can bring transparency to the innovative process.

Many interviewees also highlighted the necessity of intentional design for creating equitable AI. This approach combines the AI tool in question with the power of human insight and thoughtfulness about potential unintended benefits or harms.

This type of design is key for creating AI that prioritizes fairness and learns to create more equitable outcomes. Intentional design requires thoughtfulness about the following:

- the training data being used and what biases need to be accounted for
- whether outcomes reflect structural inequities and what can be done to remedy them
- how to prevent model drift over time
- how to make the AI or machine learning tool more transparent to users
- what the end goal of the AI use is and whether the human process it is replacing is fair or needs improvement.

Given the role federal agencies play in mortgage lending, intentional design of AI pilots could start there. But there is also an opportunity for effective intentional design to occur among fintech companies creating algorithmic tools as well as early adopters of these tools.

Another way to advance intentional design is to increase diversity among the designers of these tools and their users across the mortgage industry. Black and Hispanic people are significantly underrepresented in the tech sector and in leadership positions across industries that rely on tech.

On these counts of intentional design, smaller lenders, CDFIs, and MDIs can serve as models—and their data used to teach the AI-based models—for what it looks like to prioritize equity. Many of the interviewees we spoke to from these types of organizations had diverse leadership staff with significant representation from the communities they served. They reported that this level of individual representation helped shape their mission and lending habits, as lived experience could help inform the decisionmaking and priorities of their leaders.

People with lived experiences of racism and inequality are well equipped to understand how certain AI initiatives may entrench or worsen systemic racism. Therefore, it is important to increase diversity across these sectors to mitigate potential harms and center equity.

Another method of intentional design is adversarial debiasing, a technique used to identify and reduce bias in credit decisionmaking tools. Zest AI pioneered the approach and has patented related applications of adversarial debiasing in its credit models. According to McGill, adversarial debiasing pits two machine learning models against each other during the model training process.⁵⁵ The first underwriting model predicts creditworthiness, but the second adversary one predicts the race or other potentially protected class attributes based solely on the risk scores of the first model.⁵⁶ Competition in this game forces the underwriting model to increase parity between scores from protected and unprotected classes so that the adversary cannot accurately predict race based on risk score, resulting in a final model that is accurate and fair.⁵⁷

Pilot Programs

The federal government—Ginnie Mae and the GSEs, in particular—can also use pilot programs to determine the effectiveness of AI tools and address equity concerns at a smaller scale before the industry implements these tools more broadly.

For example, Ginnie Mae's Office of Enterprise Risk has launched a series of machine learning and AI model pilots, using different approaches to explore new ways of measuring and analyzing data.³⁸ One AI algorithm was deployed to reduce the probability of false negatives and false positives when identifying issuers that may pose enhanced risk to the program but may slip through the cracks of traditional methods of risk identification.

A potential new pilot could be used by the GSEs or the FHA to test the use of AI in mortgage servicing. These agencies do not service loans themselves, but mortgage servicing is an integral part of ensuring that investors, including those that own agency mortgage-backed securities, receive their payments from the mortgage payments made by homeowners. As these agencies bear the credit risk of the loans, they set the rules for how servicers manage nonperforming loans.

An AI-based algorithm that projects the likelihood of borrower delinquency and identifies the best risk management process could be valuable. The GSEs or the FHA could deploy the pilot and test it against current processes to determine delinquency and resulting management. Measuring how the model performs against current systems would ultimately improve accuracy and guard against racial bias. This process could continue until the algorithm was deemed ready for use.

Pilots can be useful for incorporating AI throughout more stages of the mortgage process and identifying key frictions: Why might lenders be hesitant to adopt equity-advancing AI tools? What incentives can advance the use of these tools? How well do consumers understand AI use in the mortgage process, and what are their salient concerns? What are the unintended consequences of these tools, and how can they be mitigated?

Federal Home Loan Banks (FHLBs) play a crucial role in providing liquidity and funding to financial institutions, which affects the availability and affordability of mortgage loans. Given the potential of AI to address the racial homeownership gap, it is essential for FHLBs to closely monitor its development. AI innovations can introduce new lending models, streamline processes, and improve access to credit for underserved communities. By monitoring the AI landscape, FHLBs can identify emerging technologies and partnerships that could advance fair lending practices and expand homeownership opportunities for underrepresented groups. They can collaborate with AI start-ups, their members and other financial institutions, and regulatory bodies to ensure that innovative solutions promote inclusive lending and reduce disparities in homeownership.

Furthermore, monitoring AI in the context of closing the racial homeownership gap allows FHLBs to assess and address potential risks and challenges. FHLBs can engage with AI companies to promote transparency, ethical use of data, and bias-free algorithms. They can also work with regulatory agencies

to establish guidelines and standards that foster responsible AI practices, emphasizing fair lending and equal access to credit. By monitoring AI developments, FHLBs can leverage these innovations to create a more equitable housing finance system and contribute to closing the racial homeownership gap.⁵⁹

Increase Regulatory Guidance for Use of AI in the Financial Services Sector

Federal actors have two vital roles that will fundamentally set the course for AI in mortgage finance: (1) regulatory guidance and standard setting and (2) advancing adoption via the FHA, Fannie Mae, and Freddie Mac.

Lack of clear regulatory standards governing the use of AI across the mortgage industry was a major point of concern for nearly every stakeholder with whom we spoke. There is a lack of clear standards from the FHFA, the GSEs, and HUD about the federal government's expectations regarding the use of AI. One interviewee captured a sentiment many stakeholders shared: "Anything that creates more certainty and safety from the regulatory community would help" both industry and consumer stakeholders.⁶⁰

Additionally, the CFPB has a role to play in this space. The CFPB is currently engaged in rulemaking on Section 1033 of the Dodd-Frank Wall Street Reform and Consumer Protection Act. Section 1033 addresses consumer-permissioned data, specifically requiring financial institutions to make account data available to their consumers. It also directs the CFPB to set standards for the development and format of these data.⁶¹

It will be key for the CFPB to clearly delineate "which data elements consumers have a right to access, what the standards are for private companies accessing and transferring data, and how several federal consumer finance laws should be applied to consumer data transfers" (Choi et al. 2022, 39).

Regulation in the AI space is also important to clarify standards for lenders partnering with third-party vendors (e.g., fintech companies) offering AI tools. First, data privacy and security standards are essential for protecting sensitive consumer information as it is transferred from one party to another. Consumers can experience significant harms if their data are not handled adequately.

In addition, the government, lenders, and these vendors may all have differing incentives for AI use that can affect outcomes, so it is up to federal regulators to ensure that consumers, particularly the most vulnerable, are protected.

For instance, take the residential property technology sector, which primarily provides digital tools and platforms that landlords rely on to interface with their clients. A recent research report on property technology companies found that venture capital funding creates an incentive structure that encourages these firms to prioritize investor profits over the best interests of tenants and potential homebuyers (TechEquity Collaborative 2023). Fintech companies, even those whose mission is to advance equity, are necessarily profit motivated as well. Thus, it is up to regulators to ensure that profits do not come at the expense of equitable and affordable housing access.

Because AI is a topic that cuts across the economy, other corners of the federal government may be critical areas of AI policy setting. Through the AI Bill of Rights, the current administration has already taken an important step of raising the topic of AI and advocating for addressing the potential for racial bias. Congress has also voiced its interest in addressing AI.⁶²

In addition, the depository supervisors have indicated an interest in AI and implications for their industry. The supervisors include the US Department of the Treasury, the Federal Deposit Insurance Corporation, the Federal Reserve, the Office of Comptroller of the Currency, and the National Credit Union Administration.⁶³

And NIST has taken critical steps on AI (Axelrod et al. 2022). Specifically, NIST, a federal agency, developed its framework for voluntary use and to improve the ability to incorporate trustworthiness considerations into the design, development, use, and evaluation of AI products, services, and systems. NIST notes in its framework the importance of addressing fairness in AI, including concerns for equality and equity by addressing harmful bias and discrimination.⁶⁴

The federal government—HUD, the GSEs, and the FHFA, in particular—play critical roles in advancing equitable AI in the mortgage industry. The lending industry, particularly large lenders, often seek a competitive advantage through innovation. Simultaneously, the rules and processes that these agencies implement, which bring greater standardization and scale, also fuel much of the innovation in the industry. If these agencies expand use of AI or require evidence of equity when using their AI tools, the industry will be more likely to implement equitable AI. This will keep the federal government from being a chasm that limits widespread adoption of equitable AI.

The Role of Privacy

This report has focused on the efficiency and equity aspects of AI. But AI also has the potential to undermine privacy.⁶⁵ The use of AI will likely amplify these issues. AI is optimized with a lot of data as

inputs, including personal information. In addition, as AI expands its footprint, its use could cover more processes that use personal information. And AI can create false information about people. For example, its ability to create false images and audio could affect one's personal information.

As part of the policy response to protect people and drive equity while still benefiting from the efficiency of AI, several steps are needed. Recent analysis by NFHA lays out key suggestions that can help protect privacy, such as shifting responsibility from the individual to companies and regulators, strengthening the review of algorithmic tools before their use on the public, and supporting development of an intersectional approach to design and regulate tools and models (Holloway et al. 2023).

Adopting these suggestions could be a positive first step toward maintaining privacy in the AI era. But the entire ecosystem must remain vigilant as AI evolves and adoption inevitably quickens. Continued evaluation of best practices is also critical to ensure that policy keeps up with practice.

Outlook for AI

Amid higher origination costs and greater competition, the use of AI across the mortgage industry is poised to grow. This is in part because of a massive revolution in the amount of data and computing power. In addition, while large lenders are leading the adoption of AI, fintech companies and other third-party vendors are expanding AI use across smaller lenders and areas of the federal government.

But several forces may limit growth of AI in the mortgage sector. The mortgage industry is heavily regulated in part to ensure that historically marginalized groups are protected. This also means the federal government's presence in the mortgage industry heavily informs the investments a company makes in its lending processes.

Although anecdotes indicate that incorporating protected classes and alternative data into AI-based decisionmaking models up front can close racial inequities in mortgage outcomes, federal regulations put such practices into question and for good reason. Moreover, the up-front costs of AI adoption, and the dominance of Fannie Mae and Freddie Mac in setting standards, still discourage adoption by both large innovators and smaller mission-oriented companies. And some of these lenders—MDIs, in particular—may be highly skeptical of AI's power, given separate instances of racial bias by AI-based systems.

The challenges should not result in disbanding the promise of AI for both efficiency and equity. Instead, overcoming these challenges is the path forward. Strategies such as intentional design and testing of AI-based systems, piloting AI-based systems to understand their impact in a limited environment, having a mission orientation, and adding people of color to positions of influence within the organization can center racial equity in AI and better ensure that AI produces both equitable and efficient outcomes. Doing so will ensure that AI works for everyone in the industry.

Changes will not occur automatically; the federal government must lead the way. A proactive stance can spur greater adoption of AI, but it is critical to ensure that racially equitable considerations are codified and updated. Absent federal leadership, progress on both fronts will likely be sporadic, with the need for greater equity potentially lost in profit-maximizing decisions. A strong role for the federal government can overcome the innovation chasm, provide greater clarity on the price of innovation, and more easily expand the most promising AI-based services that optimize both efficiency and equity.

Notes

- ¹ Crystal Grant, "Algorithms Are Making Decisions about Health Care, Which May Only Worsen Medical Racism," ACLU, October 3, 2022, <https://www.aclu.org/news/privacy-technology/algorithms-in-health-care-may-worsen-medical-racism>; and Associated Press, "U.S. Warns of Discrimination in Using Artificial Intelligence to Screen Job Candidates," NPR, May 12, 2022, <https://www.npr.org/2022/05/12/1098601458/artificial-intelligence-job-discrimination-disabilities>.
- ² Nico Grant and Kashmir Hill, "Google's Photo App Still Can't Find Gorillas. And Neither Can Apple's," *New York Times*, May 22, 2023, <https://www.nytimes.com/2023/05/22/technology/ai-photo-labels-google-apple.html>. Several accounts of biased AI business practices have been reported in popular media. For example, Microsoft AI's chatbot "learned" to respond using racist language gathered by social media users; racially biased facial recognition programs, such as Google's facial recognition program, labeled photos of Black people as gorillas; and a Twitter algorithm was discovered to automatically edit out images of Black faces.
- ³ Kimberly Amadeo, "What Is the Production Possibilities Curve in Economics?" *The Balance*, June 16, 2022, <https://www.thebalancemoney.com/production-possibilities-curve-definition-explanation-examples-4169680>.
- ⁴ "What Is Strong AI?" IBM, accessed October 16, 2023, <https://www.ibm.com/topics/strong-ai>.
- ⁵ "What Is Strong AI?" IBM.
- ⁶ Dan Ennis, "Barr Acknowledges AI's Promise but Warns of Bias Risk," *Banking Dive*, July 18, 2023, <https://www.bankingdive.com/news/barr-fed-housing-lending-mortgage-artificial-intelligence-bias-cra/688270/>.
- ⁷ "Advisory Bulletin: Artificial Intelligence/Machine Learning Risk Management," *FHFA Insights* (blog), Federal Housing Finance Agency, February 10, 2022, <https://www.fhfa.gov/Media/Blog/Pages/Artificial-Intelligence-Machine-Learning-Supervisory-Guidance-for-Enterprises.aspx>.
- ⁸ Connie Kim, "What Opportunities Will the ICE-Black Knight Deal Present for Lenders?" *HousingWire*, April 4, 2023, <https://www.housingwire.com/articles/what-opportunities-will-the-ice-black-knight-deal-present-for-lenders/>.
- ⁹ Liam Reynolds, Vanessa G. Perry, and Jung Hyun Choi, "Closing the Homeownership Gap Will Require Rooting Systemic Racism Out of Mortgage Underwriting," *Urban Wire* (blog), Urban Institute, October 13, 2021, <https://www.urban.org/urban-wire/closing-homeownership-gap-will-require-rooting-systemic-racism-out-mortgage-underwriting>.
- ¹⁰ Amalie Zinn, Michael Neal, and Vanessa G. Perry, "Building Trust in the Financial System Is Key to Closing the Racial Wealth Gap," *Urban Wire* (blog), Urban Institute, June 15, 2023, <https://www.urban.org/urban-wire/building-trust-financial-system-key-closing-racial-wealth-gap>.
- ¹¹ Anna North, "How School Funding Can Help Repair the Legacy of Segregation," *Vox*, February 17, 2021, <https://www.vox.com/22266219/biden-education-school-funding-segregation-antiracist-policy>.
- ¹² Drew DeSilver and Kristen Bialik, "Blacks and Hispanics Face Extra Challenges in Getting Home Loans," *Pew Research Center*, January 10, 2017, <https://www.pewresearch.org/short-reads/2017/01/10/blacks-and-hispanics-face-extra-challenges-in-getting-home-loans>.
- ¹³ Jung Hyun Choi and Peter J. Mattingly, "What Different Denial Rates Can Tell Us about Racial Disparities in the Mortgage Market," *Urban Wire* (blog), Urban Institute, January 13, 2022, <https://www.urban.org/urban-wire/what-different-denial-rates-can-tell-us-about-racial-disparities-mortgage-market>.
- ¹⁴ Spencer Lee and Annabel Burba, "How Mortgage AI Chatbot Developers Aim to Ensure Regulatory Compliance," *National Mortgage News*, August 10, 2023, <https://www.nationalmortgagenews.com/news/how-mortgage-ai-chatbot-developers-aim-to-ensure-regulatory-compliance>.

- ¹⁵ Miranda Crace, "How Long Does Underwriting Take? A Timeline," Rocket Mortgage, August 1, 2023, <https://www.rocketmortgage.com/learn/how-long-from-underwriting-to-closing>.
- ¹⁶ For one example, see "The Mortgage Automation Software to Bring Efficiency into Mortgage Processing," Expert Mortgage Assistance, accessed October 16, 2023, <https://www.expertmortgageassistance.com/automation-analytics-solutions/msuite-mortgage-automated-tool.php>.
- ¹⁷ Daniel Pang, "Incorporating Two Alternative Types of Data into Mortgage Underwriting Could Make the Process More Equitable," *Urban Wire* (blog), Urban Institute, January 4, 2023, <https://www.urban.org/urban-wire/incorporating-two-alternative-types-data-mortgage-underwriting-could-make-process-more>.
- ¹⁸ Authors' interviews with Yolanda McGill.
- ¹⁹ Authors' interviews with Tom Oscherwitz.
- ²⁰ Authors' interviews with Tom Oscherwitz.
- ²¹ Jeanne Lee, "Understanding the Mortgage Underwriting Process," Bankrate, September 5, 2023, <https://www.bankrate.com/mortgages/steps-in-underwriting-process/>.
- ²² Michael Neal and Laurie Goodman, "Appraisal Waivers Have Helped Homeowners Find Payment Flexibility amid Pandemic-Induced Economic Struggles," *Urban Wire* (blog), Urban Institute, October 16, 2020, <https://www.urban.org/urban-wire/appraisal-waivers-have-helped-homeowners-find-payment-flexibility-amid-pandemic-induced-economic-struggles>.
- ²³ Authors' interview with Interviewee 3.
- ²⁴ Authors' interview with Charu Singh.
- ²⁵ Authors' interview with Charu Singh.
- ²⁶ "About Us," Resistant.AI, accessed October 16, 2023, <https://resistant.ai/about>.
- ²⁷ Team Capacity, "AI and Automation in Loan Servicing and Loss Mitigation," Capacity, January 12, 2022, <https://capacity.com/ai-and-automation-in-mortgage/ai-and-automation-in-loan-servicing-and-loss-mitigation/>.
- ²⁸ The ECOA allows lenders to use special purpose credit programs, which permit the consideration of a prohibited basis (e.g., race) for programs designed to benefit people who fall into "economically disadvantaged classes." See Equal Credit Opportunity (Regulation B); Special Purpose Credit Programs, 86 Fed. Reg. 3762 (Jan. 15, 2021).
- ²⁹ "What AI and Machine Learning Is... And What It Isn't," Black Knight, May 24, 2023, <https://www.blackknightinc.com/knowledge-base/what-ai-and-machine-learning-is-and-what-it-isnt/>.
- ³⁰ This example is purely illustrative. It is possible that AI was not the mechanism driving this AVM.
- ³¹ Authors' interview with Tom Oscherwitz.
- ³² Authors' interview with Yolanda McGill.
- ³³ Authors' interview.
- ³⁴ Authors' interview.
- ³⁵ Authors' interview.
- ³⁶ Authors' interview with Nicole Elam.
- ³⁷ Authors' interview with Nicole Elam.
- ³⁸ Authors' interview with Gary Perez.
- ³⁹ Authors' interview with Keisha Deonarine.

- ⁴⁰ Authors' interviews with Chi Chi Wu and Odette Williamson.
- ⁴¹ Authors' interviews with Chi Chi Wu and Odette Williamson.
- ⁴² "Our Programs: Tech Equity," National Fair Housing Alliance, accessed October 16, 2023, <https://nationalfairhousing.org/program/tech-equity-initiative/>.
- ⁴³ "Our Programs: Tech Equity," National Fair Housing Alliance.
- ⁴⁴ "Fairness of Machine Learning in Underwriting Models," National Fair Housing Alliance, May 12, 2022, <https://nationalfairhousing.org/fairness-of-machine-learning-in-credit-underwriting-models/>; and "NFHTA Forum: Mining the Data: Algorithmic Bias in Housing Related Transactions," National Fair Housing Alliance, February 20, 2022, <https://nationalfairhousing.org/nfhta-forum-mining-the-data-algorithmic-bias-in-housing-related-transactions/>.
- ⁴⁵ Authors' interviews with FinRegLab.
- ⁴⁶ Adam Hoffberg and Bill Reeder, "The FHA TOTAL Mortgage Algorithm: Providing Machine Learning Analysis for Two Decades," US Department of Housing and Urban Development, Office of Policy Development and Research, June 27, 2023, <https://www.huduser.gov/portal/pdrEDGE/pdr-edge-pdrat50-062723.html>.
- ⁴⁷ "DOJ, HUD Say Fair Housing Act Extends to Algorithm-Based Tenant Screening," *InfoBytes* (blog), Buckley, January 13, 2023, <https://buckleyfirm.com/blog/2023-01-13/doj-hud-say-fair-housing-act-extends-algorithm-based-tenant-screening>.
- ⁴⁸ Rohit Chopra, "Algorithms, Artificial Intelligence, and Fairness on Home Appraisals," Consumer Financial Protection Bureau blog, June 1, 2023, <https://www.consumerfinance.gov/about-us/blog/algorithms-artificial-intelligence-fairness-in-home-appraisals/>.
- ⁴⁹ "Consumer Financial Protection Circular 2023-03: Adverse Action Notification Requirements and the Proper Use of the CFPB's Sample Forms Provided in Regulation B," Consumer Financial Protection Bureau, accessed October 17, 2023, <https://www.consumerfinance.gov/compliance/circulars/circular-2023-03-adverse-action-notification-requirements-and-the-proper-use-of-the-cfpbs-sample-forms-provided-in-regulation-b/>.
- ⁵⁰ "Blueprint for an AI Bill of Rights," White House, Office of Science and Technology Policy, accessed October 17, 2023, <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>.
- ⁵¹ "Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," White House, October 30, 2023, <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>.
- ⁵² Chris Clow, "Fed Watchdog: Fair Housing Act, ECOA Must Evolve with Realities of AI," *HousingWire*, July 18, 2023, <https://www.housingwire.com/articles/fed-watchdog-fair-housing-act-ecoa-must-evolve-with-realities-of-ai/>.
- ⁵³ "Artificial Intelligence/Machine Learning Risk Management," Federal Housing Finance Agency, February 10, 2022, <https://www.fhfa.gov/SupervisionRegulation/AdvisoryBulletins/Pages/Artificial-Intelligence-Machine-Learning-Risk-Management.aspx>.
- ⁵⁴ Casey Fiesler, "AI Has Social Consequences, but Who Pays the Price? Tech Companies' Problem with 'Ethical Debt,'" *The Conversation*, April 18, 2023, <https://theconversation.com/ai-has-social-consequences-but-who-pays-the-price-tech-companies-problem-with-ethical-debt-203375>.
- ⁵⁵ Authors' interview with Yolanda McGill.
- ⁵⁶ Authors' interview with Yolanda McGill.
- ⁵⁷ Zest AI Team, "There's a Fix to the Problem of Biased Algorithms in Lending," *Zest AI*, December 16, 2019, <https://www.zest.ai/insights/theres-a-fix-to-the-problem-of-biased-algorithms-in-lending>.

- ⁵⁸ "Artificial Intelligence at Ginnie Mae," Ginnie Mae, May 17, 2021, <https://www.ginniemae.gov/newsroom/HAPS/Pages/Post.aspx?PostID=55>.
- ⁵⁹ "Artificial Intelligence/Machine Learning Supervisory Guidance for Enterprises," *FHFA Insights* (blog), Federal Housing Finance Agency, February 10, 2022, <https://www.fhfa.gov/Media/Blog/Pages/Artificial-Intelligence-Machine-Learning-Supervisory-Guidance-for-Enterprises.aspx>.
- ⁶⁰ Authors' interview.
- ⁶¹ "Required Rulemaking on Personal Financial Data Rights," Consumer Financial Protection Bureau, accessed October 17, 2023, <https://www.consumerfinance.gov/personal-financial-data-rights/>; and Consumer Financial Protection Bureau, "CFPB Issues Guidance on Credit Denials by Lenders Using Artificial Intelligence," news release, September 19, 2023, <https://www.consumerfinance.gov/about-us/newsroom/cfpb-issues-guidance-on-credit-denials-by-lenders-using-artificial-intelligence/>.
- ⁶² "Oversight of A.I.: Rules for Artificial Intelligence," US Senate Committee on the Judiciary, Subcommittee on Privacy, Technology, and the Law, May 16, 2023, <https://www.judiciary.senate.gov/committee-activity/hearings/oversight-of-ai-rules-for-artificial-intelligence>.
- ⁶³ Request for Information and Comment on Financial Institutions' Use of Artificial Intelligence, including Machine Learning, 86 Fed. Reg. 16837 (Mar. 31, 2021).
- ⁶⁴ "AI Risk Management Framework," National Institute of Standards and Technology, January 26, 2023, <https://www.nist.gov/itl/ai-risk-management-framework>.
- ⁶⁵ Tom Wheeler, "The Three Challenges of AI Regulation," Brookings Institution, June 15, 2023, <https://www.brookings.edu/articles/the-three-challenges-of-ai-regulation/>.

References

- Akinwumi, Michael, John Merrill, Lisa Rice, Kareem Saleh, and Maureen Yap. 2021. "An AI Fair Lending Policy Agenda for the Federal Financial Regulators." Washington, DC: Brookings Institution.
- Akinwumi, Michael, Lisa Rice, and Snigdha Sharma. 2022. "Purpose, Process, and Monitoring: A New Framework for Auditing Algorithmic Bias in Housing and Lending." Washington, DC: National Fair Housing Alliance.
- Avenancio-León, Carlos, and Troup Howard. 2022. "The Assessment Gap: Racial Inequalities in Property Taxation." *Quarterly Journal of Economics* 137 (3): 1383–434. <https://doi.org/10.1093/qje/qjac009>.
- Axelrod, Judah, Alena Stern, Michael Neal, and Linna Zhu. 2022. "Comment Letter on NIST Initial Draft of Artificial Intelligence Risk Management Framework." Washington, DC: Urban Institute.
- Azhar Ali, Syed Emad, Syed Sajjad Hussain Rizvi, Fong Won Lai, Rao Faizan Ali, and Ahmed Ali Jan. 2021. "Predicting Delinquency on Mortgage Loans: An Exhaustive Parametric Comparison of Machine Learning Techniques." *International Journal of Industrial Engineering and Management* 12 (1).
- Bailey, Natalia, Alex Bloomfield, Kelly Thompson Cochran, Sarah Davies, Colin Fooks, Saurab Gautam, Hilary Griggs, et al. 2023. *Explainability and Fairness in Machine Learning for Credit Underwriting: Policy and Empirical Findings Overview*. Washington, DC: FinRegLab.
- Bartlett, Robert, Adair Morse, Richard Stanton, and Nancy Wallace. 2020. *Consumer-Lending Discrimination in the FinTech Era*. Working Paper 25943. Cambridge, MA: National Bureau of Economic Research.
- . 2022. "Consumer-Lending Discrimination in the FinTech Era." *Journal of Financial Economics* 143 (1): 30–56. <https://doi.org/10.1016/j.jfineco.2021.05.047>.
- Bayer, Patrick, Fernando Ferreira, and Stephen L. Ross. 2018. "What Drives Racial and Ethnic Differences in High-Cost Mortgages? The Role of High-Risk Lenders." *Review of Financial Studies* 31 (1): 175–205. <https://doi.org/10.1093/rfs/hhx035>.
- Black Knight. n.d. "Management of AI and Machine Learning in the Residential Mortgage Industry." Jacksonville, FL: Black Knight Financial Technology Solutions.
- Brahma, Arin, David M. Goldberg, Nohel Zaman, and Mariano Aloiso. 2021. "Automated Mortgage Origination Delay Detection from Textual Conversations." *Decision Support Systems* 140.
- Choi, Jung Hyun, Michael Stegman, Laurie Goodman, Janneke Ratcliffe, Karan Kaul, Rita Ballesteros, Caitlyn Young, Daniel Pan, and Liam Reynolds. 2022. *Reducing the Black Homeownership Gap through Underwriting Innovations*. Washington, DC: Urban Institute.
- Cochran, Kelly Thompson, Colin Fooks, Hilary Griggs, Tess Johnson, Mashrur Khan, Duncan McElfresh, Kerrigan Molland, P-R Stark, YaYa Sun, and Sormeh Yazdi. 2021. *The Use of Machine Learning for Credit Underwriting: Market and Data Science Context*. Washington, DC: FinRegLab.
- Couch, Kenneth A., and Robert Fairlie. 2010. "Last Hired, First Fired? Black-White Unemployment and the Business Cycle." *Demography* 47 (1): 227–47. <https://doi.org/10.1353/dem.0.0086>.
- Federal Reserve. n.d. "Federal Fair Lending Regulations and Statutes." In *Consumer Compliance Handbook*, 577–79. Washington, DC: Board of Governors of the Federal Reserve System, Division of Consumer and Community Affairs.
- Freddie Mac Single-Family. 2021. *Cost to Originate Study: How Digital Offerings Impact Loan Production Costs*. McLean, VA: Freddie Mac Single-Family.
- Gerken, Matthew, Samantha Batko, Katie Fallon, Emma Fernandez, Abigail Williams, and Brendan Chen. 2023. "Addressing the Legacies of Historical Redlining." Washington, DC: Urban Institute.
- Gillis, Talia B. 2021. "The Input Fallacy." New York: SSRN.

- Goodman, Laurie, Janneke Ratcliffe, Michael Neal, Jung Hyun Choi, Linna Zhu, John Walsh, Daniel Pang, et al. 2023. *Housing Finance: At a Glance Monthly Chartbook, September 2023*. Washington, DC: Urban Institute.
- Goodman, Laurie, and Jun Zhu. 2021. *The Future of Headship and Homeownership*. Washington, DC: Urban Institute.
- Holloway, Hannah, Snigdha Sharma, Samantha Gordon, and Michael Akinwumi. 2023. *Privacy, Technology, and Fair Housing: A Case for Corporate and Regulatory Action*. Washington, DC: National Fair Housing Alliance.
- McCargo, Alanna. 2017. "Sustainable Housing Finance: Private-Sector Perspectives on Housing Reform, Part III." Statement before the US House of Representatives, Subcommittee on Housing and Insurance, Washington, DC, November 7.
- Narragon, Melissa, Danny Wiley, Doug McManus, Vivian Li, Kangli Li, Xue Wu, and Kadiri Karamon. 2021. "Racial and Ethnic Valuation Gaps in Home Purchase Appraisals." McLean, VA: Freddie Mac.
- Neal, Michael, Sarah Stochak, Linna Zhu, and Caitlin Young. 2020. "How Automated Valuation Models Can Disproportionately Affect Majority-Black Neighborhoods." Washington, DC: Urban Institute.
- OIG (Office of Inspector General). 2022. "Enterprise Use of Artificial Intelligence and Machine Learning." Washington, DC: Federal Housing Finance Agency, OIG.
- Okun, Arthur M. 2015. *Equality and Efficiency: The Big Tradeoff*. Washington, DC: Brookings Institution Press.
- PAVE (Interagency Task Force on Property Appraisal and Valuation Equity). 2022. *Action Plan to Advance Property Appraisal and Valuation Equity*. Washington, DC: PAVE.
- Rugh, Jacob S., and Douglas S. Massey. 2010. "Racial Segregation and the American Foreclosure Crisis." *American Sociological Review* 75 (5): 629-51.
- Santucci, Larry. 2019. *How Prevalent Were Racially Restrictive Covenants in 20th Century Philadelphia? A New Spatial Data Set Provides Answers*. Philadelphia: Federal Reserve Bank of Philadelphia Consumer Finance Institute.
- Steurer, Miriam, Robert J. Hill, and Norbert Pfeifer. 2021. "Metrics for Evaluating the Performance of Machine Learning Based Automated Valuation Models." *Journal of Property Research* 38 (2): 99-129. <https://doi.org/10.1080/09599916.2020.1858937>.
- TechEquity Collaborative. 2023. *The Promise and Perils of Residential Proptech: Year 1 Research Summary Report*. Oakland, CA: TechEquity Collaborative.
- van Zetten, W., G. J. Ramackers, and H. H. Hoos. 2022. "Increasing Trust and Fairness in Machine Learning Applications within the Mortgage Industry." *Machine Learning with Applications* 10.
- Venkateswaran, Nitya, Jay Feldman, Stephanie Hawkins, Megan A. Lewis, Janelle Armstrong-Brown, Megan Comfort, Ashley Lowe, and Daniela Pineda. 2023. *Bringing an Equity-Centered Framework to Research: Transforming the Researchers, Research Content, and Practice of Research*. Raleigh, NC: RTI Press.
- Wu, Chi Chi. 2021. "Doing Special Purpose Credit Programs Right: Why Programs to Assist Black Communities Should Avoid Conventional Use of Traditional Credit Scores." Boston: National Consumer Law Center.
- Wu, Chi Chi. 2022. "Even the Catch-22s Come with Catch-22s: Potential Harms and Drawbacks of Rent Reporting." Boston: National Consumer Law Center.
- Zhang, Xueru, Mohammad Mahdi Khalili, and Mingyan Liu. 2019. "Long-Term Impacts of Fair Machine Learning." *Ergonomics in Design* 28 (3): 7-11. <https://doi.org/10.1177/1064804619884160>.
- Zhu, Linna, Michael Neal, and Caitlin Young. 2022. "Revisiting Automated Valuation Model Disparities in Majority-Black Neighborhoods: New Evidence Using Property Condition and Artificial Intelligence." Washington, DC: Urban Institute.
- Zou, Leying, Warut Khern-am-naui. 2022. "AI and Housing Discrimination: The Case of Mortgage Applications." *AI and Ethics* 3:1271-81. <https://doi.org/10.1007/s43681-022-00234-9>.

About the Authors

Michael Neal is a principal research associate and Equity Scholar in the Housing Finance Policy Center at the Urban Institute. Previously, he worked at Fannie Mae, where he was a director of economics in the Economic and Strategic Research Group. Before his service at Fannie Mae, Neal was the assistant vice president at the National Association of Home Builders' Economic and Housing Policy department. As a housing economist, Neal has an in-depth knowledge of housing market trends and has provided expert analysis and commentary on housing to media outlets around the country. Previously, he worked at Congress's Joint Economic Committee, the Federal Reserve System, the Congressional Budget Office, and Goldman Sachs. Neal has a bachelor's degree in economics from Morehouse College and a master's degree in public administration from the University of Pennsylvania.

Linna Zhu is a senior research associate in the Housing Finance Policy Center. She conducts data-driven quantitative research on policy issues related to US mortgage finance, housing supply, and racial inequality. Zhu has published research articles on such topics as the Community Reinvestment Act, automated valuation models, artificial intelligence risk management, and AAPI mortgage denial rates. Her research has been published in leading academic and policy journals, such as *Regional Science and Economics* and *Cityscape*. She serves on the advisory board of the Design Thinking Executive Program at the University of California, Riverside Extension. Zhu holds a BA in economics from Renmin University of China, an MS in finance from Johns Hopkins University, and a PhD in public policy and management from the University of Southern California.

Caitlin Young is a former policy analyst in the Housing Finance Policy Center. She graduated from the University of North Carolina at Chapel Hill with double majors in political science, for which she received highest honors, and in economics. She authored her senior thesis on the impact of state social welfare policy on recidivism outcomes. Before joining Urban, she interned with the Center for Community Self-Help, a community development financial institution. She is a first-year law student at Yale Law School.

Vanessa Gail Perry is a nonresident fellow in the Housing Finance Policy Center and a professor and interim vice dean at the George Washington University School of Business. Her research is focused on inequality, discrimination, and public policy interventions in credit and housing markets. Recent work has examined racial and ethnic disparities in homeownership attributable to mortgage loan underwriting guidelines, appraisal policies, and credit scoring. Perry previously served as a senior

adviser to the secretary of the US Department of Housing and Urban Development and as an expert at the Consumer Financial Protection Bureau. She also served for several years as a senior economist at Freddie Mac. In her capacity as president of the Rice Coleman Ross Group, she has been a consultant to numerous public- and private-sector clients. Perry has a BA from the American University, an MBA from Washington University in St. Louis, and a PhD from the University of North Carolina at Chapel Hill.

Matthew Pruitt is a research assistant in the Housing Finance Policy Center. Before joining Urban, he interned at Brookings Metro, the Boston Municipal Research Bureau, the National Consumer Law Center, and the Lawyers' Committee for Civil Rights Under Law. Pruitt graduated magna cum laude from Boston University with a degree in political science, where his coursework focused on urban policy and local governance.

STATEMENT OF INDEPENDENCE

The Urban Institute strives to meet the highest standards of integrity and quality in its research and analyses and in the evidence-based policy recommendations offered by its researchers and experts. We believe that operating consistent with the values of independence, rigor, and transparency is essential to maintaining those standards. As an organization, the Urban Institute does not take positions on issues, but it does empower and support its experts in sharing their own evidence-based views and policy recommendations that have been shaped by scholarship. Funders do not determine our research findings or the insights and recommendations of our experts. Urban scholars and experts are expected to be objective and follow the evidence wherever it may lead.



500 L'Enfant Plaza SW
Washington, DC 20024

www.urban.org



January 31, 2024

Chair Tina Smith
 Subcommittee on Housing, Transportation and Community Development
 Senate Banking, Housing, and Urban Affairs
 534 Dirksen Senate Office Building
 Washington, D.C. 20510

Ranking Member Cynthia Lummis
 Subcommittee on Housing, Transportation and Community Development
 Senate Banking, Housing, and Urban Affairs
 534 Dirksen Senate Office Building
 Washington, D.C. 20510

Dear Chair Smith and Ranking Member Lummis:

The National Multifamily Housing Council (NMHC) and National Apartment Association (NAA) applaud the Committee for calling a hearing entitled “Artificial Intelligence and Housing: Exploring Promise and Peril.” We appreciate the Committee exploring the use of emerging technologies, like Artificial Intelligence (AI), in the housing space and the many ways they are providing significant benefits to property owners and their residents. NMHC and NAA members have been at the forefront of technology deployments that have reshaped business operations, improved housing affordability, and benefited millions of Americans renters.

For more than 25 years, NMHC and NAA have partnered on behalf of America’s apartment industry. Our combined memberships are engaged in all aspects of the apartment industry, including ownership, development, management, finance, and suppliers partners/service providers. Drawing on the knowledge and policy expertise of staff in Washington, D.C., as well as the advocacy power of 141 NAA state and local affiliated associations, NMHC and NAA provide a single voice for developers, owners, and operators of multifamily rental housing. One-third of all Americans rent their housing, and 38.9 million of them live in an apartment home.

The Positive Role of AI and Emerging Technologies in Rental Housing

The use of artificial intelligence—or AI—has garnered significant attention in recent years, with policymakers and media outlets addressing it on an almost daily basis. Whether they are touting its beneficial qualities or warning of its potential pitfalls, it has been at the epicenter of political and social discussions across the nation and has significant implications for housing providers.

For the better part of the past two decades, the rental housing sector saw tremendous growth and opportunities emerge through new, innovative technologies. In many ways, multifamily housing communities have served as the innovation lab for America. Today, seamless connectivity delivered by Managed Wi Fi or bulk internet systems have enabled AI and related technologies in the multifamily housing sector and led to significant gains in meeting resident expectation and demand through virtual leasing and touring, enhanced resident screening, home automation, or improved property level climate resilience and predictive maintenance to name just a few. Yet with that, like most other technologies, the risk landscape has also

expanded for housing operators because of these technologies and why the multifamily industry continues to place a high priority on data security, integrity, and consumer privacy.

A snapshot of some of the early ways that AI and other technologies are being utilized in the multifamily housing context are below:

Housing Developers are using technology in critical ways to build and rehabilitate rental properties. A growing number of AI enabled platforms are cutting costs by automating design and planning and some are now even using AI to identify opportunities to develop in areas of need with favorable zoning and land use policies. Both are critical issues as we continue to work towards improved housing affordability, equity, and opportunity.

Housing Operators are using AI enabled platforms to overcome labor market challenges, provide positive rental payment reporting services, improve the accuracy and reduce bias in resident screening, and even identify cases of fraud in the application and leasing process to protect the financial viability of the property. In fact, a recent [NMHC and NAA survey and report](#) found staggering increases of application fraud, contributing to both the growth in rents and number of evictions. A vast majority of respondents (70.7%) reported experiencing an increase in fraudulent applications and payments, utilizing fraudulent documentation, financial statements and even identities in the past twelve months. AI enabled screening processes and machine learning are aiding in efforts to protect the financial health of properties and their long-term viability. Importantly, these tools also help reduce foreseeable risks to community residents and employees.

Housing Technology Suppliers are the backbone of modern multifamily property operations. From the leasing office to managing a network of renter centric PropTech platforms, solutions are being deployed to solve some of our most stubborn housing challenges like increasing responsiveness to applicant inquiries and follow-up with residents, improving sustainability and resilience, providing enhanced security for residents and communities, and driving operational costs down through predictive maintenance.

Housing Owners and Financiers are leveraging cost saving technology platforms to improve efficiencies, underwrite lending, identify investing opportunities to meet growing rental demand and work towards improving housing affordability. Through technology, owners can more quickly reallocate funding to drive production and preservation of critically needed rental housing. Even further, owners can deliver tech-driven amenities, service and solutions that residents want and increasingly demand.

Regular Review of AI and other Technologies is Needed

While the multifamily industry, like others, is in the early stages of adoption of tools and platforms that utilize AI, machine learning and other emerging technologies, there is a clear upside to their use. Yet, NMHC and NAA members also acknowledge the need for continual, internal review and oversight of AI models, processes, and outcomes to ensure that both business operations and resident impact is understood, and negative outcomes are prevented. As technology becomes more integral to every facet of the real estate industry, stakeholders must work to uphold our strong commitment to Fair Housing, Fair Credit

Reporting and other pillars of the housing and consumer policy landscape. While policymakers debate how is most appropriate to regulate AI and other technologies, it is critical that they not stifle innovation or inhibit the growth of technologies that have significant pro-consumer and pro-housing benefits. Instead, legislative and regulatory efforts should continue to promote principles that prioritize data integrity, mitigate bias, and acknowledge the legitimate business purposes for technology. To ignore those principles in favor of hyper-prescriptive approaches would impede progress, increase costs for housing, and worse, it would freeze inequities, effectively making them permanent, and prevent advancements.¹

Existing Federal Regulations Provide Strong Consumer Protection

Because the relative availability of AI systems has changed recently, the public impression is that the issues are also new, giving the impression that new rules and regulations are necessary to deal with them. This is not the case.

The current legal landscape provides strong protection against risks posed by AI, machine learning and algorithmic decision making. Additionally, research into the relative bias of AI and humans has been the subject of research for decades and has shown that even AI systems presenting bias can be corrected. Policymakers should work to ensure that additional requirements are grounded in fact rather than fear. Instead, they should consider approaches that allows for safe harbors so that the industry maintains the flexibility it needs to continue to innovate in this evolving area. The safe harbor approach is common in housing, particularly affordable housing.

Conclusion

We thank you for the opportunity to present the views of the multifamily industry as you explore the many opportunities and the risks posed by AI in the housing market. We stand ready to support the efforts of Congress in its work to unleash the power of technology to improve housing affordability, enhance property operations and meet the expectations and needs of our nation's renters.

Sincerely,



Sharon Wilson Géro
President
National Multifamily Housing Council



Robert Pinnegar
President & CEO
National Apartment Association

¹ NMHC White Paper, *Artificial Intelligence & Resident Screening: Considerations for the Rental Housing Industry*, Amy M. Glassman, et al. Ballard Spahr LLP. October 2023.

**To Err Is Automated: Have Technological Advances in the Mortgage Market Increased
Opportunities for Black Homeownership?¹**

Working Paper

**Michael Neal, Urban Institute, Housing Finance Policy Center
Vanessa G. Perry, The George Washington University
Linna Zhu, Urban Institute, Housing Finance Policy Center**

¹ The opinions expressed in this paper are those of the authors and do not necessarily represent those of the Urban Institute or the Housing Finance Policy Center.

I. Introduction

In recent years, the polarity in the political landscape affecting Black Americans has been immense. Simultaneously, many highly publicized deaths of unarmed Black citizens, e.g., George Floyd and Breonna Taylor, have elevated the Black experience in the national conversation. Many organizations and institutions have responded to these events with diversity and anti-racism initiatives, and/or increased financial commitments to communities of color. Despite this momentum, the Black-white homeownership rate gap remains at 30 percentage points, and Black households that achieve homeownership don't experience equal benefits. This reduces Black wealth and financial security today and for future generations. Despite the widespread emphasis on homeownership as the 'American Dream', Black and other minority families face many barriers to homeownership, most of which can be traced to a history of racism and discrimination across social and economic institutions.

By reducing the role of human decision-making in home-buying, automation and digital transformation in the mortgage market could advance societal goals of eradicating racism and discrimination, improving accuracy, and removing bias. To date, research suggests that automation has little impact on homeownership outcomes for Black households (Perry and Martin, 2022). One solution currently being explored is to replace human property appraisal processes with more objective, data-driven algorithms, i.e., Automated Valuation Models (AVMs).

Home valuation methods, i.e., appraisals, are critical to determining the market value of a home, the equity a mortgage borrower holds in a property, and the amount of default risk to which a lender or investor is exposed. As such, appraisals are a necessary step in the mortgage lending process. Ideally, these estimates provide an unbiased measure of asset value that investors can use to assess and manage risk. Errors in these estimates are costly for investors as well as for households and communities. AVMs are technological tools used to assess home values. Proponents of AVMs argue that they not only reduce biases and errors committed by human appraisers but that the underlying algorithms can be programmed to provide fairer and more accurate valuations. Critics of AVMs caution that these algorithms simply reproduce the same biases that are embedded in human-driven decision systems.

In this paper, we examine the role of past discrimination in the form of redlining on automated home valuations in predominately Black neighborhoods. Homes located in minority neighborhoods have historically and consistently had lower values and rates of house price appreciation than homes owned by similarly-situated White counterparts, explaining in part the persistent wealth gap between Black and White households. Recent anecdotal and empirical evidence has sparked a heated public policy debate about the role of discrimination and systematic racial bias in home appraisal processes.

Thus, we examine the extent to which AVMs convey historical discrimination, i.e., redlining. We ask whether current home values are a reflection of historical redlining and whether redlining has an additional effect on the prices paid by homebuyers. If AVMs capture the effects of past redlining, Black neighborhoods and residents are being unfairly penalized for historical circumstances beyond their control. In the case of overvaluation, prospective homebuyers risk loan denials; they also risk 'faux-equity', which means housing assets are worth less than expected. Undervaluation, on the other hand, exacerbates neighborhood racial disparities in home values and suppresses equity accumulation. Our findings reveal that past redlining affects current home values and prices. Based on these findings, we present an actionable agenda for policymakers seeking to boost Black homeownership.

II. Background

2.1 Redlining

The Home Owners Loan Corporation (HOLC) was established during the New Deal to stabilize the mortgage market by buying and refinancing mortgage loans that were at risk of default. There is widespread agreement among scholars that the HOLC maps institutionalized the practice known as "redlining" which became standard practice in government as well as private mortgage lending (Massey and Denton, 1993). Based on information collected from local officials and real estate professionals, HOLC staff developed a set of local area maps that graded neighborhoods based on risk-related factors, i.e., age and condition of housing, access to transportation, and other amenities, while also including the socioeconomic status and racial and ethnic composition of residents. Neighborhoods were assigned colors on these maps, with green indicating 'best', blue for 'Still Desirable,' yellow for 'Definitely Declining,' and red for

'Hazardous.' These maps were then used by lenders to designate areas that were considered low-risk versus higher-risk (Mitchell and Franco, 2018).

Numerous recent studies have found evidence of longstanding, negative effects of HOLC redlining on households and communities. A study by the NCRC found lower incomes, more minorities, and signs of gentrification in neighborhoods marked by HOLC as "hazardous" (Mitchell and Franco, 2018). Aaronson and colleagues (2017) found that areas that were on the lower graded side of HOLC boundaries in the 1930s experienced a marked increase in racial segregation in subsequent decades, as well as evidence of a long-run decline in home ownership, house values, and credit scores along the lower graded side of HOLC borders.

Appel (2016) found that the HOLC mapping policies are associated with a 5% decrease in 1990 house prices not explained by the area's physical features. Appel also found a 5-7% increase in vacant houses and a 3% decrease in owner-occupation in redlined areas. Townsley, et al. (2021) found that past redlining explains approximately 25% of the variation in the rate of poverty and economic inequality. Additionally, segregation and redlining predict 53 percent of the difference in median household income between neighborhoods. Neighborhoods graded 'A' have over \$130,000 in average household income, compared to \$50,000 in 'B or C' neighborhoods, and below \$40,000 in 'D' neighborhoods.'

Research has also shown that past HOLC grades affect non-economic outcomes. For example, Krieger and colleagues have uncovered links between redlining and preterm births in New York City (Krieger et al., 2020) and late-stage cancer diagnoses in Massachusetts (Krieger, et al., 2020). Huan and Sehgal (2022) found that in Baltimore, a hazardous HOLC categorization is associated with a 4-year reduction in life expectancy, and a 'definitely declining' rating is associated with a 5-year lower life expectancy, controlling for demographic characteristics. Hoffman, et al. (2020) found that redlined neighborhoods experience elevated land surface temperatures relative to their non-redlined neighbors by as much as 7 °C due to fewer shade trees and more asphalt and concrete coverage.

The present study builds on this prior work by examining the effects of the HOLC redlining practices on present-day home values and home prices.

2.2 Appraisal bias and AVMS

Proponents of digitalization in the appraisal process argue that digital technology has increased the accuracy of property valuations, resulting in lower costs and risks throughout the home buying process. For example, with digitalized appraisal inspections, appraisers can collect information about a property without in-person inspections, to be transmitted to AI automated valuation models (AVMs) that replace traditional, more subjective procedures.

While these innovations have gained widespread acceptance, appraisal bias has emerged as one of the most controversial issues in the mortgage industry. Several studies have documented systematic biases in traditional appraisals that result in lower values for Black and Hispanic homebuyers and neighborhoods (Rothwell and Perry, 2019). A recent Freddie Mac study concluded that Black and Latino mortgage applicants as well as predominately Black and Latino neighborhoods receive appraisal values that are lower than the contract price more often than White applicants, and these applied to human appraisers as well as for values produced by AVMs (Freddie Mac, 2021; 2022). Another widely cited study revealed that homes owned by Black and Hispanic individuals are more likely to be appraised at a lower value than the sales price (Folk and Chen, 2021). Research comparing traditional appraisals with those conducted by AVMs found that homes owned by White borrowers are more likely to have an appraised value that is at least 10 % higher than the AVM's estimated value compared to homes owned by Black borrowers; these overvaluations are also more likely to occur when White borrowers live in majority-Black neighborhoods (Williamson and Palim, 2022). On the contrary, other evidence suggests that AVM models are less likely to produce biased results, and as such, can be used to advance more equitable outcomes in appraisals for minority homebuyers and homeowners (House Canary, 2021).

The present study extends prior research by Neal, Zhu, and Young (2021), which examined data from Atlanta, Memphis, and Washington, DC, to determine how AVMs contribute to racial disparities in home value estimates. In particular, these authors studied the accuracy of AVMs in majority-Black neighborhoods relative to majority-white neighborhoods by comparing the AVM estimate with the actual sales price (i.e., AVM error). They found that even though the magnitude of the error (the absolute dollar difference) was larger in majority-white neighborhoods, the percentage magnitude of the error (the proportional difference for what that error meant) was larger in majority-Black neighborhoods after adjusting for lower sales prices in majority-Black neighborhoods. In addition, these authors found that Black

neighborhoods were more likely to be undervalued, i.e., have sales prices that fell below the AVM values. While they tested for several possible explanations for these disparities, they still found evidence of systematic inaccuracies in AVM estimates (Neal, Zhu, & Young, 2021).

III. Conceptual Framework and Hypotheses

Our model examines how sales prices, negotiated between prospective homebuyers and sellers, differ from the home value estimates produced by an AVM. AVMs are increasingly used in place of home valuations, i.e., appraisals, performed by human appraisers. Home sales prices are a function of objective factors as well as certain qualitative market dynamics. For example, in a hot housing market where demand for units exceeds supply, prospective buyers may have to offer higher prices to outbid competing offers. This is especially the case in areas where investors are seeking single-family properties. Alternatively, when prices are higher than the market will bear, properties may remain on the market longer, motivating sellers to reduce asking prices. Anecdotal evidence suggests that real estate agents versus discount brokers, i-buying companies, or For-Sale-By-Owner may result in pricing differentials due to differences in incentives and the negotiation process. In this study, we are not able to observe the particular mechanisms that drive home price differentials. Rather, we examine the incidence and magnitude of disparities in house prices by the racial composition of the neighborhood. However, it is important to understand these processes in order and the implications of AVM errors on Black neighborhoods and the housing market overall.

[Insert Figure 1 here]

Figure 1 depicts our conceptual framework, which shows the direct and indirect effects of past redlining on the housing market today. We propose that automated home valuations are a function of property conditions, neighborhood conditions, and past redlining. Prior evidence suggests that redlined neighborhoods experienced a lack of disinvestment and limited access to credit, which exacerbated the effects of racism on education and labor market opportunities over time. As a result, these areas are more likely to experience distressed property and neighborhood conditions. Coupled with the market value of racial segregation, lower home values result, and in turn, we expect these home values to predict house prices. At the same time, house prices are

also directly impacted by past redlining practices because comparable properties, used as inputs in AVM algorithms, are also influenced by past redlining.

The following section describes our data, analysis, and findings.

IV. Methodology and Data

4.1 Methodology

To examine the extent to which AVMs convey historical discrimination, we conduct two-stage least squares regressions. Our first-stage regressions investigate whether home values are a reflection of historical redlining, as shown in Eq. (1).

$$\ln HP_{i,2018} = \alpha + \beta \text{Redlined}_{n,1940} + \gamma \text{Black}_{n,2018} + \eta X_{in,2018} + \mu_c + \varepsilon_{i,2018} \quad (1)$$

$HP_{i,2018}$ is the sales price for a single-family property i in year 2018. $\text{Redlined}_{n,1940}$ is a dummy variable with 1 indicating the neighborhood in which property i is currently located was graded as Declining or Hazardous in the 1940s HOLC map. $\text{Black}_{n,2018}$ measures neighborhood-level majority race, with 1 indicating a majority-Black neighborhood, i.e., more than 50 percent of households are Black. $X_{in,2018}$ is a vector of controls including lot sizes, numbers of bedrooms and bathrooms, built year, exterior property condition, whether located within the historic HOLC map area, and neighborhood-level median household income. We also control for the county (μ_c) fixed effect.

β and γ are the two coefficients of interest that estimate the impact of historic discrimination on current home values. AVMs use a comparable sales approach to estimate the value of a particular property. If past redlining practices directly impact house prices, as tested in Eq.(1), then they will indirectly influence AVMs as house prices are used as inputs in AVM

algorithms. Each instance of a lower purchase price becomes a comparable sale for the next algorithmic valuation, carrying the impact of the historic redlining, which in the long run will significantly undervalue the property values in majority-Black neighborhoods.

Our conceptual framework proposes that past redlining could also have a direct impact on AVMs, as shown in Figure 1. We propose that AVMs are a function of property conditions, neighborhood conditions, and past redlining. We test this hypothesis by conducting the second-stage regressions following Eq. (2):

$$Pct_Diff_{i,2018} = \alpha + \beta Redlined_{n,1940} + \gamma Black_{n,2018} + \delta \widehat{HP}_{i,2018} + \theta ECR_{i,2018} + \rho N_{n,2018} + \varepsilon_{i,2018} \quad (2)$$

$Pct_Diff_{i,2018}$ is the percentage magnitude of AVM inaccuracy, measured by the absolute difference between the sales price and the AVM value, divided by the sales price. To disentangle the direct and indirect effects of past redlining, we include $\widehat{HP}_{i,2018}$, the predicted house price from Eq. (1) which bakes in the impact of historic discrimination. Exterior condition rating, $ECR_{i,2018}$, captures property condition. Zhu, Neal, and Young (2022) found that properties with poor ECR ratings tend to have a greater percentage magnitude of AVM error. Following Neal et al. (2020), we include $N_{n,2018}$ as a vector of neighborhood characteristics including the standard deviation of neighborhood property ages, the percentage deviation of neighborhood home values, gentrification, the share of distressed home sales, and turnover rates. β measures the additional effect of past redlining on the prices paid by homebuyers.

4.2 Data

We use four data sources in this study: property transaction data from a major data vendor, 2014-18 American Community Survey, property condition data from CAPE Analytics,

and HOLC map data from the University of Richmond. The property data includes the AVM value, sales price, transaction dates, property address, land use, property type, and building information. The ACS data provides information about each census tract, such as its household racial and income composition. CAPE Analytics data provides exterior property condition (ECR) rating at the property level based on high-resolution images captured from airplanes. The ECR covers all a parcel's visible external features, including roofs, yards, driveways, and debris. The rating is measured on a five-point scale from severe to excellent (severe, poor, fair, good, and excellent).

Our sample period is 2018. For each arm's length single-family transaction in the Atlanta metropolitan area in 2018, we match the corresponding neighborhood-level ACS data. We also include measures of gentrification, distressed sales, and household income. Following Neal et al. (2020), we consider a neighborhood to be gentrified if it meets two criteria: the tract-level income is less than 70 percent of the income in the metropolitan statistical area (MSA) and the neighborhood (identified at the census tract level) experienced at least a 10 percentage-point increase in the ratio of tract-level income to MSA-level income over the year. In addition, we calculate turnover rates as the number of home sales in 2018 divided by the number of housing stock. To capture the heterogeneity of properties within neighborhoods, we also include the standard deviation of property ages and the percentage deviation of home values.

We then matched these property records with the ECRs from CAPE Analytics based on property latitudes and longitudes, assessor parcel lot numbers, and transaction dates. In this study, we collapse the five-point ECR scale from CAPE Analytics into three categories: poor (includes poor and severe), fair, and good (includes good and excellent). Using property latitudes and longitudes, we identify the HOLC grades for each property by merging with the 1940s

HOLC map of Atlanta from the University of Richmond. The HOLC score was a four-point scale: Best, Desirable, Declining, and Hazardous. We define a redlined area if the score was Declining or Hazardous. We also create a dummy variable indicating whether or not the property is located in a neighborhood within the historic HOLC map area, as a great number of transactions occurred outside the HOLC designated areas over the past several decades.

V. Results

Figure 2a and Figure 2b illustrate the geographic distribution of single-family transactions in 2018 by HOLC score in majority-Black and majority-white neighborhoods. The Atlanta metropolitan area is highly segregated by race. Transactions in majority-Black neighborhoods are disproportionately concentrated in historically redlined areas graded as Hazardous or Declining. For example, 59.2 % and 76.9 % of transactions in Declining and Hazardous areas respectively were in majority-Black neighborhoods. In contrast, 100 percent and 72.6 percent of transactions in Best or Desirable areas were from majority-white neighborhoods.

Table 1 displays sample descriptive statistics. Historically-redlined majority-Black neighborhoods have a significantly greater percentage magnitude of AVM error, 43.1 percent as compared to 31.1 percent for all majority-Black neighborhoods and 12.5 percent of all majority-white neighborhoods. Moreover, historically-redlined majority-Black neighborhoods have greater dispersion in both home prices and property age. These neighborhoods are also more likely to experience gentrification, have a higher share of distressed sales, and have a much lower average household income. Compared to majority-white neighborhoods, majority-Black neighborhoods have lower home values. The average property value of single-family homes sold in majority-Black neighborhoods in Atlanta CBSA was \$142,554 in 2018, while the average in

majority-white neighborhoods was \$353,568. Additionally, 46.5 percent of properties in majority-Black neighborhoods have a poor rating, higher than the 34.5 percent of majority-white neighborhoods.

Column (1) in Table 2 presents the results estimating the impact of past redlining on home values. After controlling for key variables at both property and neighborhood levels, our results show that both historical redlining and neighborhood race have a significant impact on home values. Specifically, holding other factors constant, home prices in historic redlined neighborhoods are associated with a 49.3-percent² undervaluation. Similarly, all else equal, compared with majority-white neighborhoods, home prices in majority-Black neighborhoods are 64.2 percent³ lower. This supports our first hypothesis that current home values are a reflection of historical redlining.

Table 3 summarizes the results estimating the direct impact of historic redlining on AVMs. We derive the predicted house prices from the first-stage regression and include them in the second-stage regressions, starting from column (2). Column (3) controls for the property condition measured by the ECR, and column (4) adds in a set of neighborhood-level controls. The results show that poor property conditions, greater dispersion in property values, and gentrification are associated with a greater percentage magnitude of AVM error.

More importantly, the coefficients on historic redlining and the Black neighborhood variable remain significant and positive after a full set of controls in column (4). This suggests that historic redlining has a direct impact on AVMs beyond the indirect impact through home values. The results in column (4) show that AVM inaccuracy for properties in historically-

² The magnitude of the coefficient is 0.401. $(\exp(0.401) - 1) * 100 = 49.3$

³ The magnitude of the coefficient is 0.496. $(\exp(0.496) - 1) * 100 = 64.2$

redlined areas is 5.7 percentage points greater. This suggests that for a home in a redlined area with a sales price of \$148,296, the AVM error will be \$8,453 greater than for a property with the same sales price in a non-redlined area, holding all other factors constant. This supports our second hypothesis that past redlining affects the prices paid by homebuyers.

VI. Discussion and Implications for Public Policy

Our results illustrate the pernicious role that historical racism plays even today in limiting home values and undermining the accuracy of our automated valuation model in majority-Black neighborhoods across Atlanta. By contributing to relatively lower home values, historical racism has reduced the housing wealth available in these communities contributing to the broader Black-white wealth gap. At the same time, the prospect of overvaluation, which partly reflects the impact of historical racism on property prices, could fuel a destabilizing downward correction in the future (PAVE 2022) and contribute to greater volatility of house prices in Black communities relative to white ones.

These findings call into **question the appropriateness of the sales comparison approach for determining a home's value**. Intuitively, this method of determining home value relies on past property sales thereby bringing historical racism to the present period. Alternative approaches might consider the income a property generates, or building costs (Mock 2021). These should be explored for their ability to close racial disparities in property value estimation.

The federal government has already taken steps to address the appraisal challenge. Most notably, the current Administration created the Interagency Task Force on Property Appraisal and Valuation Equity (PAVE) which has released an action plan that identified administrative actions to address appraisals and AVMs (PAVE 2022). In addition, the Consumer Financial Protection Bureau created a Small Business Review Panel on Automated Valuation Model (AVM) Rulemaking to **develop regulations for quality control standards for AVMs** (Consumer Financial Protection Bureau 2022). And both Fannie Mae and Freddie Mac have identified steps they will take in their equity plans to address appraisal challenges (Fannie Mae 2022; Freddie Mac 2022). Meanwhile, the Appraisal Subcommittee of the Federal Financial Institutions Examinations Council has recently issued a request for information about the supply and

diversity of appraisers (Appraisal Subcommittee 2022). Although these administrative steps have helped to raise the issue and prompt change, **legislative action can ensure that these changes are made permanent.**

Given the role that property condition also plays in determining the performance of AVMs, **expanding and improving access to renovation financing** could help to ameliorate inaccurate house price estimates. Our analysis suggests that fair or poor property condition increases the percentage magnitude of error relative to properties in good condition (Neal, Zhu and Young 2022). And research suggests that Black homeowners are more likely to live in inadequate housing relative to white homeowners (Neal, Choi, and Walsh 2022). However, the federal government's key levers for providing renovation loans can be challenging for borrowers to navigate (Goodman and Golding 2022).

Encouraging **direct investment through small businesses** would also improve conditions in majority-Black neighborhoods. Small businesses tend to reinvest in their local communities in support of job creation. However, research has demonstrated inequities in access these businesses have to federal small business capital (Theodos and Gonzalez-Hermoso 2021). But by harmonizing access across race and ethnicity, expanded capital access would lead to increased hiring, which should reduce the likelihood of distressed sales and increase household incomes.

6.1 Limitations and Future Research

These research results should motivate continued monitoring and action to rectify the wrongs brought on by historical racism. However, there are several limitations of our research that suggest additional analysis should be pursued. For example, we rely on one automated valuation model for the basis of our results but analyzing a suite of models will help to strengthen these conclusions. In addition, we focused on the city of Atlanta, but quantifying the role of historical racism in other cities as well can help to size the scope of the problem. Although we identified a single method for illustrating racial and ethnic differences in automated valuations, the computer science community has codified several metrics that can identify racial biases in automated systems.

References

- Aaronson, D., Hartley, D., & Mazumder, B. (2017). The effects of the 1930s HOLC "redlining" maps. Working Paper, No. 2017-12, Federal Reserve Bank of Chicago, Chicago, IL, Retrieved from <https://www.econstor.eu/bitstream/10419/200568/1/1010730592.pdf>
- Appel, I. (2016). Pockets of poverty: The long-term effects of redlining. SSRN Electronic Journal, 10.2139/ssrn.2852856.
- Appraisal Subcommittee. Request for information for data development – due July 22, 2022, Retrieved from <https://www.asc.gov/Pages/ViewWhatsNew.aspx?ID=184>
- Consumer Financial Protection Bureau (2022). Small business review panel on automated valuation model rulemaking. Retrieved from <https://www.consumerfinance.gov/rules-policy/small-business-review-panels/small-business-review-panel-automated-valuation-model-rulemaking/>
- Folk, J., & Chen, K. (2021). Avoiding overvaluation risk and appraisal bias in today's uniquely challenging market' session." Clear Capital, Retrieved from <https://www.clearcapital.com/avoiding-overvaluation-risk-and-appraisal-bias-in-todays-uniquely-challenging-market-session>
- Fannie Mae (2022). Equitable housing finance plan, Retrieved from <https://www.fanniemae.com/about-us/esp/social/equitable-housing-finance-plan>
- Freddie Mac (2021). Racial and ethnic valuation gaps in home purchase appraisals, Retrieved from <https://www.freddie-mac.com/research/insight/20210920-home-appraisals>
- Freddie Mac (2022). Racial & ethnic valuation gaps in home purchase appraisals - a modeling approach. Research Note. Retrieved from <https://www.freddie-mac.com/research/insight/20220510-racial-ethnic-valuation-gaps-home-purchase-appraisals-modeling-approach>
- Freddie Mac (2022). Equitable housing finance plan, Retrieved from <https://www.freddie-mac.com/about/pdf/Freddie-Mac-Equitable-Housing-Finance-Plan.pdf>
- Goodman, L. & Golding, E. (2022). Institutional investors have a comparative advantage in purchasing homes that need repair, UrbanWire. Retrieved from <https://www.urban.org/urban-wire/institutional-investors-have-comparative-advantage-purchasing-homes-need-repair>
- Hoffman J.S., Shandas V., & Pendleton, N. (2020). The effects of historical housing policies on resident exposure to intra-urban heat: A study of 108 US urban areas. *Climate*, 8(1):12. <https://doi.org/10.3390/cli8010012>
- House Canary (2021). Reducing racial bias in home appraisals using automated valuation technology. Retrieved from <https://www.housecanary.com/wp-content/uploads/2021/12/Reducing-Racial-Bias-in-Home-Appraisals-Using-Automated-Valuation-Technology-December-2021.pdf>

- Huang, S. J., Sehgal, N.J. (2022). Association of historic redlining and present-day health in Baltimore. *PLoS ONE*, 17(1).
- Interagency Task Force on Property Appraisal and Valuation Equity (2022). Action plan to advance property appraisal and valuation equity, Retrieved from <https://pave.hud.gov/actionplan>
- Krieger, N., Van Wye, G., Huynh, M., Waterman, P.D., Maduro, G., Li, W., Gwynn, R.C., Barbot, O., & Bassett, M.T. (2020). Structural racism, historical redlining, and risk of preterm birth in New York City, 2013-2017. *American Journal of Public Health*, 110(7), 1046-1053.
- Krieger, N., Wright, E., Chen, J.T., Waterman, P.D., Huntley, E.R., Arcaya, M. (2020). Cancer stage at diagnosis, historical redlining, and current neighborhood characteristics: Breast, cervical, lung, and colorectal cancer, Massachusetts, 2001-2015. *American Journal of Epidemiology*, 189(10), 1056-75.
- Massey, D. S., & Denton, N. A. (1993). *American apartheid: Segregation and the making of the underclass*. Cambridge, MA: Harvard University Press.
- Mitchell, B. & Franco, J. (2018). HOLC “Redlining” maps: The persistent structure of segregation and economic inequality, National Community Reinvestment Coalition. Retrieved from <https://ncrc.org/holc/>
- Mock, B. What it will take to close the race gap in home appraisals. Bloomberg, Retrieved from <https://www.bloomberg.com/news/articles/2021-03-03/appraisers-acknowledge-bias-in-home-valuations>
- Neal, M., Choi, J.H. & Walsh, J. Before the pandemic, homeowners of color faced structural barriers to the benefits of homeownership, Retrieved from <https://www.urban.org/research/publication/pandemic-homeowners-color-faced-structural-barriers-benefits-homeownership>
- Neal, M., Zhu, L., Young, C. (2021). Why automated home valuation technology errors disproportionately affect majority-black neighborhoods. Washington, DC: Urban Institute. Retrieved from <https://www.urban.org/urban-wire/why-automated-home-valuation-technology-errors-disproportionately-affect-majority-black-neighborhoods>
- Perry, V.G. (2020). 2020 State of Housing in Black America. Retrieved from <https://www.shiba2020.com/>
- Perry, V.G. and Martin, K. (2022). Algorithms for All: Has Digitalization in the Mortgage Market Expanded Access to Homeownership? Retrieved from SSRN: <https://ssrn.com/abstract=4126409> or <http://dx.doi.org/10.2139/ssrn.4126409>
- Rothwell, J. & Perry, A.M. (2021). Biased appraisals and the devaluation of housing in black neighborhoods. Washington, DC: Brookings Institute. Retrieved from <https://www.brookings.edu/research/biased-appraisals-and-the-devaluation-of-housing-in-black-neighborhoods/>

Theodos, B., Gonzalez-Hermoso, J. How can we overcome inequities in who owns small businesses?, Retrieved from <https://www.urban.org/urban-wire/how-can-we-overcome-inequities-who-owns-small-businesses>

Townsley, J., Andres, U.M., & Nowlin, M. (2021). The lasting impacts of segregation and redlining, SAVI, Retrieved from <https://www.savi.org/2021/06/24/lasting-impacts-of-segregation/>

Williamson, J., & Palim, M. (2022). Appraising the appraisal: A closer look at divergent appraisal values for black and white borrowers financing their home. Washington, DC: Fannie Mae. Retrieved from <https://www.fanniemae.com/media/42541/display>

Figure 1

Past redlining affects present-day automated home valuations and house prices.

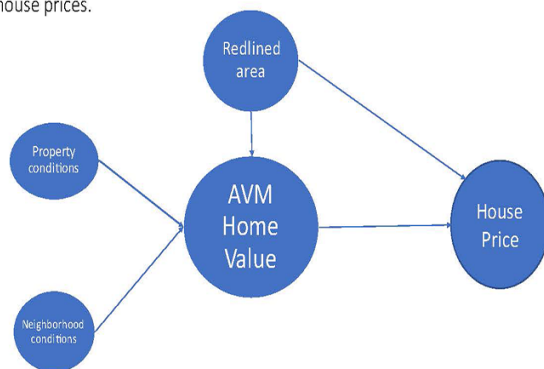
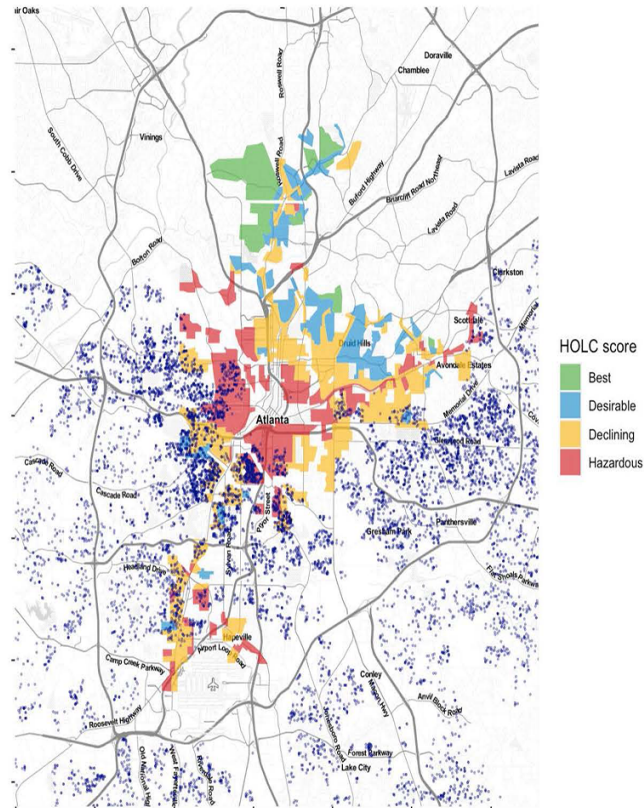


Figure 2a. The geographic distribution of single-family transactions by HOLC score, in majority-Black neighborhoods



Source: Property records data from a major data vendor, 2014-18 American Community Survey, and Mapping Inequality, the University of Richmond.

Notes: HOLC = Home Owners' Loan Corporation; Each dot represents one arm's length single-family transactions in 2018.

Table 1. Descriptive Statistics

	<i>Redlined Areas</i>		<i>Entire MSA</i>		
	Black Neighborhoods	White Neighborhoods	Black Neighborhoods	White Neighborhoods	
Percentage magnitude of AVM error	43.1%	15.1%	31.1%	12.5%	
Home Value	148,296	492,384	142,554	353,568	
Property age	75.0	72.8	45.0	35.0	
Standard deviation of neighborhood property ages	25.8	29.1	15.5	13.2	
Percent deviation of neighborhood property values	53.8%	35.0%	42.0%	33.7%	
Gentrified neighborhood	24.5%	9.5%	8.9%	2.2%	
Share of neighborhood distressed home sales	16.4%	5.4%	14.4%	4.6%	
Neighborhood median household income	32,216	90,961	48,501	95,472	
Turnover rate at neighborhood level	15.9%	7.8%	9.4%	7.7%	
Exterior Condition Rating (ECR)					
	Good	7.2%	16.1%	8.7%	13.3%
	Fair	53.0%	53.8%	44.8%	52.2%
	Poor	39.8%	30.1%	46.5%	34.5%

Source: Authors' calculations based on property records data from a major data vendor, 2014-18 American Community Survey, CAPE Analytics, and Mapping Inequality, the University of Richmond.

Notes: Redline areas are defined as areas with HOLC scores as Hazardous or Declining.

Table 2. First-Stage Regressions: House Price

VARIABLES	First-stage (1)
Redlined	-0.401*** (0.016)
Black neighborhood	-0.496*** (0.01)
Control Variables	Y
County fixed effects	Y
Observations	50,111
R-squared	0.754

Notes: In column (1), the dependent variable is the log value of sales price. Control variables include the log value of lot sizes, numbers of bedrooms and bathrooms, built year, exterior property condition, whether located within the historic HOLC map area, and the log value of neighborhood-level median household income. (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3

Table 3. Second-Stage Regressions: Percentage Magnitude of AVM Error

VARIABLES	Dependent Variable: Percentage Magnitude of AVM Error			
	(1)	(2)	(3)	(4)
Redlined	9.641*** (0.685)	10.084*** (0.676)	10.069*** (0.676)	5.654*** (0.786)
Black neighborhood	18.057*** (0.326)	3.408*** (0.504)	3.434*** (0.504)	2.168*** (0.513)
Predicted Home Price		-16.071*** (0.425)	-15.864*** (0.426)	-13.832*** (0.445)
ECR: Fair			3.664*** (0.51)	3.592*** (0.51)
ECR: Poor			3.630*** (0.529)	3.750*** (0.527)
Standard deviation of neighborhood property ages				0.038 (0.028)
Percentage deviation of neighborhood property values (%)				0.375*** (0.017)
Share of neighborhood distressed home sales (%)				-0.013 (0.010)
Gentrified neighborhood				0.044*** (0.007)
Neighborhood-level turnover rate (%)				-0.020 (0.045)
Observations	50,111	50,111	50,111	50,111
R-squared	0.066	0.092	0.093	0.104

Notes: In columns (1) to (4), the dependent variable is the percentage magnitude of AVM error. Predicted home price is derived from the first-stage regression. (***) $p < 0.01$, (**) $p < 0.05$, (*) $p < 0.1$.



January 30, 2024

The Honorable Tina Smith
Chair
Subcommittee on Housing, Transportation, and Community Development
US Senate
720 Hart Senate Office Building
Washington, DC 20510

The Honorable Cynthia Lummis
Ranking Member
Subcommittee on Housing, Transportation, and Community Development
US Senate
127 Russell Senate Office Building
Washington, DC 20510

Dear Chair Smith, Ranking Member Lummis, and members of the Subcommittee,

We write to thank you for holding the hearing on Artificial Intelligence (AI) and Housing: Exploring Promise and Peril and to submit for the record this letter outlining Zillow's work in responsible AI that benefits consumers, our partners, and employees. We also thank the Subcommittee on Housing, Transportation, and Community Development for its focus on the benefits that AI poses for consumers and its attention to mitigating risks.

AI at Zillow

Artificial intelligence has been central to Zillow since our founding in 2006 with the launch of the Zestimate feature, marking the first time consumers could find an instant, free estimate of their home's value using details about the home, property tax assessment information, and sales histories. As our use of technology evolves, our pro-consumer approach to innovation remains at the core of our mission to make home a reality for more and more people. Given the hundreds of millions of people searching for housing online, it's vital that consumers are equipped with safe, trustworthy AI tools that are used responsibly to make the process of searching for, renting, buying, or financing a home more equitable and transparent—and less daunting.

AI for All Consumers

We strive to use AI to help make renting, buying, selling, and financing easier, affordable, and safer for all consumers. Using our AI tools, we aim to help consumers



understand what's possible, empowering them to navigate the real estate process without obligation or judgment, making housing more accessible and equitable, and helping to close gaps that have persisted for too long. We're using AI and machine learning to power features like 3D Home tours and immersive floor plans, opening up homes to people who cannot be there in person to tour. We are also leveraging Human-In-The-Loop (HITL) as an important part of AI model/experience development and deployment to generate appropriate and high quality training labels and ensure high model output quality.

We've also deployed AI to help consumers search for homes on Zillow using natural language like they are talking to a friend. We continually seek opportunities to leverage AI to power numerous pro-consumer advancements in real estate, including improved property searches, enhanced property valuations, immersive virtual property tours, and streamlined paperwork and transactions.

We place a great emphasis on transparency and fairness in the housing industry, which has a long history of discrimination and bias against people of color. When built and deployed responsibly, AI can help address *appraisal bias* when leveraged for remote home appraisals, particularly in unlocking wealth for home sellers of color and making the appraisal process more accurate, transparent, and fair. AI can reduce potential biases compared to within-person appraisers, who may intentionally or unintentionally consider the demographics of the homeowner, neighborhood, or surroundings.

We've improved our models that expand *housing choice*—using AI to broaden people's search and reduce bias, which can result in consumers seeing more homes from more geographically diverse areas to ensure all consumers get a fair and accurate picture of what homes are available to them. Creating transparency and openness in the housing market is a top priority for Zillow, and we will continue to responsibly utilize AI to help ensure every buyer, seller, and renter has all the information available to them.

AI for Partners

We also believe AI can enable partners in the housing journey to provide greater value to customers. Great real estate agents are critical to helping customers get home. We strive to use AI to help agents become even more productive and successful, with tools to help scale their operations, automate mundane tasks, and access data that matters so they can focus on delivering what only they can: the human component.



AI for Employees

AI can empower employees to do more and go further. Employees are the backbone of Zillow, and they choose to work at our company in part because of our proven commitment to responsible innovation. We use AI to empower employees to continue shaping our industry with greater efficiency and excellence and to, in the process, accelerate their own career development as the nature of work evolves.

Zillow's Responsible AI Commitments

In all of our AI work, we are guided by the following principles to inform the responsible development, evaluation, and deployment of AI technologies.

- *Fairness.* We prioritize fairness in the design, evaluation, and deployment of AI systems across Zillow, understanding the distinct influence the models we use have on housing opportunities. Therefore, we test and review our training data and models to help create a fairer marketplace. We strive to continuously improve our models to help ensure that they empower everyone, including historically marginalized communities and members of protected classes, by minimizing and actively avoiding recreating existing patterns of discrimination and segregation.
- *Safety & Accountability.* We strive to apply standards and best practices for collecting, using, and validating both internal and external data used in our models and apply prevailing standards to assessing third-party models we may utilize. As AI technology continues to develop, we will work to be leaders in its responsible use in all contexts, keeping in mind both new opportunities and new risks. We aim especially for a deep understanding of our domain and the impacts our AI systems have on our customers, their privacy, the security of their data, and their housing journey. We follow rigorous privacy and security policies and standards to ensure customer, partner, and employee data remains secure. When integrating external data into our models, we conduct thorough due diligence to assess its reliability and accuracy. We actively share our progress, learnings, and areas for improvement within Zillow to foster trust and accountability.
- *Transparency.* One of Zillow's founding principles is to empower people with information in what used to be an incredibly opaque and daunting process of searching for a home. We have the same opportunity with the AI systems we are building. We believe it is important to ensure that consumers are aware of when



interacting with AI tools developed by Zillow and others and that they understand how Zillow uses these tools to improve their experience. We strive to empower customers to understand and decide whether and how their data is used to shape their experience on our platform. We are dedicated to fostering responsible AI dialogues and spearheading meaningful conversations on effectively and ethically integrating AI into the real estate industry to create internal efficiencies and new consumer tools. We actively engage with industry participants, lawmakers, nonprofit organizations, research institutions, and other stakeholders to shape and inform pro-consumer AI public policies.

- *Inclusiveness.* We build AI systems capable of understanding consumers', partners', and employees' unique needs, avoiding broad generalizations that might marginalize segments of our society. We evaluate our AI systems from diverse customers' perspectives and continuously work to improve our AI systems to serve our customers' needs throughout all stages of their home journey. We proactively identify AI models and systems where customers might be treated and prioritized differently. Through monitoring and evaluation, we test and analyze our AI technologies to mitigate the risk of bias and discrimination and provide equal opportunities and experiences for all.

Generative AI and Large Language Models in Real Estate

The many ways in which we have deployed AI to empower our customers create an obligation for us to consider and share some of the risks and pitfalls of using AI, including discrimination and bias. Given the hundreds of millions of people searching for housing online across our Web and app experiences, it's vital that consumers are equipped with responsible, safe, and trustworthy tools that make the challenging home-purchase process equitable and transparent.

Specifically, with the recent advancements in Generative AI, such as LLM technology, we have also been exploring how to leverage it to empower consumers, partners, and employees. However, applying this technology in the real estate domain requires additional optimization. In particular, off-the-shelf LLM models need to be made aware of Fair Housing and Fair Lending guidelines, and additional AI guardrails need to be deployed to uphold the standards we expect from human operators in the real estate transaction process. To share our work and learnings, Zillow published research on an [LLM Fair Housing Guardrails service](#) that combines three responsible AI strategies into a single comprehensive system. These strategies include: i) a "stop list" of potentially discriminatory or offensive language, ii) prompt engineering instructions to influence the



behavior of the LLM, and iii) a Fair Housing Compliance classifier, which is a machine learning model trained to detect potential FHA violations.

The guardrails system is a detective control and a crucial element in the overall LLM application. It can be integrated both as a pre-processing and post-processing component. For example, the system can monitor for a user's input likely to elicit an output violating fair housing requirements. This allows for the early detection and filtering of potentially non-compliant requests, helping maintain the system's integrity. Concerning post-processing, it can remove any offensive or discriminatory content to promote fairness and equal opportunity in real estate.

As part of our commitment to fostering an ethical, fair, and transparent real estate industry, we are committed to open-sourcing our classifier and the supporting data to encourage the community towards contribution and collaboration, which could lead to further innovation and expanded applications. This would also allow other organizations to leverage our work to help maintain compliance with fair housing requirements and ensure equitable treatment for all.

Thank you for the opportunity to share Zillow's commitment to and deployment of responsible AI in housing that benefits consumers, partners, and employees. We would welcome the chance to speak with you in more detail.

Sincerely,

A handwritten signature in black ink, appearing to read "Josh Weisberg". The signature is fluid and cursive, with a long horizontal stroke at the end.

Josh Weisberg
Senior Vice President, Artificial Intelligence
Zillow Group