



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:

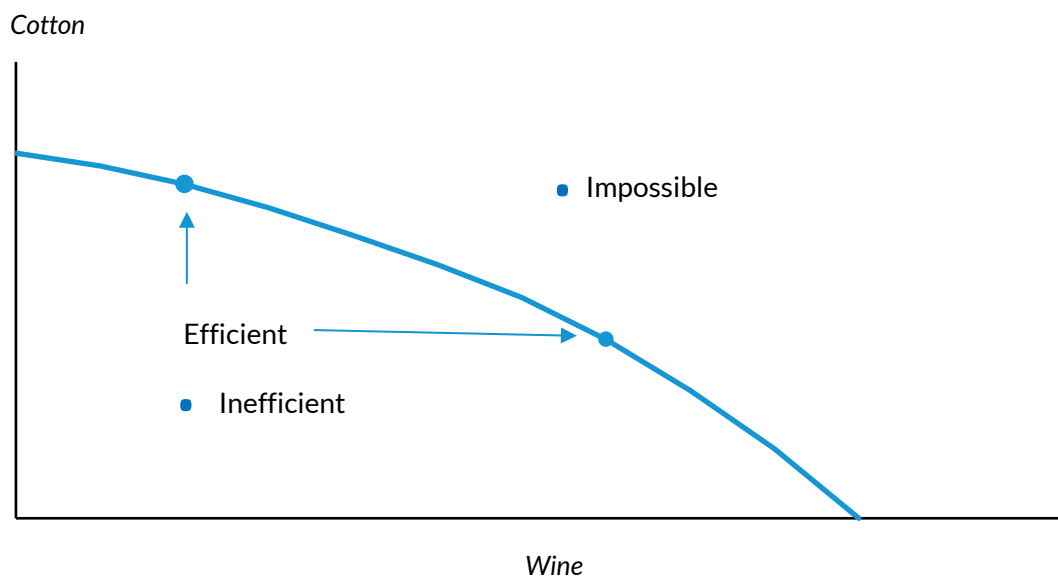
- **Intentional design.** Designers of AI models should determine whether the inputs and the underlying data that models learn from reflect the housing market’s disparate racial impacts. Prioritizing equity must happen from the start, not as an afterthought. MDIs and CDFIs provide a blueprint for detecting 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



URBAN INSTITUTE

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.

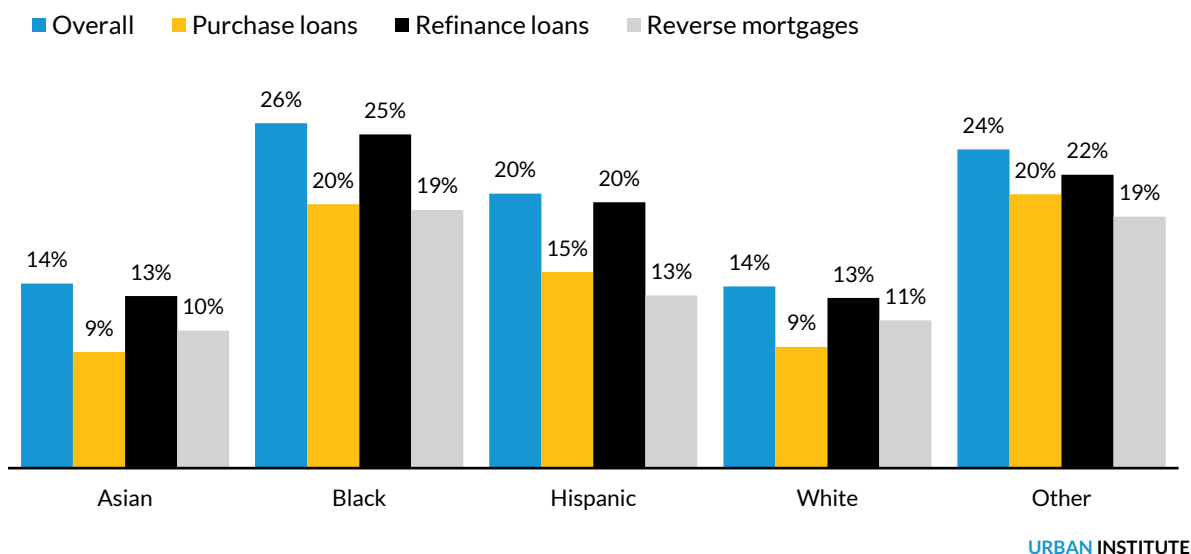
Several factors limit success for people of color throughout the entire mortgage cycle. These disparities are driven by factors that appear to be race neutral but in fact reflect and proxy for the results of a long history of racial discrimination.⁹ This discrimination, rooted in structural racism, has 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.¹³ 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: Geogfrey 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.⁵⁸ 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-a-i/>.
- 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 Foos, 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 Foos, 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 Fairle. 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, Snighda 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