Consumer Finance Monitor (Season 5, Episode 27): Using Artificial Intelligence in Underwriting Consumer Loans, with Special Guest Teddy Flo, General Counsel and Secretary of Zest AI

Speakers: Alan Kaplinsky, John Culhane, Michael Gordon and Teddy Flo

Alan Kaplinsky:

Welcome to Consumer Finance Monitor Podcast, where we explore important new developments in the world of consumer finance. I'm Alan Kaplinsky, Senior Counsel at Ballard Spahr, and I am formally the Practice Group Leader of the Consumer Financial Services Group. And I'm very pleased to be your host today on our podcast, which is a repurposing of a recent webinar that we conducted. The title of our show today is the CFPB's position regarding the use of artificial intelligence in underwriting consumer loans.

Alan Kaplinsky:

Very often throughout the program, you will hear the acronym AI and that's shorthand for artificial intelligence. This is a very important topic because an increasing number of companies that are engaged in consumer lending, both banks and non-banks are now using artificial intelligence to underwrite their consumer loans instead of a legacy system using a credit score that is given by one of the three major credit bureaus. And today we're going to really take a fairly deep dive into this phenomenon of AI and why people find it and lenders are finding it very, very interesting to use and why they seem to be gravitating in that direction.

Alan Kaplinsky:

The conventional wisdom is that AI credit scores are more predictive of an applicant's ability and intent to repay a loan than other methods and enables lenders to approve more loan applications without taking on added credit risk. One of the most controversial issues facing the CFPB is how to deal with the Fair Lending and Fair Credit Reporting Act implications of creditors using AI to underwrite consumer loans. There seems to be a great deal of suspicion at the Bureau as to whether these AI-generated credit scores comply with the Equal Credit Opportunity Act and the Fair Credit Reporting Act.

Alan Kaplinsky:

And so during today's program, we'll discuss the CFPB's pronouncements regarding its fair lending concerns about companies using AI to underwrite consumer loans, addressing questions concerning whether the CFPB will use the disparate impact theory in determining whether a lender's use of AI to underwrite loans violates the Equal Credit Opportunity Act. We'll also look into whether the CFPB's recently released guidance sheds any real light on the reasons for adverse actions that companies using AI to underwrite loans should provide.

Alan Kaplinsky:

And finally, we'll look into whether the CFPB will be using enforcement to attack AI loan underwriting programs, and importantly, should companies worry about whistleblowers, because that was one of the issues that got raised by the CFPB in its recent pronouncement regarding the need to provide adverse action notices under the ECOA. And we will toward the end of the program, be focusing on that issue, and why is it that the CFPB made mention of it? It typically doesn't mention it in other pronouncements that it makes.

Alan Kaplinsky:

Our speakers today will discuss their experience with CFPB examinations of lenders' use of AI to underwrite loans and what they've learned from that, as well as to provide practical advice for companies using AI to underwrite loans, including preparation for CFPB exams, and will also address the use of AI in some other areas other than underwriting, although we're not going to get into that in any detail today. Probably that will be fodder for a later program regarding AI. But we're going to touch on how AI is becoming a very important tool throughout the entire lifecycle of the making of a consumer loan from marketing all the way through underwriting and collections and servicing.

Alan Kaplinsky:

So let me introduce our panelist speakers today. And first, our special guest, Teddy Flo, a former colleague of ours at Ballard Spahr, who a few years ago went on to become General Counsel and Secretary of Zest AI, a company that provides software, enabling lenders to use AI to underwrite loans. Teddy works across the company to execute its strategy and ensure the compliance, litigation contracting, and intellectual property operations run smoothly.

Alan Kaplinsky:

Let me now introduce two of my colleagues from Ballard Spahr. First, John Culhane. John's a partner in our Consumer Financial Services Group and a member of the firm's Fair Lending Team. And John has done a lot of work on the topics that we're going to be talking about today.

Alan Kaplinsky:

And then finally, my pleasure to introduce Mike Gordon, who recently joined us as a partner. He's a former Senior Official at the CFPB with two decades of experience in consumer financial services law. In addition to being a former regulator, Mike's also been the General Counsel of a fintech lending platform. And it should be noted that Mike was on the startup team that established the CFPB in 2010. And he maintained several senior roles at the agency, including senior counsel to the director at that time, Richard Cordray.

Alan Kaplinsky:

So this is how we're going to proceed today. First, turn it over to John and have John give us the legal framework, I guess, I often use cliche, he will level set, he will tell you what the relevant laws and the relevant pronouncements of the CFPB have been so far with respect to this important topic.

Alan Kaplinsky:

Then we will go to a Q&A, and I have developed a number of questions that I'm going to propose to Teddy. And we will have some back and forth discussion about a whole range of things that I'm sure are on the minds of many of you, if you are considering the use of AI.

Alan Kaplinsky:

And then finally, and appropriately, we'll be turning to Mike Gordon, who will tell you about the whistle-blowing concern. And Mike having been at the CFPB has a lot of experience in dealing with complaints that are brought to the CFPB by so-called whistleblowers. So without further ado, let's go to John to level set.

John Culhane:

Thanks, Alan. So I'm going to get us started by talking about some of the developments that occurred in the prior administration at the CFPB. So we'll have a sense of how the conversation has changed as we've moved forward. And the most important document under the prior administration was this blog post by Patrice Ficklin and some of her colleagues, it was called an Innovation Spotlight. And the topic subject of the spotlight was providing adverse action notices when using artificial intelligence or machine learning models.

John Culhane:

And it was actually very positive from an industry standpoint, the blog basically positive that the existing builtin flexibility for adverse action notices could be compatible with the use of artificial intelligence models. And in that regard, the blog post was referring to comments in the official staff commentary to Regulation B that say things like no more than four reasons are necessary. You don't have to describe how or explain why a particular factor adversely affected an application, or if you're using a scoring model, even explain how that variable relates to credit worthiness. So there's a fair amount of flexibility there. And then there's some guidance also about how to provide adverse action notices with the use of scoring models.

John Culhane:

To the extent that that wasn't sufficient, the blog post referenced a number of policies that existed at the time as ways that industry could use to resolve any uncertainty as to how to provide adverse action notices by dealing with the CFPB. So those included the Trial Disclosure Policy, TDP, which specifically mentioned adverse action disclosures as a disclosure policy that the CFPB would consider tweaking, the No-Action Letter Policy, which has sort of had its rise and fall under the new administration. And what used to be the Compliance Assistance Sandbox Policies, always that the then administration felt that any uncertainty or any questions as to how to provide adverse action notices when using AI or machine learning models could be addressed. And as I said, this was a very positive blog post. It really supported the use of this technology, particularly with the possibility that using the technology would allow for expanding access to credit and benefiting consumers.

John Culhane:

So shortly after the blog post went up, the CFPB then had a tech sprint on electronic disclosures of adverse action notices. And we discuss this in our Consumer Finance Monitor blog, but it didn't really get a whole lot of traction or a whole lot of publicity. One of the topics that the CFPB indicated they would like to have addressed in the tech sprint was how to refine the methods for translating model factors into adverse action notices. But that really wasn't a big topic in the tech sprint. Most of the participants in the tech sprint focused on making notices more clear, improving their educational content. The participants divided into teams and conducted various projects. And some of those projects did attempt to improve algorithmic credit decision notices, but none of them were really particularly endorsed by the CFPB.

John Culhane:

Now, there still are some links on the website to some videos that the CFPB put up about the tech sprint and about some of the top scoring teams and particularly ones that scored high for creativity and innovation, ones that scored high for effectiveness and impact, and ones that scored high for market readiness. But again, the CFPB didn't really take the tech sprint and use it as a way of discussing for the industry, how they might proceed to provide adverse action notices when using artificial intelligence or machine learning models.

John Culhane:

Now about the same time that these things were going on, the CFPB had put out this gigantic request for information on the Equal Credit Opportunity Act and Regulation B, had a very broad array of topics, just about everything that was a hot issue and continues to be a hot issue at the time. And one of the topics on which the CFPB invited comment was, of course, how to comply with adverse action requirements when using complex models. They didn't say a lot about what they had in mind, but they did invite comment. And remember, the time period during which this is going on is under the prior administration. So there's not really much of a battle about how to comply. The battle is really over how much instruction to take from the Patrice Ficklin blog post.

John Culhane:

The NCLC comments on this, but they didn't really wade into the specifics. They just urged the CFPB to ensure that use of models would be compliant with adverse action notice requirements, again, without going into specifics. The Center for Responsible Lending didn't really address this at all. They focused their comments on other issues that were teed up in the

request for information. And both the Consumer Bankers Association and the American Bankers Association supported the guidance in the Patrice Ficklin blog post, which, as I mentioned, seemed very favorable and seemed to endorse the ability to comply with the requirements of Regulation B using the guidance in the existing commentary. Now, the American Bankers Association did recommend that the CFPB provide a small entity compliance guide for adverse action notice issues, both to cover the use of complex models, but also to cover many of the issues that come up generally with adverse action.

John Culhane:

So there's not a lot of discussion here. There's an endorsement of the Patrice Ficklin blog post, and then we have a change of administrations. And we have a new administration come in and we have Rohit Chopra, and he spends maybe the first several months of his administration to the extent that he's talking about the use of artificial intelligence or machine learning, declining the possibility of digital redlining. So he makes comments in connection with a settlement agreement the CFPB enters into. He raises the issue when testifying before the House and before the Senate. He mentions it when he's interviewed, I think by The Washington Post. He doesn't say a lot in any of those contexts, in any of those forums about adverse action notices.

John Culhane:

And then on May, I guess it's May 26th, the CFPB, having started this policy of issuing circulars, drops the circular directed specifically to the use of complex algorithms. And it's very, very intimidating. It basically tees up a question in a way that suggests that it's not really possible to comply. So the question that the CFPB says that it will address is, when creditors are making credit decisions based on complex algorithms, and here's sort of the sandbagging of the question, that prevent creditors from accurately identifying specific reasons for denying credit or taking other adverse actions, when they do that, how do they comply? Do they need to comply with the ECOA's requirement to provide a statement of specific reasons?

John Culhane:

And then the responses are somewhat self-evident. Yes, creditors need to comply with the adverse action notice requirements. There's a footnote, and we'll talk more about this, saying that the creditors need to be using an interpretable model, or if they are using a model that provides for post-hoc explanations that approximate the performance of the model, they have to be able to validate the accuracy of those explanations. And some fairly obvious statements. Complex technology is not a defense to compliance.

John Culhane:

And then, I guess, really the key part of this circular, if your technology is too complicated and you can't give reasons for adverse action, then you shouldn't be using that technology at all because you're violating the Equal Credit Opportunity Act and Regulation B. But there's no guidance here really about how to do anything. There are just these platitudes that if you use artificial intelligence or machine learning, you still have to comply. And then, the close of this is a call for help. The CFPB asks tech workers to blow the whistle on violations. So, that's where we are right now. Let's go to the Q&A part, Alan.

Alan Kaplinsky:

Yeah. John, thank you very much. Teddy. Let me welcome you again. It's great to have you with us.

Teddy Flo:

It's good to be here, Alan. And thank you for inviting me.

Alan Kaplinsky:

Okay. Before we get into some of the harder questions, let me throw you a softball or too. The softball question is, are lenders really using AI technology for underwriting loans? And if so, do you have any idea of the percentage of lenders or the number of loans dollar wise that are being underwritten through AI? What say you?

Teddy Flo:

I sure hope so, Alan, because that is our business. We help lenders of all sizes, banks, credit unions leverage the power of machine learning to reduce loan losses and also increase their approval rates, and also to close the gap in approval rates between protected groups and white folks. We're really proud of that work and our customers see a ton of value there. To get more directly to the point, there are a number of folks using this. Large banks have enormous data science teams who are exploring this question and whether or not they're actually using the models today, it is certainly on the roadmap for all the customers that we talk to. You have folks like Zest who produce machine learning powered scores for financial institutions. And then you have other lenders like Upstart, for example, who also leverages the technology.

Teddy Flo:

John Culhane mentioned the tech sprint. If you were to watch those videos, I think each of those individual companies mention their use of AI in underwriting the financial products that they're offering. And just to give you a sense of numbers, at Zest, if you were to just look at our credit union business, and to take the credit unions that have signed with Zest or that are in contracting right now, those credit unions serve 17.2 million Americans. So you're looking at 6% of the US adult population is being impacted by just tiny Zest. And so this is a technology that is growing in use and helping a lot of people.

Alan Kaplinsky:

Yeah. Teddy, what types of AI are lenders using?

Teddy Flo:

That is a good question. I can tell you what we use. AI is really a broad term that encompasses any number of things, including the evil AI that we see in the Terminator movies and in the Matrix. That is not the sort of AI that we are using at Zest. We use an AI algorithm to produce the risk scores, but the algorithm is locked, meaning that it will not change on its own before it has been thoroughly vetted and reviewed by compliance and risk professionals. So once the model is deployed and being used to underwrite loans, it sort of stops changing the way it thinks about consumers until you take it offline, make changes, and then put it back into production. And that's one of the key ways to make sure that AI and ML that's being used is safe and compliant.

Alan Kaplinsky:

Okay. I mentioned in my introduction, some of my thoughts as to what I perceive the benefits of using AI are compared to using conventional legacy credit scores from one of the three credit bureaus. Well, maybe you can amplify on that somewhat.

Teddy Flo:

Certainly. Yeah. The lenders that use Zest models routinely see increases in approval rates between 15 and 30% with no increase in credit risk. Other lenders who've had more of a problem with losses will see decreases in loan losses between 30 and 50%, again, with no decrease in their approval rate. And then the lenders that leverage the AI-powered fair lending technology will see increases in approval rates for women and people of color that are extremely significant. One of our larger credit unions saw an increase in approval rates for women of 22%, which was more than the increase in approval rates for men when they implemented the technology. Another saw increases in approval rates for African Americans by 30%, which was substantially more than the increase in approval rates for the white folks that were applying. So the technology has extraordinary power to reduce risk, expand access to credit, and also close the credit gap. And so the benefits of using the technology are really tremendous. And there are certainly risks that we're going to talk about today that need to be addressed, but responsible users of the technology do so.

Alan Kaplinsky:

Yeah. Teddy, you said there are some risks, and I assume you're referring principally the legal risks. And John talked a little bit about the legal framework that exists and some of the pronouncements that have been made by the CFPB, not a lot there.

What to your way of thinking are the key regulatory areas of concern? What laws worry you? Maybe that's not the right way to put it. Maybe I should say, what are the laws that you focus on the most?

Teddy Flo:

Certainly, yeah. There's a number of requirements that we are attentive to, but really three big ones that we think about every day. The first is ECOA and Reg B's requirement to produce adverse action reasons. As John mentioned, there is flexibility in the guidance and in Reg B to do it, but fundamentally the reasons you provide consumers for taking an adverse action need to be accurate. And we'll talk more about how we get to an accurate result later, but that's really critical. And it's something that actually we pointed out to the CFPB two years ago. We submitted a comment letter on this very issue in response to one of the CFPB's requests for information. So the first, as I mentioned, is the requirement to produce accurate adverse action reasons.

Teddy Flo:

The second is the Equal Credit Opportunity Act's prohibition on disparate treatment. It's really important when you're using a machine learning model that may have 150 or 1,000 variables that it considers to make sure that none of those variables are highly correlated with protected class status. And that also relies on accurate explainability technology.

Teddy Flo:

And then the third area that we focus on is ECOA and Reg B's prohibition on their disparate impact rule. And it's also important to make sure that when you're using an AI model, or really any algorithm, that we are not causing unlawful disparate impact, and that we are doing everything we can to make sure that there's less discriminatory alternative underwriting model. And that's really where Zest AI's fairness technology has focused is on that advanced LDAs or less discriminatory alternative for LDA searching. So those are the three big areas that we focus on, but of course there's privacy concerns, data security, and the like. But I think for today's purposes, those three are-

Alan Kaplinsky:

Yeah. Sure. And I assume there are some state laws that need to be focused on as well. Am I right?

Teddy Flo:

There are state law equivalents for each of these requirements. Most of our customers are large national banks or GSEs or credit unions. So the main focus would be the federal statutes. You're right, there are certainly state laws to consider.

Alan Kaplinsky:

Yeah. I mean, are there, to your knowledge, any state laws that focus specifically on AI and its use in credit underwriting, or are they more just mini ECOA laws and many FCRA laws?

Teddy Flo:

Most of the ones that we focused on are the mini ECOA and FCRA, but a number of states are actually considering legislation, particularly focused on algorithms and algorithmic fairness. And we monitor those and we do communicate with state regulators and lawmakers about them.

Alan Kaplinsky:

Yeah. I mean, I would think that looking further down the road, it could be a major headache if a whole bunch of states start legislating in the area, particularly, they can start getting into various algorithms and various factors that can be used. It could be quite a challenge to comply with all of that.

Teddy Flo:

Whenever you have a number of state laws on a single issue, it does create some compliance issues. But stepping back a moment, by and large at Zest, we celebrate notices like the circular, that call for strict compliance standards and that call for accurate explanations because there are ways to do it technically. Our technology, for example, we rely on academic articles published in journals, such as Nature, for example, where Darwin published his first paper for the theoretical justification for our explainability methods. So we rely on the most accurate and up-to-date technology. And then we perform empirical validation of the accuracy of our explanations. And so there are ways of doing this right and there are ways of complying that everyone should be focused on, if we're going to do this in a way that serves consumers and that uses algorithms that are safe and creates safe and sound banking practices.

Alan Kaplinsky:

We've talked about ECOA and the FCRA, we haven't yet mentioned UDAAP, unfair and deceptive acts and abusive practices, the sort of umbrella statute in the Consumer Financial Protection Act used by the CFPB. The CFPB, how does their use of the unfairness doctrine under UDAAP, how does that impact machine learning?

Teddy Flo:

I think it's a very interesting move by the Bureau to use the unfairness component of UDAAP to address fair lending issues, and it'll be very interesting to see how that evolves. At the moment it creates a number of open questions that we are thinking about at Zest and that we talk to folks like you and law firms and consultants about, particularly, for example, like the applicability of the explainability and fairness requirements to other areas of lending such as fraud models. It's typically been understood that fraud models could consider race and gender in making a fraud determination. How would the use of the unfairness doctrine impact that sort of historical understanding?

Teddy Flo:

Or you could perhaps look at marketing models. AI has been used for marketing models for a long time, but folks have not seen the need to use explainable AI for their marketing models, except if they have their own internal risk standards, but they haven't seen it as a compliance requirement. Does the use of the unfairness doctrine change that? That's unclear. And so we'll be watching closely how the Bureau decides to tackle these and other open questions. Right now, it doesn't have sort of a direct impact on ML itself. But as I mentioned, these open questions could change that or how these open questions are answered, could change that.

Alan Kaplinsky:

Okay. All right. Let's turn to this issue of explainability. And I noted that we've gotten a lot of questions already dealing with explainability and interpretability. In the recent circular that the CFPB issued, they noticed the importance of providing accurate adverse action reasons, even when underwriting is using complex algorithms. It's not an excuse as John Culhane said. What do you think prompted them to do that?

Teddy Flo:

That is a great question. And I am no expert on reading tea leaves. My guess, if I were guessing, I think the CFPB in conducting their supervision work in the industry may have seen some actors out there who take the position that you can use any old explainability method for AI models and it's fine. And so I think the circular may have been aimed at them. For example, there's a recent study put out by FinRegLab, where they looked at particularly different methods for explaining machine learning models and the ability of the various methods to produce accurate adverse action reasons. The conclusion was, a really common sense conclusion, but it's now been empirically validated that you need to use the right tool for the job. Certain explainability methods are okay for certain types of logistic regression or machine learning models, but for more advanced models, you need a more advanced explainability method. And that was sort of the conclusion that was reached in

the FinRegLab study. And I see the circular as supporting and amplifying that conclusion that you need to be using the right tool for the job.

Alan Kaplinsky:

Right. And then this issue of interpretability. When the circular used that word interpretable, what do you think it was getting at and is that different than explainability?

Teddy Flo:

There's a great question. Different scholars and academics will use the word interpretable differently. There isn't a well or broadly accepted definition of that term, as it is applied to machine learning models. If you go to GitHub, which is sort of an opensource library, there are a number of resources there where they talk about machine learning and some of those resources, what we at Zest describe as explainability, they will use the word interpretability to talk about some of those same methods. And so that's one thing that I think would really benefit the industry is sort of a coming together of financial services providers to develop well-understood definitions of the concepts of interpretability and explainability, so that we all can know that we're doing things correctly.

Teddy Flo:

You'll read some academic papers or listen to some industry participants and they will talk about interpretability to mean that human eyeballs have to be able to examine the algorithm and be able to intuit how a risk score will change based on inputs into the algorithm. That's sort of what some folks mean when they talk about interpretability. We put out a white paper that's available on our website where that definition of interpretability is really untenable. Aside from the fact that most humans are not that great at arithmetic and would have a hard time doing the numbers on their fingers, looking at the algorithm, there's a number of factors external to the algorithm that you need to know in order to understand how that algorithm is going to behave. And so that sort of overly simplistic definition of interpretability, we think it can't be what the CFPB meant, but it isn't exactly clear either.

Alan Kaplinsky:

Footnote one to the circular says, and I'm quoting here, "While some creditors may rely upon various post-hoc explanation methods, such explanations approximate models and creditors must still be able to validate the accuracy of those approximations, which may not be possible with less interpretable models." What does that all mean?

Teddy Flo:

That was a good question. So let me actually take a step back and talk a little bit about explainability and some of the more advanced methodologies. And that will sort of, I think, inform the answer to your question. If you were to review the FinRegLab study, one of the findings that you would see is that Shapley-based explainability methods, and I'll get into a definition of that in a moment are really good at explaining simple algorithms and very complex machine learning algorithms. Those are the methods that we use at Zest, and that we think responsible users of AI deploy when they're explaining their models and producing adverse action reasons.

Teddy Flo:

But let me talk a little bit about what that is. Shapley values or Shapley-based explainability is based on the work of a guy, no surprise to you, by the name of Lloyd Shapley, who is a Nobel Laureate, whose work has been examined by Scott Lundberg, for example, at Microsoft, whose article was published in the journal Nature and determined to be the most accurate way of explaining how an algorithm works. And what Lloyd Shapley did a number of years ago, this was in the '70s, was he was looking at this question of how to attribute value to various players in a cooperative game. So if you were to look at soccer or basketball, there's the guy that puts the ball into the goal and scores the point, but he didn't get there without a lot of passes and blocks and what have you. And so what Shapley was trying to understand was how is it that we can understand how all the

other players contributed to the final score of the game, taking into account all of the various interactions. And his work was determined to be the only way of accurately answering that question.

Teddy Flo:

And what Scott Lundberg found was that those Shapley methods were perfectly in explaining how machine learning models work. Because, if you think about a machine learning model, the model is sort of the game. The risk score produced by the model is very much like the score in the soccer game or the basketball game. And the various players on the field work exactly how the various variables in the model work. And so when you take the algorithms and the equations and the theory that Lloyd Shapley developed and apply it to machine learning algorithms, you get really accurate explanations about how models work. And so what you'll see if you're looking at a fully explained machine learning model is a table that lists every single feature in the model and a contribution, how each feature contributed to the risk score that the model produced. The contributions sum up to the model score. So it creates a very intuitive understanding about which features contributed to that final risk score. That's what we use at Zest. And that's what responsible vendors use to deploy the technology.

Teddy Flo:

There are other methods, however, LIME, for example, that don't rely on Shapley's work and we think what the CFPB was referring to in footnote one. Because LIME, what it would do is, if you were to picture your model as a curvy line on a page, and you're trying to explain sort of how that model is behaving, what LIME does is it takes it and creates a number of straight lines that sort of follow along the path of the curvy line. And then it explains how the straight lines are behaving. And so it's explaining an approximation of the true model, and that's what we think footnote one is referring to primarily. And perhaps there are other methods that rely on approximations as well. At Zest, we think those methods aren't as accurate as they could be. And it is important to perform that theoretical validation.

Alan Kaplinsky:

Yeah. Let's circle back to fairness. How do you think the CFPB can address the issue of reducing disparate impact?

Teddy Flo:

Yeah. We think it's important that lenders and the CFPB focus in on really the third prong of the disparate impact test. The first prong is that first you will look at a model or an algorithm to see if it produces disparate impact. The second, you would see if there is disparate impact. Is it business justified? Are there differences in approval rates between protected groups justified by some legitimate non-discriminatory business purpose? And then third though, you would look at whether there's a less discriminatory alternative to the algorithm that you're using. And really for a long time, there haven't been great tools out there to perform that sort of LDA searching.

Teddy Flo:

At Zest, we developed a proprietary methodology, we have a couple of them that use machine learning to perform that less discriminatory alternative search by searching through billions of alternative models, to find one that produces similar business outcomes, but that may perhaps be more fair. And it's very often the case that we can find using those sorts of search methods, models that are substantially more fair than the original, the benchmark model, but that are no less accurate or the drop in accuracy is statistically insignificant.

Teddy Flo:

And so that's one of those important tools that we think it's important for the industry to begin adopting and because we think the CFPB is focused on it as well. Once the CFPB in supervision starts using a technology like that, that performs advanced LDA searching, they're going to uncover their lending problems at lenders that lenders who hadn't gotten ahead of the issue will not have seen themselves. And so that's one area that we think it's really important for the industry to begin to focus.

Alan Kaplinsky:

I've asked some very specific questions. I want to ask a very general question of you. Let's say you're in our audience today and listening to the presentations that we've made, and we've convinced you that your bank or your lender ought to start using machine learning. Where do they go? Where does it all begin? How do they go about implementing AI?

Teddy Flo:

Certainly. There's a couple of ways to do it.

Alan Kaplinsky:

I guess they go to you, right?

Teddy Flo:

Well, you could call us direct, that's right 1-800... No, I'm just kidding. We don't have an 800 number. But yes, you could definitely reach out. We'd be happy to help you, but there are a number of frameworks you could use. I think you could develop a technology in house. That tends to be extremely costly because you have to have a lot of expensive data scientists. You have to do the appropriate compliance work to adopt responsible AI framework within your organization. And then you have to thoroughly vet the models and document them. And that's one thing we actually haven't talked about is the documentation piece, which can be extraordinarily burdensome to document even an old-school logistic regression model. But documenting an AI model, if you don't know what you're doing, can take a really long time given the number of features involved.

Teddy Flo:

But anyway, so you can create your own in-house system, where you have the data scientists and you have the compliance professionals and you have the risk teams that can do all that needs to be done, or you can find a vendor like us, and there are others out there that do this as well. But really, it's important to make sure the vendor you're working with understands these compliance issues inside and out, because it's easy to get it wrong if you don't know what you're doing. So I think those are really the two main options on how to go forward.

Teddy Flo:

But if I can put in one other sort of plug, there are a number of nuances and details about AI explainability that we probably won't have time to cover today. At Zest, we produced an 18-page white paper shortly after the circular came out, that is available on our website, that goes into great detail on a number of these issues. If you're having trouble sleeping, feel free to stop by the website and download a copy for your perusal.

Alan Kaplinsky:

Yeah. Let me ask you one other general question. And that is, where do all these variables come from, these factors? Are they obtainable on the legacy credit bureau reports? Is there any other sources that are used? I mean, where does that come from?

Teddy Flo:

Yeah. At Zest, we rely on FCRA compliant data sources, but before we get to what we do rely on, let me talk about what we don't rely on. You may have heard myths in the industry that folks are using Facebook data and internet search history and all of that. Maybe some people are, we don't consider that to be appropriate. We don't consider that to be appropriate for any number of reasons. So we rely principally on the three main credit bureaus to get the data and used to power the model.

Teddy Flo:

And one in particular that's considered sort of an alternative data source is LexisNexis, and they have a number of useful attributes that can be used to amplify the accuracy of models. And they would say this, you don't use every single one of the features that they offer because they're offered for different purposes, but the ones they offer for credit underwriting can be extraordinarily useful in amplifying the power of the models. And so we derive all of the features in the machine learning models from one of those four data sources and perhaps some data that the customer may have on its own to produce the risk scores.

Alan Kaplinsky:

Got it. Got it. Okay. Well, Teddy, is there anything that I forgot to ask you, something, another point that you want to make that our listeners ought to be aware of?

Teddy Flo:

I think we've covered the highlights here. As I mentioned, there are a number of details and nuances to the discussion. We have that white paper, it's a good resource, and I'm always happy to chat with anyone about these issues in the future.

Alan Kaplinsky:

Sure. Before we turn it over to Mike, that's going to talk about this scary subject of whistleblowers, I'm going to ask my colleague, John Culhane, John, do you have anything that you want to add to what Teddy has said?

John Culhane:

The only thing that I want to add to what Teddy has said is that I think if you're using AI or machine learning models, what you should take from the CFPB's notice here is that they are going to expect you to explain the reasons, explainability is almost going to mean justifying the reasons that are generated by the model to the CFPB. So you really have to be focused on that. And as Teddy has discussed in detail, you have to be very clear about the methodology that's being used to identify those reasons. And you have to be able to talk about that to the exam team or the fair lending team that comes in. I think the burden is going to be on the institution to explain that what it's doing is compliant. It's not going to be that the examiners are going to come in and tear the model apart. I think that's a very important message here.

Alan Kaplinsky:

Yeah. Yeah. Thank you, John. Okay, let's go to Mike. Mike, what can you tell us about whistleblowers?

Michael Gordon:

Thanks, Alan. So I'll keep my remarks brief in light of the time here, but I think it is worth highlighting the risk from whistleblowers and what the Bureau's attitude has been. So to level set, when Dodd-Frank was passed, it included a provision protecting whistle blowing, if it's to the Bureau or to other agencies with respect to alleged violations of consumer finance laws. It did not however include a bounty or an award-like system that the SCC has for whistleblowers, which in the SCC's mind has been a very effective program. I'm not sure all the reasons and the legislative history for this, whether it was an oversight or intentional, or what the behind the scenes haggling was on this one, but without that incentive, you'd imagine a whistleblower program is much less robust. And even under Director Kraninger during the Trump Administration, the CFP proposed adding an award program. She sent the draft legislation over to the Hill for just such a program, which has been introduced in both the House and Senate by Democrats. But I don't think has much bipartisan support, but I wouldn't necessarily expect that to advance anytime soon.

Michael Gordon:

Aside from the statutory basis and protections for whistleblowers, the Bureau has sort of relatively infrequently discussed the topic. I mean, they're big on consumer complaints generally, obviously, and have a huge machine to bring in and analyze

consumer complaints and make sure they're processed. But as to whistleblowers specifically, they've called them out in a couple instances and they do seem to revolve around issues of technology. For example, when Cordray took over, there was an initial announcement that the whistleblower program exists and inviting people to submit. But then in December of last year, the Bureau announced kind of a updating of the program and specifically in that announcement called for tech workers to make reports to the Bureau. And it noted that these algorithms can remove discriminatory impacts, but also can embed them. And as Alan mentioned earlier and others have noted, the Bureau seems generally suspicious of the use of artificial intelligence and algorithms, pejoratively referring to them as black box. And just generally when they speak publicly today about these kinds of technology solutions, the tone is very suspicious and negative.

Michael Gordon:

So why would they focus on technology workers for their appeals to whistleblowers? And I think there are a few reasons for that. First of all, it's hard for any agency to understand what's going on in these algorithms. It's constantly changing in the marketplace. Generally speaking, these are proprietary systems and models that are closely held and agencies don't have real-time visibility into how they're built with few exceptions, if they're examining a big bank or something that has the use of it and they care to dig into it, I suppose they can ask. But generally speaking, there's not a lot of visibility, nor is there necessarily a lot of agency expertise to understand the models, even if they have access to them. This isn't necessarily a criticism of the Bureau itself, but I think throughout the government, I think this is a problem. As technology changes, it's hard for the government to keep up in terms of expertise and knowledge. And so it's an area where you can see why it would be particularly helpful for the Bureau to have an insider, explain the model on how it's impacting consumers.

Michael Gordon:

The final thing I would say is, it's possible that the Bureau is looking for whistleblowers to help it identify targets for supervision. It made an announcement earlier this year that it's going to use a risk-based approach to supervising entities', that it doesn't currently have authority to supervise. And if they get a whistleblower, that's willing to walk them through allegations of disparate impact about a model, I can see that quickly turning into a decision by the Bureau to do a full-blown exam of that company. So those are my thoughts on whistleblowers generally, and the Bureau, what their attitude is towards them.

Alan Kaplinsky:

Okay,.well, we've come to the end of our program today and I want to thank again, our very special guest, Teddy Flo. I want to thank my colleagues, John Culhane and Mike Gordon. And in particular, I want to thank all of you who downloaded our podcast show today and took the time to listen to it. And I would encourage you to do that every Thursday, because we release a new show every Thursday. The final thing I'm going to say is our podcast show was selected last year as the second best law firm podcast show among the 200 largest law firms in the country. And we're very proud of that achievement. So with that, I want to wish everybody a good day.