Testimony of Makada Henry-Nickie,
Fellow, Race, Prosperity, and Inclusion
Brookings Institution

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

Hearing on
"Equitable Algorithms: Examining Ways to Reduce AI Bias in Financial Services"

February 12th, 2020

Chairman Foster, Ranking Member Hill, and distinguished members of the Task Force on Artificial Intelligence. Thank you for the invitation and opportunity to testify before the Committee on equitable algorithms and algorithmic bias. I am Makada Henry-Nickie, Fellow at The Brookings Institution, my research covers consumer financial protection and labor market impacts of new technologies. My comments today will focus on AI-driven benefits, algorithmic bias in financial services, and algorithmic oversight, my comments are my own and do not reflect any official Brookings position. I hope my contribution furthers the Committee's understanding of the opportunities and challenges facing consumers, regulators, financial institutions, and the scientific community as the integration of AI into financial products and services accelerates. Forward-leaning congressional leadership on the implications of artificial intelligence is critically important to ensuring that emerging technologies interact with consumers responsibly and deliver inclusive benefits to all members of our society.

Market Interest in Artificial intelligence is Soaring

Artificial intelligence has permanently reshaped the financial services marketplace and altered consumers' behaviors, preferences, and their expectations of financial institutions. Consumers, digital natives in particular, are increasingly open to engaging with their financial institutions through AI-based interfaces. According to Adobe Analytics, 44% of GenZ and 31% of millennials have interacted with a conversational interface (chatbot); surprisingly, these generational segments overwhelmingly preferred interacting with a conversational bot to a human representative.[1]

Shifts in consumer preferences underpin the business case for further AI investments. A two-year study of Bank of the West customers revealed that once customers opened digital bank accounts, they were nearly 60% more likely to embrace other products including, credit cards, mortgage

[1] Adobe Analytics: How Different Generations Bank (2019). Available at: https://theblog.adobe.com/adobe-analytics-research-how-different-generations-bank/.

refinancing, and personal loans and increase annual revenues.[2] While systematic evidence of the industry's return on investment is sparse, periodic surveys such as the Narrative Sciences reported that the financial services industry's investment in AI technologies grew by a remarkable 60% between 2016 and 2017.[3] It's clear from the sector's increasing investments in artificial intelligence that deep learning and machine learning technologies will continue to have a substantial influence on the consumer financial market.

Beyond the success of conversational interfaces, incumbent banks, nonbanks, fintechs, and insurance companies employ deep learning across a diverse array of business verticals. Nascent AI anomaly detection models have been successfully deployed to detect various types of fraud, including payments, transactions, and loans. Tech startups, such as TrueAccord and Collectly, are experimenting with machine learning software to reinvent the debt collection experience.[4]

AI has similarly penetrated traditional domains such as target marketing, lead generation, and credit decisioning. Though, loan underwriting algorithms, in particular, have received intense public scrutiny fueled in large part by egregious cases of algorithmic discrimination. Most recently, a software developer revealed that Apple's branded credit card, issued by Goldman Sachs, was gender-biased after Apple denied his spouse's application for a credit line increase, despite her higher credit score and similar income.[5] The tech company's misstep is not altogether uncommon. In March 2019, the Department of Housing Urban Development (HUD) sued Facebook for discriminatory marketing practices. HUD alleged that Facebook's allowed advertisers marketing housing services, including mortgage lenders and rental agents, to selectively curate target audiences and exclude protected groups. The tech giant's Custom Audiences and Lookalike tools allowed advertisers to explicitly exclude certain groups based on protected characteristics such as "women in the workforce" or "foreigners" in violation of the Fair Housing Act (FHA).[6]

Instances of algorithmic discrimination transcend financial services, as notable and disconcerting examples can be found in other domains from Amazon's biased hiring algorithm to Google Photo's image classifier that associates blacks with images of gorillas.[7] Discrimination or bias that systematically disadvantages minorities was a recurrent theme in each successive instance. What's clear from these illustrative cases is that artificial intelligence algorithms can adversely impact minority groups and exacerbate disparities. Consequently, it is of paramount importance that policymakers, regulators, financial institutions, and technologists critically examine the benefits, risks, and limitations of artificial intelligence and proactively design safeguards against algorithmic harm, in keeping with societal standards, expectations, and legal protections.

---

[2] Panno, Kelsey, S&P Global, "Study Finds Digital Banking Adoption Leads to More Valuable Customers." Jun. 22, 2016. Available at: https://www.spglobal.com/en/research-insights/articles/study-finds-digital-banking-adoption-leads-to-more-valuable-customers

[3] Narrative Science. (2018). Research Brief: The Rise of AI in Financial Services. Retrieved from https://narrativescience.com/wp-content/uploads/2018/11/Research-Report_The-Rise-of-AI-in-Financial-Services_2018.pdf

[4] Chin, C. (2018, September 4). Silicon Valley Wants to Use Algorithms for Debt Collection. *Wired*, Retrieved from https://www.wired.com/story/silicon-valley-algorithms-for-debt-collection/

[5] Vigdor, N. (2019, November 4). Apple Card Investigated After Gender Discrimination Complaints. *New York Times*, Retrieved from https://www.nytimes.com/2019/11/10/business/Apple-credit-card-investigation.html

[6] The United States Department of Housing and Urban Development, on behalf of the Assistant Secretary for Fair Housing and Equal Opportunity v. Facebook. (March 28, 2019). Retrieved from https://www.hud.gov/sites/dfiles/Main/documents/HUD_v_Facebook.pdf

[7] Madonik, R. (2018, January 11). When it Comes to Gorillas, Google Photos Remain Blind. *Wired*. Retrieved from https://www.wired.com/story/when-it-comes-to-gorillas-google-photos-remains-blind/

Artificial intelligence can improve the financial lives of consumers

The story of AI in financial institutions is not all bad; innovative fintechs have made salient contributions that make financial services more inclusive and more accessible for consumers. AI-enabled financial applications that enable consumers to accumulate savings incrementally or micro-invest are often overlooked in financial inclusion dialogues. Still, these emerging tech-enabled solutions can embody valuable financial inclusion tools. Through micro-savings apps, fintechs and legacy financial institutions have empowered millions of consumers to save more and to do so automatically. For example, Digit—an AI-enabled micro-savings app—uses machine learning algorithms to analyze checking account transactions and identify micro-savings patterns that enables its users to save consistently. Since its launch, Digit's users have saved over $2.5 billion; the company reports that its "auto-save" algorithm empowers users to save an average of $2,000 annually.[8]

In credit markets, a combination of machine learning models and high-dimensional alternative data is generating cautious interest in the ability of algorithms and Big Data to expand access to credit and decrease historical disparities. According to a 2018 study of German consumers, digital factors such as phone operating system, type of internet service or format of an email address can reliably predict loan default.[9] In the U.S., practical applications of alternative data do not routinely include such controversial factors. Instead, alterative data are more likely to include variables not traditionally incorporated into conventional credit underwriting models. For instance, rental payments, utility bills, or deposit transaction histories can fill critical gaps in assessing an applicant's ability to repay a loan and predict the likelihood of default.[10] The Consumer Financial Protection Bureau's (CFPB) publicly released findings from its No Action Letter (NAL) review of Upstart Network Inc's lending algorithm and use of alternative data. Upstart is the first CFPB-approved fintech lender sanctioned to use alternative data in its underwriting models. According to CFBP's audit, loan approval rates increased by nearly 30% for some customer segments, and the price of credit was dramatically lowered as APRs decreased, on average, between 15 and 17%. What's more, the Bureau reported that Upstart's data showed no evidence of fair lending disparities for members of protected classes.[11]

Additionally, findings from two recent studies show that algorithmic lending, directly and indirectly, expands access to credit. On the one hand, algorithmic peer-to-peer (P2P) lending indirectly enhanced access to credit by creating positive credit report signals resulting in increased loans from traditional banks to P2P borrowers.[12] Meanwhile, a UC Berkeley study found that algorithmic lending substantially decreased pricing disparities and eliminated underwriting discrimination for African-American and Hispanic borrowers.[13] Crucially, these emerging research

---

[8] Celebrating Digit's Impact and Funding. (2019, September 30). Retrieved from https://blog.digit.co/

[9] On the Rise of FinTechs – Credit Scoring using Digital Footprints Tobias Berg, Valentin Burg, Ana Gombovi, and Manju Puri NBER Working Paper No. 24551 April 2018, Revised July 2018 JEL No. D12,G20,O33

[10] FinReg Lab (2019) The Use of Cash-Flow Data in Underwriting Credit. Available at: https://finreglab.org/wp-content/uploads/2019/07/FRL_Research-Report_Final.pdf

[11] Patrice Ficklin and Paul Watkins, *An Update on Credit Access and the Bureau's First No-Action Letter,* Consumer Financial Protection Bureau (Aug. 6, 2019). Available at: https://www.consumerfinance.gov/about-us/blog/update-credit-access-and-no-action-letter/

[12] Balyuk, Tetayana (2019) Financial Innovation and Borrowers: Evidence from Peer-to-Peer Lending. Available at: https://www.fdic.gov/bank/analytical/fintech/papers/balyuk-paper.pdf

[13] Bartlett, Robert et. al (2019) Consumer Lending Discrimination in the Fintech Era. Available at: https://faculty.haas.berkeley.edu/morse/research/papers/discrim.pdf

studies show that despite the risks, algorithmic models have the potential to provide benefits to consumers.

Algorithms propagate bias

While artificial intelligence has the potential to deliver tremendous benefits that improve consumers' financial lives, algorithmic decision-making is inherently risky and susceptible to bias. The question of bias is less than straightforward and can be frustratingly complex and difficult to disentangle. Bias from a technical perspective affects an algorithm's accuracy or ability to make correct predictions about the real world based on its experience with training or example data.[14] But the definition, though precise, does not capture an intuitive, societal interpretation of the biased exclusions and unfair outcomes described in news stories about Google, Facebook, and Amazon that or discriminatory conduct that violates civil rights and equal opportunity statutes.

Tracing the sources and transmission mechanisms of bias is critical to informing the design of technological and policy solutions to reduce biased outcomes. Machine learning research has established an unambiguous link between biased outcomes and flawed training data. Class imbalance bias stemming from minority underrepresentation or selecting sample data with distorted distributions, such as systematic discrimination, can introduce selection bias into the modeling process.[15] These sources of bias might explain why the UC Berkeley study found that while algorithmic lenders did not discriminate against minority applicants in their underwriting decisions, they systematically charged them higher interest rates. This result is inconsistent with CFPB's NAL fair lending conclusions that algorithmic lending was associated with more equitable pricing. On the contrary, the Berkeley study confirmed that algorithmic lending perpetuated discriminatory pricing practices—Hispanic and African American borrowers paid 5.3 basis points more in interest than their white counterparts. [16]

In the final analysis, machine learning algorithms capable of producing equitable outcomes were not sophisticated enough break the logically flawed statistical correlation between race and credit denials or supplant the biased effects of decades of explicit racist housing policies. Algorithmic bias has tangible opportunity costs, by the researchers' estimation, minority borrowers pay an estimated $765.0 million in excess interest payments annually, instead of saving or paying down student loan debt.

Machine learning bias is fluid and can shift in response to changes in underlying data or design processes and hence requires a flexible and vigilant ecosystem of safeguards to ensure that artificial intelligence delivers on its full potential. CFPB's declaration of Upstart's success as a potential equitable algorithm should not be regarded as an endorsement of a bias-free endeavor. Dataset shifts from non-stationary data distributions or changes in a neural network's activation function can potentially bias an algorithm over time. This hypothetical outcome is not entirely implausible, experimental alternative data such as educational background variables are early in their deployment and need critical ground-truth datasets to benchmark accuracy. Public education datasets have documented coverage gaps in certain variables such as college major.

[14] Jindong Gu: "Understanding Bias in Machine Learning", 2019, 1st Workshop on Visualization for AI Explainability in 2018 IEEE Vis; [http://arxiv.org/abs/1909.01866 arXiv:1909.01866].
[15] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez: "Fairness Constraints: Mechanisms for Fair Classification", 2015; [http://arxiv.org/abs/1507.05259 arXiv:1507.05259]
[16] Ibid

Machine learning bias is neither inevitable nor final. And algorithmic bias is not benign. Algorithmic decision-making has enormous systemic social and economic consequences for affected racial, gender, and sexual minorities; these effects should not be ignored or trivialized.

Algorithmic Oversight is Indispensable for Consumer Financial Protection.

AI-enabled financial technologies are relatively nascent and primarily involve weak or narrow forms of AI. However, financial institutions are increasingly experimenting with advanced deep learning neural networks that are second to none in fitting high volumes of data to extraordinarily accurate predictive functions. Lamentably, deep learning's opaque "Black Box" effect even challenges AI experts when asked two fundamental questions: How? And. Why? These questions are central to human intuition and our cognitive ability to understand and negotiate our environment.

Beyond philosophical musings, our regulatory system rests firmly on a framework that assesses accountability through a causal lens, to which the answers to the questions of *how* and *why* are crucial for the system to function and effectively serve and protect American consumers. In a causal system, explanations have semantic significance and assist in making connections between reckless judgments or honest mistakes and unfair outcomes. Regulators need to be clear-eyed about an institutional agent's intent to assess the extent of its liability; without clear, rational explanations and clear causal connections between discriminatory outcomes and decision processes, the accountability framework becomes unstable and dysfunctional.

Understandably, AI's black-box effect underpins a growing chorus of calls for intuitive AI explanations between model correlations and biased outcomes. However, explainable AI is not equivalent to the type of transparency we need to redress harms caused by algorithms or identify positive lessons to inform the development of equitable algorithms. Achieving an unbiased and impartial algorithm is improbable because machine learning forces the system designer to choose a tolerable balance, based on her preferences or optimization goals.

A systemic solution that mitigates the harms of biased algorithms continues to escape the legions of AI researchers are aggressively exploring technical solutions to the challenge. Instead, Congress should focus on strengthening, maintaining, and growing the resiliency of the federal consumer oversight framework. Specifically, this task force should take action to strengthen and improve the model governance architecture. Recently, the Government Accountability Office (GAO) concluded that SR 11-7 a mission-critical model risk management framework is subject to review under the CRA.[17] While the effect of GAO's opinion is not immediately apparent, CFPB's precedent makes it clear that a critical regulatory gap will emerge and potentially weaken regulators' capacity to supervise financial institutions adequately. The task force should encourage CFPB to develop a parallel consumer-focused model governance framework, in light of the proliferation of algorithmic decision-making and marketing tools. Finally, the taskforce should vigilantly monitor the progress of HUD's proposed rule changes to amend the disparate impact standard.[18] The proposed rule

---

[17] GAO (2019, October 24) opinion on the "Board of Governors of the Federal Reserve System—Applicability of the Congressional Review Act to Supervision and Regulation Letter 11-7", Retrieved from https://www.gao.gov/assets/710/702190.pdf

[18] Housing and Urban Development Department Proposed Rule, "HUD's Implementation of the Fair Housing Act's Disparate Impact Standard" Aug. 19, 2019 https://www.federalregister.gov/documents/2019/08/19/2019-17542/huds-implementation-of-the-fair-housing-acts-disparate-impact-standard

introduced five new criteria for establishing disparate impact burdens that, in principle, serve to provide a safe harbor to institutions using algorithms to exploit vulnerable consumers and exacerbate historical disparities.