

BANK POLICY INSTITUTE

ANNUAL CONFERENCE

THE FUTURE OF CREDIT UNDERWRITING: ARTIFICIAL
INTELLIGENCE AND ITS ROLE IN CONSUMER CREDIT

Washington, D.C.

November 19-21, 2019

ANDERSON COURT REPORTING
500 Montgomery Street, Suite 400
Alexandria, VA 22314
Phone (703) 519-7180 Fax (703) 519-7190

PARTICIPANTS:

Moderator:

ERIC MOGILNICKI
Partner
Covington & Burling

Panelists:

ALBERT D. CHANG
Counsel, Office of Innovation, Consumer
Financial Protection Bureau

ANNIE DELGADO
Chief Compliance Officer
Upstart

CAROL EVANS
Associate Director
Division of Consumer and Community Affairs, Federal
Reserve Board

MEREDITH FUCHS
Senior Vice President
Chief Counsel of Regulatory Advisory
Capital One

* * * * *

P R O C E E D I N G S

MR. MOGILNICKI: Okay. Good morning everyone. I'm Eric Mogilnicki. I'm at Covington, Washington, D.C., where I lead the Consumer Financial Services Practice.

You are in the panel on *Artificial Intelligence and its Role in Consumer Credit*.

Artificial intelligence, as you know, is all around us. It's part of a whole range of bank functions, from customer service at the frontend, to A&L and cyber detection on the backend, and used properly artificial intelligence could transform credit underwriting, improving credit quality and credit access.

However, our current regulatory framework predates the use of artificial intelligence, and is at times an uncomfortable fit with innovations that move us past familiar benchmarks. So today's panel looks more at how financial institutions and regulators can work together to find a way forward.

The Bank Policy Institute has already been hard at work on this issue, one product to that is this

artificial intelligence discussion draft which is a joint BPI/Covington product, and so I highly recommend. There are copies on the back table, and it's available, like everything these days, online.

Before I introduce our panelists, there's one definitional issue. When people talk about artificial intelligence, are often talking about different things. So, when we say artificial intelligence today we'll mean all of the following three things. One is the use of computer systems to make intelligent decisions that traditionally required human intelligence. But also machine learning, which is the ways software algorithms learn over time, and optimize performance without further human intervention.

And third the use of big data and alternative data, those are the deeper and different sets than those that are traditionally used to calculate credit scores. And we'll distinguish among those facets of artificial intelligence as we go along, but when we say AI, we'll mean all three, unless otherwise noted. So, we've collected the very best in non-artificial intelligence to guide us through the thickets today.

I'll introduce our panelists, starting with Annie Delgado on my furthest right. Annie is the Chief Compliance Officer at Upstart, a leading AI lending platform designed to expand access to affordable credit. She's a leader in creating data-driven technology-based compliance programs to manage regulatory risks associated with the use of non-traditional credit models and variables. As you'll hear, her company, Upstart, is the only company that today has been able to work through the no-action letter process with the Consumer Financial Protection Bureau.

Speaking of which, our next panelist is Albert Chang, who serves as Counsel in the Office of Innovation at the CFPB, where he advises the Innovation Director, he's Co-Chair of the TechSprint Initiative, and implements programs that facilitate consumer-friendly innovation. Albert joined the Bureau in 2013 as a Brandeis Fellow in the Office of Fair Lending, but since that time he's been a Policy Assistant in the Office of Card and Payment Markets; Counsel in the Office of Supervision Policy, and Co-Chair of the CFPB's Alternative Data Working Group.

Beside Albert is Meredith Fuchs. Meredith Fuchs is Senior Vice President, Chief Counsel of Regulatory Advisory at Capital One. In that capacity Meredith leads the Company's horizontal regulatory legal function, supporting all lines of business. And Meredith joined Capital One after a distinguished career at the Consumer Financial Protection Bureau, serving various, serving various times as General Counsel, Acting Deputy Director, and Chief of Staff to Director Richard Cordray.

To my immediate right is Carol Evans. Carol is an Associate Director of the Division of Consumer and Community Affairs at the Federal Reserve Board. Carol provides strategic direction and legal expertise in several areas including fair lending, unfair or deceptive acts of practices enforcement, and consumer protection policy and outreach. Carol and her team are currently studying the implications of FinTechs for consumers. And if you're interested in learning more about the regulation of AI, as soon as you finish BPI whitepaper, I recommend Carol's article on *Keeping FinTech Fair*, thinking about fair lending and UDAP

risks.

So, with those introductions out of the way, let's start, if we could, by better understanding where the regulators are in regulating the use of artificial intelligence in consumer credit.

So, Carol and Albert, starting with you Carol, can you describe how your agencies are facing the issues presented by the use of AI?

MS. EVANS: Sure. Thanks, Eric. And first I will start with a disclaimer, and I think that Albert will have a similar one as well. The views I'm sharing today are my own, and don't necessarily represent the Federal Reserve Board, the Board of Governors, the Federal Reserve Banks, or anybody but me.

So with that out of way; at the Federal Reserve when we think about artificial intelligence, and machine learning, and big data, we think about it through the lens of safety and soundness, and also through the lens of consumer protection. And although I'm from the consumer protection side of the House, I coordinate very closely with my colleagues on the safety and soundness side, and we do the safety and soundness

and consumer protection as synergistic, we don't view them at odds at all. But really, what's good for one is good for the other.

And we think about the issues that we're talking about today, artificial intelligence and machine learning, and alternate data, and big data, we think that, combined, these two trends can create potent benefits for consumers and for banks. Used wisely they can result in better underwriting models, more accurate underwriting, lower cost at underwriting, and potentially increasing financial inclusion by better serving those who may be closed out of mainstream credit.

That said, those same trends that can create these benefits, the machine learning and big data, also if not implemented thoughtfully can undermine financial inclusion. And we'll talk about some of that later.

So, we view our role as the regulator, is how can we supervise institutions in a way that allows the benefits from these trends to flourish while still clear eye on the risks, and as we do that, we do really encourage discussion with all of you, and very much view

our interactions with a whole range of stakeholders, industry, FinTechs, and consumer advocates, as an open dialogue in a way that will help us perform our roles better.

MR. MOGILNICKI: Great. Albert?

MR. CHANG: So, as Carol has reminded me, I should also give a similar disclaimer. So, I work for the Bureau, I'm obviously not the Bureau, and so the views that I express here are certainly mine, and not necessarily those of the CFPB.

In terms of the CFPB's approach to these issues I think we've been sort of intentional about trying to understand both the risks and the benefits of the use of alternative data and machine learning in credit underwriting. I say benefits, in part because in 2015 our Office of Research found that 45 million Americans lack sufficient information in their Credit Bureau report, to produce a traditional credit score. So they don't have access to the same types of products that you or I might have access to as a result of having a traditional credit score.

And so fraction of that 45 million, certainly

not all 45 million, but some fraction of that 45 million could in fact be credit worthy, and could in fact be -- determined to be credit worthy if nontraditional data is actually used to evaluate their credit worthiness.

And so that's part of the reason why in 2017 we issued a request for information looking at the benefits as well as the risks of alternative data and machine learning, that's also, I think part of the reason why staff in 2017 issued the new action letter to Upstart, so that we would have a front row seat to some of these innovations related to alternative data and machine learning.

We also hope, through our new compliance assistance sandbox, which gives a legal Safe Harbor in exchange for testing these types of innovations, and providing data as well certain consumer safeguards during the determined testing.

We are hopeful that the compliance assistance sandbox, will give us this front row seat again to the use of AI in credit underwritings, so that we understand what the benefits are, what the risks are, and whether they actually materialize in real-life testing. And I

think all of that information will give us a better sense of how to encourage the responsible use of these technologies.

MR. MOGILNICKI: So, it needs to be used wisely, thoughtfully, responsibly. Annie and Meredith, do you want to tell us how you're meeting those marks?

MS. DELGADO: Sure. So, first by way of a quick introduction to Upstart, for anyone who is not familiar, Upstart is an online lending platform that partners with banks to offer a fully-digitized consumer loan product online. And we have banks of various sorts of sizes and geographies on our platform today, and those banks have collectively originated close to 400,000 loans through the software.

And a core part of the offering that we provide to banks is the access to our underwriting and pricing engine, which does leverage machine learning and alternative datasets to more accurately and fairly, in our view, price consumer credit. So, sort of to underscore what Carol had said, I think it's important to talk about the risks associated with these types of technologies, and we're certainly going to have a lot of

conversation about that today.

But at the outset I just want to underscore that a core part of the conversation is also to really look at the benefits. And in the work that we've been doing since Upstart started, you know, seven years ago, we've really observed a lot of increased assets to credit, a lot more fair pricing, and a lot more improvements on customer experience by leverage these types of technologies. And it's something that we hope to sort of continue to see grow, and to continue to see adopted throughout the industry.

MS. FUCHS: Sure. Thank you. I want to thank TCH and BPI for hosting this conversation, and happy to talk a little bit about how we are beginning to use artificial intelligence and machine learning. But at the outset I want to sort of recognize that the way people talk about this and think about this kind of varies between a vision, or some sort of utopia, in which there's perfect and complete information, and making sure we make perfect decisions, unbiased decisions, and the sort of the SkyNet or help in 2001 where machines are taking over, and they are taking the

place of humans, and making decisions.

As all things I think the reality is something in between. The machines, and the data, and the alternative data, help us get closer to making good decisions, and better decisions, so that the right credit goes to the right people. But we do still need to have human intervention, we can just let go of it fall, and so with the right approach I think we can actually deploy this for good, instead of the destruction of humanity.

And I think that's very much the way, you know, we are thinking about it. Certain aspects, certain things are, you know, automatable, and when they're automatable we are certainly already taking advantage of it.

But you asked specifically about underwriting. I think that our willingness to kind of let go is sort of an inverse proportion to the materiality of thing that's being decided. And when it comes to something like underwriting that's certainly something which is important and where there is risk attached.

And so, the way we look at it is, you don't

sort of unleash the machine to make a decision, humans are always involved in the decision-making, the models and the machines allow humans to set criteria, gather more data than they could before, identify correlations that could have been -- they couldn't identify before. But we still set the outcome criteria.

So, as we think about it, you know, things like advertising, marketing are more automatable. There are things like back office operations that are more automatable. But, you know, as we get to something like underwriting, we are not just letting -- you know, unleashing the machine without human guidance and consideration.

MR. MOGILNICKI: Okay. Thanks. I think it is true that banks generally are experiencing AI's development has been lagging in credit underwriting, as opposed to, say, customer service and other areas where they feel more comfortable pursuing a technologically advanced pathway.

I'd like to turn to one of the -- there are two issues really that are front and center on AI and credit underwriting, one is fair lending, and one is

explainability. I'd like to turn to each of those today, starting with fair lending. And fair lending is a real concern because folks are -- there are genuine concerns that in an effort to increase fairness in credit underwriting, what AI actually does is create the potential for disparate impact, through the actual -- through proxy discrimination.

The machine will find a way to figure out who is who in ways that we're not comfortable with, and assess whether or not people were offered credit, and with results that we would not count (inaudible) if it was a person making the decision. But to start the discussion it's only fair to start with some of the advantages for fairer lending through artificial intelligence.

And I know, Annie, you've had actual results, thanks to the Bureau's no-action letter. So, would you start us by talking about the potential benefits for fairness to consumers from the use of artificial intelligence?

MS. DELGADO: Yes, absolutely. So, when Upstart was first sort of getting off the ground, we

wanted to really understand the problem statement of what we were trying to tackle, and so we did a study with a national credit reporting agency to try to understand what is the delta between prime customers and people paying back their loans.

And what we learned is that there about 45 percent of Americans who have access to what we would call prime credit, and there's 83 percent of Americans who have never defaulted on a loan. And so that is a big problem statement, there's a lot of people that are being left out of the system today, or are being overpriced by the system. And so that is sort of the core mission of what we want to try to be solving in our day-to-day work at Upstart.

And our view is that the way to solve this is by leveraging additional data sets, and using those data sets in nontraditional ways as specifically through machine learning and artificial intelligence, as opposed to a traditional regression model, in order to bring in more information about prospective applicants, and then evaluate the application in a more holistic fashion, and a non-binary fashion.

So, it's important, and I think we touched on earlier, there's sort of this issue of machine learning, and there's a separate issue of datasets, but I argue, is that those things really should be combined together because when you're looking at issues of fairness there are certainly categories of data that can be used -- or that can be unfair if used in the wrong way.

But if they're using a statistically rigorous manner, such as the sort of more sophisticated modeling techniques, they can observe nuances and correlations that actually bring more people into the game relative to what exists today.

And so that's kind of the core of what we do at Upstart. And I think a big part of this that's important to mention is that these systems can really reduce the incidents of bias and lending generally because if you think about a traditional sort of ecosystem where the decision -- loan decisions are sort of spread out all throughout the country in different underwriter's hands, there's thousands and thousands of different people making these decisions every day.

Whereas, if you're talking a machine learning

model and you have a more limited group of people that can control the inputs and monitor the outputs of that system, you can have a much more sort of condensed framework to understand what's going on. And of course to control the fairness of the system, you teach a machine to not have bias, it's much harder to change society.

And so that's sort of the belief Upstart, and we've had very good results so far. Earlier this year the bureau published a blog with an update to our no-action letter, as part of our no-action letter process we an annual access to credit study to observe how the Upstart model is performing relative to our traditional model, in terms of improving access to credit.

And this year, our results have shown that our model is approving 27 percent more applicants relative to a traditional model, and those applicants are getting priced at about 3.5 percentage points lower in terms of their interest rate. So we're really proud of those results, and it's something of course that we are going to continue to monitor year over year, but very promising results so far.

MR. MOGILNICKI: That's perfect, Annie.

Thanks for sharing. I know on the flip side, Carol, you've done a lot of thinking about how we can make sure this is fair. And I wonder if you would share some of the pitfalls that you and/or the regulator see with this use of AI in private underwriting.

MS. EVANS: Sure. And first I do want to say that I concur with Annie, that there's a lot of benefits out there. And when I talk about the risk, it's not to somehow undercut the validity of those benefits, but rather is to ensure that we're all working together to make sure that FinTech delivers on financial inclusion. And I think the best way to do that, is have a very clear view of the kinds of pitfalls and the risks are, because that same combination of new data and potent analytical techniques if not used well, could risk automating discrimination or unfairness.

And on the risk side o the ledger, there's two issues that I want to draw your attention to. One is, I'm thinking a little bit about alternate data, and I really encourage everybody to unpack that term alternative data, and not treat it as a monolithic

category.

Some kinds of data that are being discussed and employed, bear obvious relationship credit worthiness, and underwriting, and may be highly aligned with traditional underwriting principles. And other kinds, on the other hand, may have no obvious relationship to credit worthiness, and may potentially be correlated with race.

So, for example, on the type of data that are often unrelated to credit worthiness, we know that some lenders are exploring this in cash flow data, and there's been some recent research by FinRegLab about a study of some non-FinTech lenders that were using it, but overall saw some positive results that these kinds of data were predictive, and that they appeared to be consistent with fair-lending principles.

And then on the -- kind of the other side of the spectrum, also stories about some lenders exploring who your Facebook friends are, and your credit worthiness, to determine your credit worthiness, and some of these kinds of data do raise concerns that they could further calcify financial inequities, and that

they may undermine fairness.

So I think the unpacking alternate data is really helpful and useful. If you think about this, I also kind of urge us all to think critically about the phrase "data driven" everybody wants to be data driven these days. And, you know, when I started first doing fair-lending work, that was also what automated underwriting systems and credit scoring was all about, is being data driven.

And instead of the risk of human underwriters taking their own bias into consideration when making decisions, our models were supposed to use more objective data, such as, you know, how a consumer had actually his or her bills, and really make objective decisions. And I think, overall, those trends really have increased financial inclusion.

And sometimes when we're talking about this data like, what your hobbies are, or how you browse on the Internet, you can call models based on that as data-driven, but it's not data driven in a way in which you usually thought about it, and I think we all want to be cautious of not using data, that we've actually tried to

guard against in the underwriting process, and suddenly it's sort of automating with potential biases. So, that's something I think we all want to be sensitive to.

And I don't want to suggest that some data should be clearly off limits, that's not prohibited by law, but data that may reflect past discrimination, or may be correlated with race, national origin, or other protected characteristics really warrants additional (inaudible) and scrutiny, and we think that's really important.

The second issue I just want to call to your attention, is what it means to be on the Internet these days, I'm sure we all remember seeing that iconic New Yorker cartoon, that said, "Nobody knows that you're a dog on the Internet." Well, now nobody is anonymous on the Internet by and large, and there's not really a single version of the Internet.

With increasing frequencies, when we all are on the Internet, what we observe maybe curated for us, based on the information websites have about us. And there can be some really great benefits to this, and we're more likely to be showing information that we care

about, and information that we don't care about.

But on the other hand, there are some risks, because if information can be curated for us, and we may be shown different things, there is a risk that this could result in minority communities being shown different information potentially resulting in steering or other redlining concerns.

And many of you have probably seen further Facebook litigation, and there're two facets of that litigation. One is concerns that advertising users using Facebook could target audiences in very granular ways that could be proxies for protected characteristics such race and national origin. And Facebook settled last spring with a number of civil rights organizations, including ACLU and International Fair Housing Alliance, to try to impose some safeguards to limit these targeted advertising based on (inaudible) characteristics.

But there is also a high charge against Facebook that sort of also discussed this target, advertising, but raised a new concern, which is even if an advertiser intends to reach a broad audience that high charge alleged that it's possible that the Facebook

algorithms that determine what ads each person sees, could still result in narrowing that audience who sees an ad based on protected characteristics.

And even though this litigation is about Facebook, I think some of these concerns also may be worth consideration when we talk about platforms.

And so it's a highly evolving issue, but I think it illustrates that we all want to be sensitive to how information can be curated and targeted, and make sure that it's done in a way that really enhances financial inclusion, and that doesn't undermine it.

MR. MOGILNICKI: Meredith, how did these, the pros and cons of using AI sort of resolve themselves when and institution has to make decision when to use AI?

MS. FUCHS: Well, I thought, you know, Carol identified a number of really important issues that we've been thinking about. So, one is there's certainly bias risk using alternative data, and machine learning, because we know the data is not always clean or complete, et cetera. And so those are things that we are thinking about and working on.

But I also think we are fooling ourselves if we think the existing system is completely without bias. I mean, yes, there are certain things that we rely on that we think are traditionally related to the likelihood of our ability to repay, but there's things that are black boxes in the existing system, like FICO score or something like that.

So, in a sense I think we should not forget that AI potentially will help us see bias more, and improve and fix it to get it out of the system. So, in that sense I think it's very exciting.

I also think another thing that Carol highlighted today that I think is a very severe problem, particularly for institutions and entities that don't have, you know, large compliance staff, is complacency where you let the machine just run. And I think whenever there are things that require quick decisions under pressure, there's going to be a desire to just allow the machine to run, because it's been set up.

And that's a thing that I think regulators do have to absolutely keep an eye on. A couple of areas though that are challenging for us as a business that it

would be helpful over time as the regulators develop views on, is understanding sort of like, what is a proxy that, you know, for a protective class that we have to be worried about.

Because the machines are going to create correlations that are different than the ones that we use today, and it may not be as straightforward. And so understanding is this something that is risky, or is this something that is okay. Or, if the considerations are balanced in a certain way, it's acceptable, but if they're balanced in a different it's not going to be acceptable. That's the kind of thing that would help us take advantage of the opportunity of AI.

I also think, you know, you mentioned the Facebook litigation, I think that another area that is challenging is using third parties in general, and how do we deal with them in a way that we're not taking on liability for things that we may not be able to fix.

I mean, I will say we are by our nature conservative as a bank, and so we are very, very careful in trying to use, you know, like social media for advertising and marketing, but at some point there may

be some things that we're not able to control, and we don't want to be closed out of that as that marketplace for us to reach the public.

So, how can you help us understand what's acceptable and what's not acceptable there?

MR. MOGILNICKI: Well, so if only there was some way for the CFPB to describe its decision not to take action against the particular practice. Upstart famously received the one and only, I believe, no-action letter from the CFPB relating to their underwriting program, and to the application of the ECOA and Reg B to it. So, appropriately enough, Albert, and Annie would you work together to describe how that process came about and what the letter says?

MS. DELGADO: Sure. So for a sort of initial context, Upstart had the alternative viewpoint that we should from day one try to be very proactive and engaged with regulators, which I think is an uncommon approach for FinTechs, our Regulatory Council was employee for. I heard someone joking yesterday that FinTechs often hire regulatory council's employee (inaudible), but I'm really grateful that Upstart did not do that.

So we engaged with CFPB very early on before we ever even started -- launched the platform, just through regular office hours to try to bring to them, hey, this is what we're trying to do. We are not trying to fly under the radar, we are -- you know, we believe that this technology can help and here is how. And so we started having those conversations very early on in our tenure, and so when the no-action letter process was announced it was sort of a natural evolution of things that we were already thinking about and discussing with the Bureau.

And I think for us, there was a clear business need to sort of submit an application, and that business need laid in the uncertainty that some of our investor market had about how regulators might view what we're doing. And so we always said, you know, we have this rigorous compliance plan, we are doing this testing, and the reaction from investors of course, well, what are your regulators thinking.

So we decided, well, let's go and submit this application, and find out what our regulators think about this. And so, that was sort of the genesis behind

why we started that conversation, and the specific area of regulatory uncertainty that I want to make sure is clear is the area where we felt we needed some more certainty, is some information about how to measure compliance with ECOA, so not to go to the regulator and say, can you sign off on this variable or this modeling technique or that type of thing.

But more practically, how do we demonstrate compliance with ECOA, this is what we're doing, let's work through -- let's work through that part of the puzzle. Because I think that's where still a little bit of uncertainty in the Reg itself.

MR. CHANG: And just to speak somewhat more generally, whether it was with respect to the Upstart no-action letter application or, you know, future applications under a new sandbox policy. There are distinct aspects of that application process, so I think the most -- I think one of the most critical aspects is to be very clear about the product or service in question, and how it might interface or impact consumers.

And I think a second aspect is to be

transparent about the potential benefits and risks of that product or service. Third, you know, if you're applying for a program that involves providing some kind of testing data to the CFPB, it's important to think carefully about the methodology that would be used to evaluate whether or not those asserted benefits and risks are actually -- are realized during the term of the testing.

And then last but not least and, you know, Annie touched on this, it's really critical to think carefully about the regulatory uncertainty that you link the sandbox approval, or the no-action letter to address. And to the extent that that regulatory uncertainty can be framed in a -- I think narrowly-tailored manner, that just makes it, I think, so much easier for an agency like CFPB to process that application.

And I think Annie sort of pointed out how that regulatory uncertainty was framed in the Upstart request. That Upstart is public, and you'll see that the framing was around how demonstrate compliance, around the Equal Credit Opportunity Act, with respect to

the use of alternative data. Obviously there are other ways in which you could frame that uncertainty in a broader way, but I think that would -- that would have made a fairly lengthy process even longer.

MR. MOGILNICKI: So, Albert, the no-action letter process has been streamlined in the last year or so, so it's still easier now to get a no-action letter processed. Have you seen the uptick in applications? Or do you want to advertise for one right here?

MR. CHANG: It did take some time to process that first application, we'd like to be more efficient in terms of processing these applications. So one change to the no-action letter policy, and one aspect to the compliance assistance sandbox, is that once we receive a complete application, and we're certainly hopeful that we'll have conversations with potential applicants before we receive a complete application. But once we receive the complete application, we'd like to make a decision on that application within 60 days.

So, we'd like to move this process along more quickly. There are other aspects of the program that we've added to try to increase some demand in these

programs. So, with respect to the sandbox like I said earlier, there's a legal Safe Harbor involved. And then with respect to the no-action letter, instead of the no-action statement being a staff recommendation, it is a statement from the Bureau with respect to no-action letter.

So we're hopeful that those new aspects will be stoke demand, and again, allow us to have a front row seat to some of these innovations.

MR. MOGILNICKI: So with that, let's turn to the other topic I mentioned, which is explainability. You know, one of the advantages of artificial intelligence in credit underwriting is they can make subtle and complicated decisions about who should get credit.

One of the problems with artificial intelligence in vetted underwriting is that it makes subtle and complicated decisions about who should get credit. At times, as Carol indicated, it may use logic that humans can't follow, right. There may be correlation without causation; or causation that would never occur to a human reviewing the results of an AI

system.

Federal law, on the other hand, include ECOA, Reg B and the Federal -- and the FCRA, all require that explanations be provided to consumers about why adverse action was taken on their loan applications. So, there's a disconnect in some ways between the old system providing clear explanations and the use of AI where the explanations may be less clear, less capable of explanation in a few words.

So, let me start with you, Meredith, and ask, sort of what's that bank's perspective, or your bank's perspective, on the intersection of using AI and having this adverse action notice regime that's left over I think, where it did not anticipate the role of AI in credit underwriting?

MS. FUCHS: Well, I mean, if you start with what's the purpose of adverse action notices, I'm not sure I agree with the premise that, you know, it's sort of leftover from the past. It may need to be adapted to the future. But I think, you know, banks live in a world where we regularly have complex algorithm that we use that are not easily understandable to the public.

You know, so when we think about certainly credit underwriting even today, it may not be so clear for folks. And we've come up with a way to explain it to meet the goals of the adverse action notices both, you know, for the purposes of being able to identify and make clear that we're not discriminating, but also people can understand what they can do to improve their ability to get credit in the future.

So, I don't think it's as novel at the moment as it could be, and we're operating within the current system, and Albert can certainly speak to the flexibility that the current system allows.

Again, one area that I think would be helpful, you know, from regulators, is as people start to use some of the flexibility for adverse action notices not come up with certain new ways to describe things, to demonstrate that that's acceptable. And maybe that's a good use for the -- you know, no-action letter process.

So I don't think it's necessarily been something that is holding us back, I don't think it's very possible that if we get to a point where we are really allowing much more complex decisions to be made

by computers without the human intervention along the way, then it could get challenging, because think people have a hard understanding -- I mean Carol spoke to it earlier -- things that are not intuitive.

So we are going to need to come up with new ways of explaining that, and again, the regulators I think will be very, very critical in helping us do that.

MR. MOGILNICKI: So, you know, I agree that there temporary solutions to this problem, but I think probably not a permanent solution given how fast AI is developing, how much data is being developing. In the time we've been sitting here, there have been 100 million Google searches, right, all of which are part of their permanent history, especially the folks in the back who are doing them as we speak. (Laughter)

And I know you've probably all read about quantum computing, and how computing is taking a vast leap forward in the ability of the computers to sift through data rapidly. So, I am a little concerned that we will solve today's problems, but not be looking out five, ten years from now when the black box is deeper and darker that it is now.

So one hard question I think, Carol, and so I'll pose it to you is: what do we do? What gives when it turns out that we can have better predictions or better explanations but not (inaudible)?

MS. EVANS: Well, I look forward to that world. I don't think we're there yet. You know, I -- so as with Meredith I agree that these aren't necessarily new issues, those who have been engaged in having to create notices for FICO decisions, and other custom underwriting decisions, have been thinking through these issues for quite some time.

And I think that we can learn from the evolution and credit scoring automate underwriting, how we can ensure transparency, and still leverage all the benefits of the technology. I think that the principle of transparency is an important one. I think the principle of fairness and nondiscrimination are also critically important, so I'm very hesitant to kind of jump on the bandwagon that says, gosh, you guys have certainly a cool new toy, and you don't think I could hit my principle, so the principles have to give.

That strikes me as bit of a dangerous place to

be, and I don't see evidence that we're there yet. And I also just want to -- you know (inaudible) question, a note which is the whole rhyme and reasoning of, we're smart people, I know you understand the models, but trust us, they don't think that's worked that well for financial services in the past, so I don't really encourage jumping into that abyss today.

And I also, I don't see the need, I don't -- I haven't seen any reason why we can't adhere to these principles that are so important including transparency. And Meredith had some really interesting remarks about, we may be able to leverage that, and technology helped to get us there, and there's a lot of really great thinking going on, about how to make AI more explainable.

And I think that had increased that transparency, so I would urge our energy to be spent there as opposed to thinking about ways why we should -- arguments to dispense with the principles of transparency.

MR. MOGILNICKI: Okay. So, let's direct that energy to the CFPB, now the CFPB interprets Reg B, they

have sample notification forms. What can you, what you do to make adverse action a little more AI-friendly?

MR. CHANG: Let me just say, I'm really glad that we're talking about adverse action notices. I think for far too long, those requirements were seen as secondary to the anti-discrimination provisions of the Equal Credit Opportunity Act. And I think Eric, as you had pointed out, they're really core to that anti-discrimination purpose.

In fact just, you know, thinking about the news from last week in terms of allegations of algorithmic discrimination, if you're facing those allegations, I think the first step is to understand why the algorithm or the decision-making process led to two different outcomes for apparently similarly situated applicants. And if you don't have that ability to explain that decision, then it's awfully hard to rebut that allegation of discrimination.

But in terms of the adverse action notice provisions themselves, and AI, I think Meredith pointed out that there's a lot of built-in flexibility in the regulation, it's not terribly prescriptive, its rules

around adverse action notices, there isn't a single methodology, to identify principled reasons for denial, when you're using credit scoring model.

There is no requirement to use that sample notification form, or the sampled reasons that are in the commentary to Regulation B, and in fact, the legislative history to ECOA include a statement I think from the Senate that the Notice provision is intended to work, you know, in a sensible and flexible way. And so I think there is room for innovation with respect to AI, that's compliant with adverse action notice requirements.

Whether there is incentive on the part of financial institutions to innovate in terms of disclosures, I think that's not clear. I think you all probably have a better sense of that than I do, but if there is an incentive that the CFPB can provide.

You know, we recently revised our trial disclosure sandbox program, and that provides a legal Safe Harbor to test a trial disclosure that is intended to improve upon existing disclosures. And that policy is very explicit that we consider adverse action notices to

be one type of disclosure that can be tested under that policy.

MR. MOGILNICKI: Thank you, Albert. I would suggest that on behalf financial institutions that the level of comfort with providing adverse action notices that aren't listed in the sample notification forms is really quite low. And so it's true that in theory, financial institutions have the flexibility to go above and beyond, or to do something different. In practice we won't, so we ask the Bureau to help us by expanding on this. Go ahead.

MS. EVANS: One interesting thing Albert and I were talking about right before the panel, was that the trade associations also have the ability to request -- is that under the trial disclosure program or just the -

MR. CHANG: All three of them.

MS. EVANS: All three of them?

MR. CHANG: Yeah.

MS. EVANS: So there may be an opportunity to think about, you know, working with the Bureau to try to address them or not.

MR. MOGILNICKI: Annie, could I ask you about some best practices in this area?

MS. DELGADO: Yeah, absolutely. So, I think when you're talking explainability, or when you're talking about fair lending, or a myriad of other topics, but the key need is for a rigorous compliance program with testing and controls. Obviously as a compliance person I'm going to say that about most things, but it's been my experience that having a sophisticated framework where there is ethic about ongoing monitoring on the inputs and the outputs of your system, you can mitigate a lot of these risks.

So, in terms of, you know, best practices that we have Upstart we have a regular cadence for testing and observing these things. We have predefined thresholds for how we think about them, and we have -- we bring a number of different sort of areas of expertise to the table when we are looking at these results.

People from data science, people from credit, people compliance and legal. So that we can really have a comprehensive -- a comprehensive assessment of what's

going in and what's coming out of the model, and not a place as they said, both two issues of fair lending, but also to explainability, and how candidates' denial reasons are selected and used.

And then of course also, you know, observing the outputs, doing traditional testing, right. They tell people all the time we do very fancy, you know, data science and modeling with our credit modeling, we have a very traditional compliance program. And I think that that's really important because at the end of the day you need humans looking at the output of what's going on, and making assessments about whether they're compliant.

MR. MOGILNICKI: Great. We'll ask for questions in a moment, but some have arrived through the computer, they may have been artificially generated, I'm not sure. (Laughter) But let me pose a few of the ten to the panelists before we wind up. Here's one: To the extent AI can actually broaden the access of borrowers to prudently underwritten credit, what role can it or should it play in CRA compliance?

Let me broaden the question then. AI is a

tool that we've taking about as using credit underwriting, but isn't it also an effective tool to use for things finding AML risk, or compliance risk generally, or extinguishing CRA risks, as the question asks?

MS. EVANS: Yeah. I think there are so many used cases for AI, if you just look at -- you can walk into a financial institution or a bank and look at what's going on, and you can probably apply AI in so many different areas, and there's probably, you know, as time progresses even more areas that are going to start to leverage these things.

So, I think definitely fraud detection, AML, CRA, fair lending, servicing and collections activities. I mean there's just -- customer experience, marketing, there's so many different aspects where AI can bring benefit for consumers.

MR. MOGILNICKI: Okay. And it may take AI to police AI at the end of the day.

MS. EVANS: And Eric can I add?

MR. MOGILNICKI: Please.

MS. EVANS: Community investment part, I can

just say that we think there's also synergy between CRA and consumer protection, and safety and soundness, so all the trends that we're discussing up here today, we think will inherently benefit the CRA, and helping serve their communities, which is really what the Community Investment Act is all about.

MR. MOGILNICKI: Okay. Another one from the audience: How much of the eventual regulatory conversation will be ML -- machine learning I assume -- model choice by the firm? Sometimes the machine learning model choice will produce a result that is accurate from a predictability perspective but not transparent, and vice versa?

And I think I'll broaden that question, you know, there is model -- the sort of model risk management guidance that may or may not be suitable in light of the advent of AI. So, anyone have thoughts about model choice and how that plays into the risks and benefits of AI?

MS. FUCHS: I guess I would say, I mean I think that, you know, our traditional framework for thinking about models is sort of two versions, you know,

something that is not very material, and doesn't require a lot of oversight. And then the other side of that, which is, it's material and it requires a lot of oversight.

I think that one of the things we have observed is that, you know, there is more of a spectrum of decision-making, and more of a spectrum of uses for artificial intelligence, and so when we think the kind of oversight we're putting into it, it may -- and that's like an important regulatory discussion we need to have.

It may make sense to have, you know, some changes to the way we approach model risk oversight. You know, so for example, like we talked about, you know, Annie gave like a whole list of all the possible benefits of using machine learning and artificial intelligence, and I agree, it was a great list, because there's no question we can do things faster for our customers, and our customers are going to get better experiences in a number of ways.

And those things, once you figure out how to do it, they are not very risky. And there is the other side, we talked a lot about underwriting today. But

there's a lot of stuff that's in between there where we can create benefits, and it's not of the level that is like underwriting, where we are very, very concerned about making sure we have explainability, and that we can have the adverse action notices populated with something that's meaningful. So I feel like there's a range of things there that we might -- it would be good to have a regulatory discussion about.

MR. MOGILNICKI: Yeah. One of the points the whitepaper -- did I mention the white paper -- (Laughter) -- it's available in the back. It makes it that the model risk management guidance that we have now really predates AI, and it looks at models, the way we've traditionally looked at models, which is a machine.

And you look at how the machine is constructed, and you look at the inputs and the outputs and you expect, like a good machine, for the inputs to yield the same outputs over time. That doesn't work for AI. AI is not a machine, it's an organism that changes over time, that evolves, that gets better over time, and so you can't expect to kick the tires in January, and

expect the same results in June. And that's a real problem for -- particularly line-level examiners who may not be particularly familiar with AI, in reviewing credit decisions should be made by those institution that are using AI in credit underwriting.

I'll keep looking for questions here, but I wanted to, before our time is up, ask a closing question for all four of you, for Meredith and Annie. What do you wish regulators better understood about how you're using and why you're using, or why you might artificial intelligence to add to underwriting.

MS. DELGADO: So, I actually think regulators do understand, you know, my experience in speaking with various entities is that there is a lot of interest in these technologies, there is a lot of interest in innovation, generally. And so I think that -- you know, the regulators have done an amazing job of being sort of innovation-friendly, and trying to forward policies like the ones that Albert discussed, and that promote innovation.

I think that what regulators need to understand is the challenge that banks in actually

adopting the technologies, and then understanding how existing guidance may or may not apply. So, you know, there is model guidance, there is third-party risk management guidance, and that framework is very solid, but when a bank is actually wrestling with the decision of: How am I going to use this technology? Should I partner with the FinTech? Should I build my own thing?

Really what they're thinking is like: What is my examiner going to think about this? And how is my examiner going to interpret this guidance? And so I think that that's where it would be useful to get a little bit more clarity from the regulators, so that the examiners have just a more consistent playbook when it comes to some of these issues.

MR. MOGILNICKI: Thanks.

MS. FUCHS: You know, I would agree with that. I think Annie makes a good about, you know, the agencies are clearly signaling in many ways including, you know, your participation in this panel that they want to enable banks to innovate. And there is no question that, you know, we have to innovate, that if we fail to innovate we are not going to be able to serve our

customers, and the whole industry is going to be impacted negatively.

So, that's a really important focus. I think that making sure that the communication with the on-the-ground supervisory team is really good, so that they feel comfortable assessing what we're going, and assessing its safety and soundness along the way, is really going to be an important thing to help us be successful in doing some of these new things with AI.

MR. MOGILNICKI: I would just add, from my perspective, and this relates to Eugene Ludwig said at one of the earlier sessions, which is: There is a real unlevel playing field right now between non-banks and banks when it comes to the use of AI in credit underwriting. It is not good for consumers or for credit, if there is one set of standards for non-banks, and another set of standards for the bank down the street.

In fact it may, for banks create a real spiraling problem, which is that if you accept the premise that AI makes you more efficient about deciding who gets credit and how it's priced, if non-banks were

allowed to use tools that banks either aren't allowed to use, or hesitant to use because of the supervisory apparatus that is over their shoulder, then you're going to have a real problem that consumers are protected either -- both too much and too little in the use of AI and credit underwriting.

So, I would urge, and particularly that the Bureau, which I think because it regulates banks and non-banks have a special role to play here, to focus on the need to make sure that the same standards apply across the board. That's the only way consumer is going to have same level of protection, and regardless of what website they happen upon when they're looking for credit.

MS. FUCHS: Can I just -- But the reality is it's a challenge for the Bureau to do that as well. I mean for the banks we have supervisory staffs in our institutions all the time, and the regular exams, and choose things like, how do you oversee a model is something that is a regular discussion with our regulators. If you're not a bank then you just don't have that level of oversight.

So it's kind of at the other end, when the outcomes are bad, is really the only time that people are going to be looking at, you know, what we're doing.

MR. MOGILNICKI: So, let me ask the same question in a different to Albert and Carol; which is what -- this is your last few minutes to tell this distinguished audience exactly what you wish they knew about how the regulation is thinking about AI, and how they might proceed safely in this space? Albert, do you want to start?

MR. CHANG: Sure. I think with respect to risks it's important to quantify those risks, and to think carefully about how to manage those risks, at the same time I think it's also important not to be paralyzed by those issues. I was on D.C. FinTech Week Panel, last month, and one of the audience's questions was related to AI ethics, and I think the person asking the question was just presenting the self-driving car scenario, where cars come hurtling down the road, and has to make a split second decision as to whether to sort of plough into the pedestrian, or veer into oncoming traffic and sacrifice the car, and perhaps

driver.

I mean, I think in the credit underwriting context, you don't need to make that, you know, split-second decision, right. If your AI estimate is very different from you traditional system's estimate, there's an opportunity there to, you know, kick that decision out for some kind of manual review. There's certainly a sacrifice on speed but I think consumers would appreciate getting that decision right, rather than emphasizing speed in that situation.

MR. MOGILNICKI: So, Carol, as always, the Fed has the last word. (Laughter)

MS. EVANS: Thanks, Eric. I'll answer two parting pieces of my own advice, which is first of all, make sure your consumer compliance, or your compliance folks are working with your business staff early and often. As I think Annie did a fabulous job of illustrating what kind of (inaudible) and thoughtful issues can be raised by a compliance officer.

And the second I just want to say, is diversity and inclusion matters in this discussion. About a year or so ago, you all might recall an article

about Amazon that they were developing -- or hiring algorithm for tech folks, that was trained on a pool of résumés that appealed to mostly male. And as a result the algorithm learned to discriminate against women. And I should say, this was not -- (inaudible) was not deployed, Amazon didn't go forward with it.

But that said, when I hear stories like that it makes me wonder, hmm, you know, the fact that tech workers are mostly male, that wasn't a hidden secret, it was a secret of intellectual property. Who was in the room? Who was at the table when this model was discussed? And who wasn't?

So, to both reap the benefits of this technology and all this innovation, and also to be aware of the risk, I really encourage you to think hard about risk and inclusion in your organizations.

MR. MOGILNICKI: Great. Please join me in thanking our panelists for a wonderful discussion.

(Applause)

* * * * *

CERTIFICATE OF NOTARY PUBLIC

I, Carleton J. Anderson, III do hereby certify that the forgoing electronic file when originally transmitted was reduced to text at my direction; that said transcript is a true record of the proceedings therein referenced; that I am neither counsel for, related to, nor employed by any of the parties to the action in which these proceedings were taken; and, furthermore, that I am neither a relative or employee of any attorney or counsel employed by the parties hereto, nor financially or otherwise interested in the outcome of this action.

Carleton J. Anderson, III

(Signature and Seal on File)

Notary Public in and for the Commonwealth of Virginia

Commission No. 351998

Expires: November 30, 2020

ANDERSON COURT REPORTING
500 Montgomery Street, Suite 400
Alexandria, VA 22314
Phone (703) 519-7180 Fax (703) 519-7190