

Bari A. Williams – Proposed Testimony of use of Artificial Intelligence in Financial Services

To Chairwoman Maxine Waters, The Task Force on Artificial Intelligence of the House Financial Services Committee:

February 12, 2020

I am Bari A. Williams, an attorney and startup advisor, born and raised, and still live in Oakland, CA, working in technology transactions, with a focus on artificial intelligence (“AI”), privacy, and commercial contracts. My educational background includes a BA from UC Berkeley, an MBA from St. Mary’s College of California, and a Masters in African-American Studies from UCLA. In my career, I’ve worked for Facebook, Stubhub, and All Turtles, which is an AI startup studio, akin to an incubator.

In my work in the tech sector, I’ve been exposed to many interesting use cases for technology that provide convenience, efficiency, and optimization to our lives. But one nagging question always lingers – at whose expense are these gains made, and how do we solve for the negative impacts of some of the most pervasive uses of technology?

AI provides a unique example of this. To begin – what is AI? It is essentially someone’s bias, via datasets, baked into code that can determine the ads one sees, one’s credit worthiness, employment prospects, school admissions, housing opportunities, and criminal justice implications (i.e. facial recognition technology, gunshot locaters such as ShotSpotter, predictive policing such as Hunchlab, and predictive sentencing technology).

- (1) How are data sets, proprietary algorithms, and models are deployed and used within financial services, and what are ways to improve their deployment in financial services?

There are five main issues with AI, particularly in financial services: (1) what data sets are being used – who fact checks the fact checkers; (2) what hypotheses are set out to proven using this data – has the narrative that is being written been adequately vetted; (3) how inclusive is the team creating and testing the product – who are you building products with; (4) what conclusions are drawn from the pattern recognition and data that the AI provides – who are you building products for, and who may be harmed or receive benefit, and; (5) how do we ensure bias neutrality, and what is the benefit of neutrality.

Data sets in financial services are used to determine home ownership and mortgage, savings and student loan rates; the outcomes of credit card and loan applications; credit scores and credit worthiness, and insurance policy terms. It affects other outcomes, such as credit card fraud prediction. The danger in this is a repeat of redlining, the discriminatory practice of ensuring Black homeowners were confined to specific areas of a city and that their credit worthiness led to higher interest rates. The problem is that the data sets that are being used by these companies are “stale,” meaning they are dated and old. The older data has these remnants of credit worthiness during redlining, including income earned (which is already a huge disparity for people of color), and additional debt incurred.

The largest segment of AI use in financial services is anti-fraud. We see this not just in banks, but in any company that deals in consumer transactions, like StubHub, for instance. There are tech companies that make this software, such as Sift, to identify potential fraud risks. The problem is that if you are using a stale data set that skewed in favor of a certain demographic for “fraud potential,” it is already flawed, and the pattern recognition will be a self-fulfilling prophecy. Additionally, data is which is already coded based on someone’s personal bias, as data by itself doesn’t discern any conclusions. Data sets are often *chosen* to support a specific hypothesis, not to be neutral. This isn’t just a notion. In 2017, per [data from the Home Mortgage Disclosure Act](#) showed that 19% of Black borrowers and almost 14% of Latinx/Hispanic borrowers were turned down for a conventional loan.

There are several ways to improve the deployment of this technology in financial services. It starts with companies owning their power in implementing these technologies, and being deliberate about auditing the systems. Companies must proactively look for and identify bias in their AI by asking themselves these five questions:

1. Ensure all data groups have equal probability of being assigned to favorable outcomes.
2. Ensure all groups of a protected class have equal positive predictive value.
3. Ensure all groups of a protected class have predictive equality for false positive and false negative rates.
4. Maintain equalized odds ratio, opportunity ratio and treatment equality.
5. Minimize average odds difference and error rate difference.

Additionally, two techniques that can also drive fair outcomes include leveraging statistical techniques to resample or reweigh data to reduce bias, which is like a visual of giving someone a box to stand on if they are short to make them the same height and have the vantage point of someone privileged with more height. The second technique includes adding a “fairness regulator,” which is a mathematical constraint to ensure fairness in the model, to existing algorithms. The fairness regulator is akin to my fifth question about what does it mean to be bias natural. It is important to remember that not all biases are bad – some are actually beneficial and seek to right wrongs. There may be ways to award additional “points” or level the playing field with historically discriminated against marginalized groups to ensure parity with interest rates, financial advice, and credit assessment.

That said, nothing beats having diverse teams make and test these algorithms prior to use in the market. That will help ensure data sets are also diverse, and there is no disparate impact when testing. Lack of diversity in tech is an ethical issue, not just one about ‘doing the right thing.’

(2) What emerging methods or ways AI can be used to decrease discrimination and bias?

There are elements of AI that can be used for good. The concepts of transparency and fairness are not mutually exclusive, but to the contrary, are closely related. One solution uses mathematical methods in two separate products that provide explanations – the ability to identify what’s driving the disparity, and fairness – thus, provides the ability to reduce the disparity.

Some emerging methods that AI is used for productive outcomes includes AI that actually identifies bias via identifying pattern recognition with disparate impact. An example of this is seen with Zest, a tech financial services company, which has created a product, ZAML Fair, that reduces bias in credit assessment by ranking an algorithm's credit variables by how much they lead to biased outcomes, and then muffles the influence of those variables to produce a better model with less biased outcomes.

If more banks, and even consumer facing companies that use credit (i.e. retailers, credit reporting bureaus, etc.) would utilize a tool like this, we would see greater impartiality in financial decision making, which would save consumers billions of dollars, and may actually benefit companies by demonstrating efforts to be equitable, thus encouraging more business with consumers.

- (3) What ways existing laws and regulations can be applied to provide more transparency while still preserving data privacy and maintaining strong cybersecurity standards?

As I tell my clients, and my kids... a rule is only as good as its enforcement. To that end, it is imperative that the government use the laws already enacted to create greater parity and transparency.

The Fair Housing Act can be applied to ensure more fairness in the use of AI in financial services. For example, if a [mortgage lending model](#) finds that older individuals have a higher likelihood of defaulting on their loans then decides to reduce lending based on age, there is a legal claim for this to constitute illegal age discrimination and housing discrimination.

Additionally, greater enforcement of discrimination based on disparate impact on the basis of any protected class is illegal under the US Equal Credit Opportunity Act of 1974. The excuse that “the model did it, not a human” won’t work, when humans are coding the models. Per a [2018 study conducted at UC Berkeley](#) found that both traditional face-to-face loan decisions and those made by machine learning systems charged Latinx/Hispanic and Black borrowers interest rates that were 6-9 basis points higher. Lending discrimination costs these US minority borrowers \$250-\$500 million per year in extra mortgage interest. The study concluded that algorithms have not removed discrimination, but may have shifted the mode, and also made it more efficient.

Using greater enforcement of the laws on the books, calling for transparency into the data sets used to train these algorithms, as well as understanding the technique utilized by any human involvement in decision making after assessing an AI output could be very effective.

- (4) Are there any regulatory and legislative proposals to strengthen federal oversight of algorithmic decision-making and AI technologies utilized by financial institutions?

In addition to the suggestion of greater enforcement of the US Equal Credit Opportunity and The Fair Housing Act, I have submitted, along with this written testimony, an attachment as Exhibit

A, that is a proposed “AI Bill of Rights,” which would be guidelines that companies must meet in order to deploy this technology. No longer should it be “ship it fast,” which is a Silicon Valley ethos to get a product to market as quickly as possible, oftentimes to usurp a competitor, but instead the industry should adopt the medical ethos of “do no harm.” The premise of my recommendations is attached in Exhibit A, but of note, they include guidelines for requiring transparency and auditing of data sets being available to consumers and/or the govt. to discern equitable inclusion for unbiased results, and for all terms of service and information on data collected and its use be written in *plain English*, not legalese.

Additionally, the larger problem has been that technology is constantly iterating and improvising and improving, while law has no kept pace at the same rate. To that extent, the expertise needed to understand the tech production process is likely not going to happen in government, but perhaps creation of a hybrid model where there’s an institution that enables expertise to be cultivated, while also understanding the process of what it takes to turn a proposed bill into law. So, when companies are planning to implement AI systems, it's not just “ship it fast, and we’ll see what happens and fix it on the backend.” You actually have to verify the claims of your system, there is transparency around data sets used, and a company can answer five key questions: (1) what are you building and how, (2) what information are you using as the foundation of your system, (3) who are you building it for, (4) who are you building it with (diversity in tech and testing matters), and (5) is there any disparate impact when testing your system. Similar to the FDA model, it's acknowledged that not all drugs work for everybody. Not all technology is inclusive. One size does NOT fit all, so there are limitations.

The impact of technology has been a gift and a curse. The advent of this fourth industrial revolution has brought great convenience, but at a great expense – privacy, data collection and use, and less human interaction at the behest of automated decision making. In our quest to provide greater efficiency and convenience, we have been lax to look at who is left behind and how. If we aren’t careful, we will automate greater discrimination into the tools that we use everyday, and further exacerbate the legacy of lack those in marginalized communities. I implore the committee to do a deeper dive into how they can both enforce and enhance the US Equal Credit Opportunity and The Fair Housing Act, in addition to adopting the AI Bill of Rights as attached as Exhibit A.

Thank you for the opportunity to speak with the Committee.

Exhibit A – Proposed AI/Technology Bill of Rights

1. Make all terms of service and privacy policies in plain English (or whatever the applicable language may be). Even better if they can be written in a Q&A format.

The rationale for this rule is the ability to easily read and understand how information is collected, why, and how it will be used, and any ability for a consumer to delete or opt out of said collection.

2. Transparency and auditing of data sets should be made available to consumers and/or the govt. to discern equitable inclusion for unbiased results.

The rationale for this rule is to ensure that data sets are vetted for accuracy, any bias, or to make suggestions for other data sets that may complement or correct for errors in the data sets used to train the AI algorithms. It is important to do this to ensure there is fairness from the start.

3. No product should be embedded/integrated into products without testing for inclusion (i.e. differently abled, LGBTQ, employment/housing decisions, rural/city implications, economic ramifications) and no disparate impact. If a product shows a disparate impact upon a marginalized population within [TBD%] range, the product should not be sent to market.

The rationale for this rule is to ensure that products are not rushed to market without testing for adverse effects on one group of people of a protected characteristic more than another. It ensures parity of the ability of usage of the product, and that there will not be negative impact no already marginalized groups.

4. To ensure that AI products do not have a disparate impact on marginalized populations, companies shall do beta testing with people from marginalized groups, though also ensuring their privacy and protection v. limited data collection, anonymized data, and encryption at any and all points in the process where available.

The rationale for this rule is to ensure that testing is done with marginalized communities before a product is shipped. It is a complimentary rule to #2. Additionally, this supports more diversity in tech. If a company does not have sufficient employee population to engage in this testing, it encourages outsourcing to diverse suppliers who can assist with this beta testing with focus group organization, or doing it in in-house. This solves for both the company testing issue, and supports more diversity in tech, which produces an ethical and inclusive product.

5. AI For Good - The ability to "opt-in" to data collection, and sufficient notification (again, in plain English) before data is collected and used. Again, this information should be written in plain English.

The rationale for this rule is to ensure that consumers and individuals have the ability to control their own information. Currently, people are passively giving away what would be deemed “proprietary information” about themselves, and this affords the ability to control what is learned about a person and how it is applied. This is especially important regarding financial services when credit reports and scores are part of an employment applications and background checks.

6. If there are certain devices that an app needs access to, which wouldn't be intuitive (i.e. camera and gallery access for a food delivery app), a plain English explanation of why this access is necessary, and an example of how it will be used shall be provided. Example: "We need access to your camera/contacts to facilitate _____, and it will be used for [XX] duration in the following ways: _____."

The rationale for this rule is to give people all the information they need to make an informed decision before they decide to download and/or use an app. Oftentimes, people do not read the required permissions necessary to use an app, and instead just “scroll and accept.” Unfortunately, that means that an app may have access to features and attributes of your phone that it doesn't actually need to effectuate the service, but instead to just surveil you and collect information. By having these notifications written in plain English, it becomes clear what is needed, and what isn't, and a consumer can make a more informed decision about using that app, another, or none at all.

7. Optionality of Features - Provide the ability for consumers to opt-in to SOME features of an app, but not all. As the previous example notes, if we do not want to provide access to our cameras, contacts, or some requested access, though not necessary to deliver services consumer wants, then the app may still work, but will not provide full functionality possible had all permissions been given.

*At the moment, a consumer's choice is to provide all access requested, or not have use of the app. Depending on model of phone, alternative options, etc., this could demonstrate disparate impact.

The rationale for this rule is to ensure that consumers have choice when using an app, and can decide what permissions or access the app will have to their phone and its features. Currently, consumers do not have that option, and it means in order to use an app, one often has to consent to overbroad uses and access to features of a phone, such as the camera, contacts, and voice data. These access rights are often unnecessary to effectuate services, but are just a means of data collection and to surveil a consumer, often used to market new services, sold to another company or data broker, or to assist with building a new product. Customers shouldn't help create wealth and IP for a company without their explicit knowledge and consent, not just for convenience and lack of knowledge.

8. Data Portability and Right to be Forgotten - Consumers should have the right to retrieve, correct, or delete personal data controlled by any company that has access to such data (with correct permissions and rules around certain categories of data - i.e. medical records in emergency, criminal records that are not expunged, etc.). Along with this right,

much like GDPR, if a consumer decides to delete all of their records from an app or platform, that should ensure their information is wiped clean.

* This is especially true if algorithms are being created to look at credit worthiness, if there are insurance or medical decisions being made off of erroneous or dated historical data, etc.

The rationale for this rule is akin to the California Consumer Privacy Act, which affords a consumer the ability to audit and remove their data, or to opt-out of certain features. This would be very helpful when dealing with financial services, should excessive credit show up on a report which is requested for employment, or information that may bias decision making, such as rental history and locations. By giving a consumer power over their data collection and use, it restores trust in the companies that consumers decide to do business with, as they have the information needed to make an informed decision, but not enough to negatively impact decision making.

9. Affirmative Action for AI - After the 2008 housing crisis, 44% of Black Americans have a credit score below 650, after being steered to sub-prime loans. That said, some functions of AI that are currently used as using historical data that has bias baked into the code (i.e. housing data that incorporates redlining features, or predictive policing that includes historical crime data that disproportionately target Black and brown people, primarily in low-income areas as seen in Ferguson and San Francisco) should find a way to seek equity and parity in analysis of marginalized groups to not further negatively impact them.

The rationale for this rule is to correct for past wrongs, and to ensure greater equity. AI has the ability to build in extra scripts to award additional inputs and "points" to people of certain profiles for parity when looking at financial services, health disparities, or housing. This should be considered when making enterprise to consumer-based AI technology that could have far reaching, lasting effects on an entire community's wealth prospects.

The pace of law lags that of technology, as the latter drives innovation and the former waits to see the results before passing legislation or creating policies. If companies are forward-thinking in their application of predictive analytics, AI, and machine learning they can make these technologies inclusive without the need for new laws or regulations.