

# **“Artificial Intelligence and Housing: Exploring Promise and Peril”**

**Testimony of Nicholas Schmidt  
Partner and Artificial Intelligence Practice Leader at BLDS, LLC and  
Founder and CTO of SolasAI**

**Before the United States Senate Committee on Banking, Housing, and Urban Affairs,  
Subcommittee on Housing, Transportation, and Community Development**

**January 31, 2024**

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

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

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

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

## **Defining Artificial Intelligence**

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

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

In practical terms, ML and AI represent a class of mathematical algorithms that learn patterns and rules from data. These learned rules may then be applied to new data to inform real-world decisions. Thus, it is important to remember that the rules AI develops are based on mathematical (i.e., not human) insight, but that those rules are developed on historical data that encode many types of human biases. However, as I discuss below, despite the computer developing the rules, human decisions affect how the rules are developed and systems are used. Understanding this is essential for writing effective legislation.

### **Human Decisions Drive Algorithms**

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

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

---

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

1. Choosing what the model will define as a delinquency. A data scientist might define delinquency as 60 days or 90 days of non-payment. For people with less income security, but who are likely to be able to pay their bills over a longer period, choosing 60 days instead of 90 days may make the difference between being provided a loan or being rejected. Importantly, this decision is an entirely human-based decision.
2. Choosing which data will be used to predict delinquency. The computer only makes its rules based on the data it is provided. The person building the model might only include data clearly related to delinquency (e.g., the existing level of debt, past payment history, etc.), or they may include data that is not clearly causally related to delinquency (e.g., education or purchase history). The choices that the modeler makes will affect the fairness of the model, its accuracy, its reasonableness, and its reliability. While algorithms will choose how to weight different data (and possibly exclude certain data altogether), choosing which data to provide is an entirely human-based decision.
3. Choosing the type of algorithm that will build these rules. Model developers have many options regarding what model architecture they will use to develop the model (i.e., how the model will learn from the data). These include architectures like deep neural networks, gradient-boosted trees, or traditional linear regression. This decision has direct implications in terms of the transparency of the model and whether the model's decisions will be explainable. Some kinds of models, such as regression models, are inherently explainable. This means it is possible to know exactly how the model arrived at a decision.<sup>2</sup> Others, such as neural networks, are

---

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

not transparent. The decisions surrounding model architecture have practical and legal implications, as I explain *infra*.

4. Choosing whether the algorithm is sufficiently accurate for the task. No algorithm is perfect. One issue is that a model might be very effective for some people, but not others. For example, it might be highly accurate for high-credit quality individuals but be far less accurate for those with lower credit quality. A person working for the lender will ultimately make a decision whether that trade-off is acceptable before putting the model into production. More generally, whether to test for different forms of inaccuracy, how to balance the varying costs of inaccuracy, and what minimum accuracy requirements are required are all human decisions.
5. Choosing how to implement and use the algorithm. Generally, an AI or ML model does not make a decision on its own. Frequently, only a subset of applicants will be scored by the model; which applicants are scored is a human decision. Once applicants are scored, what that score means and how it is used must be determined. For example, will a cutoff be used, or will the specific model work with another model or subjective rules to make a decision? Will the score be considered in light of other variables or factors? All of these are human decisions. In fact, when all of these decisions are examined together, it is clear that the entire model system, of which the model itself may only play a relatively modest role, is largely made up of interlocking human decisions that result in an ultimate decision.

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

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

---

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

defined and reasonable regulations would give companies more clarity on how they can safely use AI systems.

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

### **The Use of AI Systems in Housing**

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

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

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

However, before discussing how the history and practices of regulating and monitoring the use of algorithms can be extended and made better in the world of AI, I will discuss a particular use case of algorithmic decisioning - automated valuation models - because they

represent a concrete example of many of the benefits and challenges inherent in the expanded use of algorithms and AI.

### **Case Study in the Use of Algorithms: Automated Valuation Models**

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

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

AVMs' obvious and clear benefits are their speed, availability, and price. The National Association of Realtors (NAR) reported that, in 2022, the average cost of a home appraisal was \$500, with 9% of appraisals costing more than \$800.<sup>7</sup> Once an AVM is built and running, the cost of calculating an appraisal is virtually zero for the model owner (though borrowers could still be charged for this service). NAR additionally reported that the median response from realtors indicated that it takes 14 days for a lender to return a

---

<sup>4</sup> <http://tinyurl.com/3w5s2dak>.

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

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

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

completed appraisal, with 4% of realtors reporting that it typically takes them more than 30 days to receive a completed appraisal. In comparison, obtaining an appraisal from an AVM is virtually instantaneous; even if the lender performs subsequent reviews of the AVM-based appraisal, an AVM-based appraisal is likely to be available substantially faster than a human-based appraisal. It is likely that the use of AVMs has already ensured that home sales have closed faster, instead of languishing or being rejected.

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

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

To summarize, the discriminatory effect of AVMs is likely to be mixed. Because an AVM does not have access to or incorporate information about individual sellers or borrowers, it cannot discriminate based on those factors. However, it will still incorporate price-affecting discrimination in the data it sees. As a result, we can view AVMs as a tool to help fight discrimination in home pricing, but not one that is complete or without significant residual problems.

My work, and the software we have developed at SolasAI, measures bias and discrimination in models such as AVMs. As explained above, data may be biased due to

---

<sup>8</sup> <http://tinyurl.com/2jt4vtrs>.

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

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

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

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

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

---

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



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

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

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

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

### **The Benefits of a Non-Prescriptive Regulatory Environment**

A principles-based and less prescriptive approach to AI regulation can encourage innovation while ensuring the responsible development of AI. It recognizes the dynamic nature of technology and compliance and provides flexibility for continuous improvement and adaptation.

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

responsible model building. A less rigid but strong regulatory environment encourages such advancements.

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

### Key Principles to Consider for AI Regulation

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

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

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

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

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

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

### **Using Existing Regulations and Frameworks to Guide Further AI Regulation**

Congress and regulators will not need to devise laws and regulations from scratch to achieve the objectives that I have laid out. There are many regulations, standards, and frameworks with a proven track record of setting standards for human decisions related to AI, holding relevant actors accountable for those standards, and supporting the development and deployment of these technologies. Importantly, in many (but of course not all) cases, the industry has welcomed these for providing clear and reasonable standards. I discuss these at a high level below.

- **SR 11-7**<sup>11</sup>, a supervision and regulation letter from the Board of Governors of the Federal Reserve System, constructs accountability mechanisms and organizational structures to ensure adequate and risk-based governance of credit modelers. While the document highlights the importance of technical processes such as testing and monitoring, its primary focus is on principles such as effective governance structures (e.g., independent validation teams with high stature and strong financial incentives), risk management executives with independent reporting chains, and documentation requirements.
- **NIST SP 1270**<sup>12</sup>, a special publication from the National Institute of Standards and Technology (NIST), describes technical, process, and cultural problems and solutions relating to AI bias. It highlights that many aspects of data and AI systems are strongly influenced by human behavior and decisions, and suggests that

---

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

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

approaches from model risk management (e.g., SR 11-7) coupled with more novel approaches, such as structured feedback activities (e.g., bug bounties or red teaming), human-centered design, and information sharing are strong mitigants for managing bias in AI systems. Another prominent theme of NIST SP1270 is that basic scientific rigor in AI needs to be improved.

- The **disparate impact, disparate treatment, and proxying framework** is a legal doctrine that has been developed over the course of decades under the Fair Housing Act (FHA), the Equal Credit Opportunity Act (ECOA), and Title VII of the Civil Rights Act of 1964. This framework sets forth requirements for measuring discrimination that could be used in any decision tool (including AI or ML models) and provides a conceptual framework for mitigating any meaningful disparities found through conducting an LDA search. Many AI tools that affect consumers in the housing market are likely covered by this framework via the FHA. Setting the expectation that this framework would be extended to all high-risk AI use cases throughout the housing industry would ensure that companies move towards adopting fairer models.<sup>13</sup>
- The **NIST AI Risk Management Framework**<sup>14</sup> puts forward four risk management functions for organizations: (1) Govern, (2) Map (understanding the risk of AI systems in their operational contexts with less emphasis on their development), (3) Measure, and (4) Manage; and seven trustworthy characteristics for AI systems: (1) Safe, (2) Valid and Reliable, (3) Accountable and Transparent, (4) Explainable and Interpretable, (5) Privacy-enhanced, (6) Fair with Harmful Bias Managed, and (7) Secure and Resilient. Governance guidance is largely aligned with the risk-based principles laid out in SR 11-7 but introduces additional governance concepts from data privacy, information security, and more recent academic research. The AI Risk Management Framework has two distinct strengths: it acknowledges (1) that many AI risks arrive from real-world problems, not computer code bugs; and (2) that overlaps and connections between risks, risk controls, and governance must be recognized. It does all this while orienting governance toward traditional risk-based principles focusing on human accountability.

---

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

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

## **Conclusion**

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

Regulating AI so that it is used responsibly and safely and so that it can continue to benefit consumers and industry is a complex and crucial task. Materiality, fairness, accountability, and transparency provide a framework of principles that can serve as a touchstone for evaluating the effectiveness of proposed regulations. Further, leveraging existing regulations and frameworks such as SR 11-7, NIST SP 1270, the disparate impact framework, and the NIST AI Risk Management Framework, regulators and legislators can create an environment that fosters innovation and protects against potential harm.

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

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