

Consumer Finance Monitor Podcast (Season 9, Episode 36): AI in Financial Services – Consumer Protection Challenges in the Age of Artificial Intelligence

Speakers: Alan Kaplinsky and Delicia Hand

Alan Kaplinsky:

Welcome to the award-winning Consumer Finance Monitor podcast, where we explore important new developments in the world of consumer financial services and what they mean for your business, your customers, and the industry. This is a weekly show brought to you by the Consumer Financial Services Group at the Ballard Spahr law firm. And I'm your host, Alan Kaplinsky, the founder and former practice group leader for 25 years, and now senior counsel of the Consumer Financial Services Group at Ballard Spahr. And I'm delighted to be moderating today's program.

For those of you who want even more information about the topic that we're going to be discussing today, or for that matter, anything else in the world of consumer finance, don't forget our blog, which also goes by the same name as our podcast show, namely Consumer Finance Monitor. We've hosted our blog since 2011, so there's a lot of relevant industry content there.

We also regularly host webinars on subjects of interest to those in the industry. So to subscribe to our blog or to get on the list for our webinars, please visit us at ballardspahr.com. And if you like our podcast, please let us know about it. You can leave us a review on whatever platform you have used today to access our podcast show.

Also, please let us know if you have ideas for other topics that we should consider covering or guests that we should consider inviting as speakers on our program. So let me introduce our topic today. This is a topic we've been covering in a lot of depth in the last couple of years because it's becoming of increasing importance. And what I'm going to be talking about and what we're going to be talking about, and when I introduce my guest, what my guest will be talking about is the use of artificial intelligence, or we'll usually call it AI, and what it means for consumers, markets and regulation.

AI is already making decisions that affect people's financial lives every day, whether it's determining credit worthiness, flagging transactions as fraud, setting prices, or even providing financial advice. But as powerful as these systems are, they're often, unfortunately, opaque, difficult to challenge, and in many cases, not designed with consumer protection as the primary objective. A recent report highlights a striking disconnect. While approximately 75% of Americans are concerned that AI in financial services could lead to bias or unfair treatment, only 8% believe that current laws protect them. That gap between technological innovation and meaningful accountability is exactly what we're going to explore today.

And I am pleased at this time to introduce a true expert in this area, namely Delicia Hand. Delicia is a nationally recognized expert in digital civil rights, AI governance, and consumer protection. She serves in a leadership role at Consumer Reports where she focuses on ensuring that emerging technologies, particularly AI, advance consumer welfare rather than undermine it.

She is the author of an important new report by Consumer Reports titled AI and Consumer Finance: A Landscape Analysis of Market Forces, Evaluation Frameworks, and Regulatory Gaps. The report examines how AI is transforming consumer finance and emerging from that transformation and what policymakers, regulators, and industry participants should be doing about it. So, Delicia, a very warm welcome to our show.

Delicia Hand:

Alan, thanks for having me. We're very glad to be here.

Alan Kaplinsky:

Terrific. Well, I've got a lot of questions for you today that I know you're going to know the answers to, but they're the questions that I raised in my own mind, and I try to put myself in the head of a consumer and what a consumer might be

thinking of this phenomenon that we call AI. And that's how I developed these questions. So my first question is what prompted Consumer Reports to undertake this major study of AI at this particular moment?

Delicia Hand:

Sure. Well, honestly, the timing felt urgent, but it also built on years of work that we've been doing at Consumer Reports. We built probably about nine years ago now something called the Digital Standard. When it became clear that our everyday home products, our TVs, phones, et cetera, are connected and that there were some baseline privacy, transparency, security standards that needed to exist and didn't really.

And so we've been doing that. I left the CFPB for over four years ago and saw that there was a real opportunity to build on that digital standard to think about what standards ought to exist for fintech products. And so we got on the road and had been evaluating peer-to-peer apps, banking apps, buy now, pay later, all the latest sort of app-driven fintech products. And so when we evaluated banking apps around 2023, we kept running into same issue, which is we had a sprinkling a handful of questions about traditional machine learning issues, but we were in a moment where Generative AI was emerging and being used and AI was clearly embedded in how products and companies were making decisions, but there were no consistent consumer-centered standards to evaluate against it.

And over the last couple of years, you've just seen that continue to evolve. So for us, it really became clear that we were in a moment of a lot of evolution and we kept seeing piecemeals of frameworks addressing the consumer protection problem, but really in isolation. Add on top of that, that regulation is always in catch-up mode. Industry and others innovate and deploy, and then regulation responds to a crisis. So we were seeing that, and that's what led us to undertake this analysis.

Alan Kaplinsky:

Yeah. Well, let me ask you one quick question, first of all. While you were at the CFPB, were you in the research area? Was that your area that you were in?

Delicia Hand:

No, I wasn't actually. So when I was last at the CFPB, I was in a leadership position and I oversaw our consumer complaints division as well as our external engagement. And so that provided a real opportunity for us to be, frankly, getting a lot of consumer data about where consumers struggled in the marketplace, but also meeting with different companies, particularly on the fintech side who were engaging with us.

And so I think that was partly the inspiration for the work that I'm now doing at Consumer Reports, which I recognize that inherently because of my regulatory hat, we couldn't have truly consequence-free conversations. And so there were limitations to what we were able to learn and vice versa, what we're able to bring forth.

So was able to take some of those learnings and then think about, well, what if we could have a no consequences type deep inquiry into, well, what's happening in the design and the production of these different products? And what if I was able to go direct to industry with those findings and then give them an opportunity to actually make corrections? So that's the business model that we have here that I wasn't really able to do being a regulator.

Alan Kaplinsky:

Got it. So your main survey findings are striking, Lisa, and I mentioned it during my introduction, but let me repeat it. 75% of consumers express concern about bias and unfairness with respect to AI, yet only 8% believe that existing laws adequately protect them. What explains that lack of confidence, Delicia?

Delicia Hand:

Yeah, I think the gap tells you something important. So consumers are always willing to embrace new technologies. They offer convenience. And I think you also see that there's increased access. But I think one of the things that 75, 8% tells you is that

this is not just blind acceptance. They're also paying attention, and they know that this is emerging technology and maybe that not all the protections are keeping a pace.

Also, when these new technologies or products that are optimized, have AI embedded in them are rolled out, consumers are going to run into challenges. So they're not necessarily just anxious about AI in the abstract, they're actually experiencing it. So when a consumer engages with a chatbot and is trying to resolve something, but they can't get out of that doom loop, that's an example, earlier example. So I think it's really pointing to the fact that there is a little bit of disconnect there in terms of when these technologies are out there and the rules of the road in terms of new consumer protections are just fine and consumers can see that in their own experience.

Alan Kaplinsky:

Right. So one of the most provocative themes in the report is your argument that AI adoption is not primarily being driven by consumer demand. Explain what you mean by that.

Delicia Hand:

Yeah, it was fairly provocative. I think one of the things we're trying to point out potentially is that we heard that consumers' top priorities for improving financial services are data security, lower costs, better consumer service. But then we also asked, "Well, what benefits do they actually associate with AI adoption?" And those same priorities actually ranked much lower.

The one thing that they associated most with AI was faster decisions, but then only 18% of people said that speed was a priority for them. And I want to be clear, I'm not necessarily saying we slow down AI. I don't think there is any slowing that down just to the inherent nature of the technology and the opportunity here is real. You can lower costs, there's better fraud detection, more personalized services. Those are things worth achieving. I think what we were realizing is that there's speed happening without necessarily due diligence, and that is a trend we've seen before.

So if adoption is being driven by other things, cloud provider revenue incentives, executive signaling, rather than actual demonstrated consumer benefit or consumer protections also scaling, you end up deploying systems before they're really ready and you run the risk of locking in bad outcomes at scale. And so I think that's one of the things that we're trying to unpack a little bit. And so that was just one of the themes that we identified.

Alan Kaplinsky:

Yeah. And I guess a part of the speed is not just the speed of AI and how quickly it can do decision making, but the fact that there's such intense competition to develop the best, the fastest, the most intelligent model to become number one because everybody recognizes that not every company is going to survive this race that they have.

And I guess it's probably fair to say that the speed and the competition results in consumer welfare or consumer issues, type issues, consumer protection issues not being the highest priority. It is important. I'm sure any of these companies that are involved in this race to become number one would say, "Of course, there's got to be consumer protection." But ultimately it's speed and it's the wanting to win that race because whoever wins that race, and there may not be more than one or two that ultimately survive the race, it's going to be a financial windfall probably like we've never seen before. I mean, that's my reaction. Do you have a reaction to that, that factors I mentioned are resulting in giving less prioritization to consumer protection?

Delicia Hand:

Yeah, absolutely. And you don't even have to take my word for it. If you've ever spent time working with a general purpose chat tool, you will have the experience that I think people that's now more common, which is it's really good at some things and speed is one of them. Processing large amounts of information is another. And it's also fundamentally not good at some very simple things that humans frankly have learned to master. So context is one where in certain situations, even if you're in a rush, you're always checking just to make sure that you're right. So we have learned to not hallucinate.

The tendency might be there, but if we want to keep our jobs, we want to make sure that information that we give to our bosses is fact checked and verified. So that's now part of our processes. This is something that the chat tools still do,

regardless of how much more complex they are and developers and others far more capable than myself are still working to address.

So you see these innate challenges with the tools that are now on the marketplace, but you don't see institutions slowing down to integrate them into everyday life and increasingly complex functions in our daily lives. And so I think there is a clear financial incentive to encourage broad adoption, even premature adoption, sometimes before we have enough certainty about the impact on consumers. We're seeing that all the time. And ultimately-

Alan Kaplinsky:

A good example would be Anthropic's development of the, I think it's called Mythos or Mythos where it almost felt to me like they had created a Frankenstein monster that they couldn't get control of, and it was able to perform functions that really weren't intended by what they were doing. I mean, that's probably Exhibit A, right?

Delicia Hand:

Exactly. Yes.

Alan Kaplinsky:

So our audience may not be completely familiar with what I'm talking about with respect to the model that Anthropic developed where all of a sudden they invited the government to see what they had created and they were looking for input on how to control what they created. If you don't mind, maybe just briefly tell our audience what happened there.

Delicia Hand:

Yeah. So Mythos was or is a very powerful and right now restricted access artificial intelligence model that can rapidly identify and generate software vulnerabilities specifically. And so that can be used for good things. And also it can be used for bad things. So upon announcing the model, Anthropic, for example, revealed that Mythos had autonomously, so that's something to keep note of, discovered thousands of hidden software bugs and flaws in everyday operating systems that are in use. And then it had the potential to do that across different ecosystem.

So this is both an AI that had the capability of patching software, so detecting them to patch them. But I think it became really clear the same force developed for good could also be weaponized. So they've launched a project which is a limited program for just a small group of tech and cybersecurity firms to be able to engage with it. But it really sort of, I think, awakened folks across the globe as to the danger of putting something out there without the guardrails and without frameworks to deal with improper and nefarious use cases.

Alan Kaplinsky:

Yeah. Yeah. Okay. So the report introduces something called the 12 Axes of Evaluation for Assessing AI Systems in Financial Services. I'd like to have a couple of questions. One is what inspired that framework? And second, what do you mean by axis of evaluation?

Delicia Hand:

Sure. I mean, I think what has inspired it is I've been thinking ever since I left the CFPB, so this is even before the GenAI moment, I've just intuitively been wrestling with this notion of, okay, we've got this rich landscape of various consumer protection laws, the ABC land, Alphabet Soup land of consumer protection laws that we live with that have mostly been developed in response to something.

They're all almost sort of backward looking. A crisis happens or enough consumers get harmed we respond and then the world moves forward. I was at the CFPB for 10 years, that was birthed in a certain moment, but then the laws don't change. So what I saw while I was there was fintech was just being born when the CFPB was opening its doors. And you had this 21st

century agency, which was always looking at rules and law established long before the new technologies that were in the marketplace.

So there's always been this little bit of, always had this focus on that disconnect and the asymmetry and the way in which we regulate the marketplace that's continually moving forward, but with rules that were decided decades ago. And so in the GenAI moment, I thought, well, it'd be good to see what the experts out there think is happening with respect to consumer protection.

Folks argue that the existing consumer protection laws are enough. So is that actually true? And so fundamentally that's what I was seeking to do in pursuing this initiative, which is to say, "Well, are the baseline consumer protection frameworks still relevant? Are there gaps? If there are gaps, how then do you go to collect evidence of what AI might be doing in these gap areas?"

And then what we'll be putting out probably in the next couple of weeks is a recommended standard. So what are those new emerging consumer protection standards that should exist whenever AI is used in financial services? So that's why we landed on this set of axes. They tether to baseline consumer protections, but there are a new set of emerging issues that come up or there are new issues that tether to traditional frameworks.

Alan Kaplinsky:

Right. So there are 12 things that are part of this framework. There actually are 12 things. Could you just give us an example of some things that are in there?

Delicia Hand:

Yeah. So I've had a couple that don't really get a lot of attention because it's sort of exposed fact, but redress mechanisms and consumer agency. There's a lot of focus on bias and fairness, which absolutely matters. But when AI gets it wrong and it does, what can you actually do about it? So I kind of go back to the chatbot doom loop, that's an example of rejust pathway where consumers might get stuck and not be able to connect to a human. And so the burden of correction sometimes when the system fails almost falls entirely on consumers. And on the agency point, 77% of the Americans we surveyed said that they want the right to opt out of AI-driven decisions on major financial matters. What does that option look like today if you're a consumer going online to get a personal loan? And so I think those are the questions that make redress even more complex in this environment, which is we're moving towards systems that operate in near real-time ways.

Those are now going to become agentic systems that may take these same actions on a consumer's behalf. And so then agency matters a whole lot in terms of how you define it as well as the redress mechanisms. So those are just a couple of examples of the axes that we have identified or specifically defined with respect to the use of AI in financial services.

Alan Kaplinsky:

Okay. One of the things I found very fascinating in the report is this word that is used that I had never heard of before called sycophancy. I hope I pronounced it correct. AI systems telling consumers what they want to hear rather than what is really accurate. Why is that especially dangerous in financial services?

Delicia Hand:

Yeah. I mean, so in a non-financial context, someone telling you what you want to hear can be annoying, disappointing, et cetera. But in a financial context, it can cause real harm. ... saying if a AI chatbot is validating a risky investment, then the outcome of that is a consumer will suffer real losses. And when you're dealing with a human advisor, there are equivalent guardrails for providing financial advice. So I think that's one of the things that is a huge risk. And the truth is the fix isn't just instructing the AI not to agree with everything you say, because the systems are often optimized for engagement. And so it's not just a setting that you can flip on and off. The sycophantic nature of these model systems is embedded in how the model was trained, what it's rewarded for. And so that's why we've identified it as a consumer protection problem, not just a product design flaw.

Alan Kaplinsky:

These AI systems are like yes men, right? Is that a good way to put it? That if a consumer says, "Well, I'm thinking of making an investment in ABC company and it really looks like a good investment to me and I like this about it, I like that about it, but it's really a horrible investment and any financial advisor worth their title as being a financial advisor would advise the consumer of that.

AI is going to be inclined to go say, "Yeah, yeah, that sounds like a good idea." And in its written response, will focus more on the positive things than the negative things. Is that really the problem? The AI system wants consumers to like them. If the system is always disagreeing, say, "No, that's not good. That's bad for you." Maybe not that it's human nature not to like something that disagrees with you all the time.

Delicia Hand:

Yeah. Unless you're specifically prompting the AI to not do that, then it is more likely to just do things that will continue your engagement with it. So if you get a bunch of nos, bad ideas, you're probably less likely to continue in that conversation. It's an oversimplification because very likely they exist now and we'll see more of them where these model systems get developed and trained for specific functions like that. The developers will have put in guardrails. I'm not necessarily saying that they will have overcome the sycophantic nature of these systems, but there's awareness of it and efforts to try to address this in model and system design.

Alan Kaplinsky:

Yeah. So let's talk about AI in credit underwriting and lending. It seems like that's an area that it's been going on for quite some time. There are companies that focus on that. So what are the biggest consumer protection concerns that you have and that the report enumerates when AI is involved in credit underwriting and lending?

Delicia Hand:

So AI allows lenders to incorporate a lot more value variables into underwriting decisions. So behavioral data, alternative data sources, patterns that traditional credit scoring would never capture. And so the conventional framing in terms of risk is typically about disparate impact and bias in particular. So where proxy variables are correlated with protected categories such as race or income producing discriminatory outcomes. And so that's a real concern and the legal framework matters. But I think the more interesting question for this moment is those frameworks exist, but we still have the challenge. So have the consumer protection laws actually solved for the original problem they were designed to address? Are more people in homes, for example? And are the communities that ACOA and the Fair Housing Act were designed to protect actually seeing the access that was promised? And there's certainly risk there, which is the data that you put in, if it's not complete, if it's bias, so many different things in terms of the data.

Garbage in, garbage out. And so you've seen that consistently, whether you're using machine learning models or Generative AI. But we do have an opportunity to think about, is there an opportunity to actually focusing on improving the outcomes in ways that benefit consumers rather than just sorting people into categories faster? The risk is that we will get the same things wrong at scale. The opportunity is maybe we get genuinely personalized and fair outcomes for people who were previously locked out regardless of what category they fit in.

Alan Kaplinsky:

Yeah. And I would guess that with the finalization of the CFPB rule not so long ago, essentially eliminating the use of disparate impact analysis under ECOA and Reg B, that exacerbates the problem. Because I would think very unlikely that AI underwriting would be guilty of overt intentional, I should say, discrimination against a covered class. It's more likely that if you use disparate impact, it might have a disparate impact or effect on a covered class. Would you agree?

Delicia Hand:

I mean, yes. With the CFPB finalizing their rule, really you're looking at the effects tests is now the basis for ECOA enforcement. And so that is clearly a sea change. And what that means, I think practically with respect to the AI models, you move from a situation where no one intended to discriminate, but the model produces racially or economic skewed results through proxy variables. That's been removed, at least the CFPB level. So I think it remains to be seen, frankly, how others like state and the courts react to that. But I think in terms of looking backwards, so clearly that framework is now significantly weakened, but I still think it makes sense to ask the question, have our existing consumer protection laws actually delivered on their original promise? Are more people in homes? Period. And can we, through these new and emerging tools, eliminate significantly reduce risk, achieve personalization that actually benefits all consumers so that we have people who were previously locked out of lending entirely or folks with thin files or irregular income? Can AI change that? So there's risk and there's also potential opportunity.

Alan Kaplinsky:

Yeah. Yep. Let's move to another topic that seems to be getting a lot of attention these days. And what I'm referring to is algorithmic pricing, or sometimes it's referred to as surveillance pricing. And in your report, you talk about that as being a case study in market structure risks. Maryland just enacted, I think, the first of its kind statute regulating algorithmic or surveillance pricing. And in fact, I'm going to be soon doing a podcast show just on that topic. But are we moving toward a world where every consumer effectively receives a personalized price? I mean, is that what's really what you're fearful of?

Delicia Hand:

I think we're already partway there. So alongside this landscape analysis and developing frameworks for how AI is showing up and impacting consumers, we've also on our investigative journalism side, been looking at pricing specifically in the general consumer retail space. So we, over the last year, for example, we've looked into the grocery space and looked at things like Instacart or Kroger. And so personalized pricing is already there.

People are being shown different prices for the same items, which is the effective finding of our Instacart investigation last year, which is if you're using Instacart, at the time, Instacart has since then updated its practices, but you were definitely being charged different prices based on where you lived.

Also, our reporting on Kroger found that there's definitely a difference between what's listed on the shelves and then what you ultimately received at checkout, which is tied to the use of automated pricing system.

Alan Kaplinsky:

Are you saying, let's make sure I understand it, Delicia. I thought this was more of an online phenomenon, but you're saying this is even going on in stores where you'll go into a store and you'll see a price, say in a grocery store, a price for a certain item, but when you get to the checkout counter, somehow you're charged something else?

Delicia Hand:

Absolutely. And our two recent studies showed this. So look, there's always been price differentiation, whether you're using coupons or loyalty and rewards programs. Now, algorithmic pricing has just automated some of that, but also made that a whole lot smarter. So the rewards that you might see and receive and that might pop up in your app such that you see something on the shelf, but then you go to checkout, you may have a whole set of different rewards that are available to you based on where you live, your prior shopping patterns, et cetera, than I might see. And so you can't really observe it. You can't trust necessarily the price that is on the shelf. It's all going to be increasingly personalized based on-

Alan Kaplinsky:

Is it usually higher than what you've seen on the shelf or could it be higher or lower?

Delicia Hand:

100%. Our Instacart investigation, for example, showed that there are instances where consumers are paying up to 23% more for a product just based on having used the app compared to other consumers. So now Instacart has now committed to improving its practices, but I think that the key issue is that you can't observe pricing anymore in the same way that you used to, or you can't necessarily compare or meaningfully contest it when it's all digitized. So it's not necessarily just in online spaces. Now put that into the context of price discrimination and financial services. This may mean that you're charged something different for insurance based on your personal characteristics, how fast you drive, how long you've had insurance, et cetera. Or you might get-

Alan Kaplinsky:

How gullible you might be. Is that something else that could be that the seller may know?

Delicia Hand:

There are a host of different criteria which are now being used for determining pricing. And we're not necessarily advocating for uniform pricing, but fair pricing. So in theory, we want to ensure that personalization, there's transparency about it and there's actual oversight. And so that you're not necessarily moving to a system that reflects maximum willingness to pay, but some transparency about the criteria and the way in which prices that consumers see and are offered. What's that about and what the differentiation is based on? And that's what you're seeing in some of these algorithmic pricing legislation that is moving through various states, Maryland being one of the most recent ones.

Alan Kaplinsky:

Yeah. One of the things I've noted, I go to a lot of sporting events, I go to a lot of theater, and so I'll go online. And this concept of dynamic pricing has been with us for a long time now where the price of a ticket, let's say to see a Broadway show, can vary from hour to hour for the same seat.

Another thing I've seen is a lot of these ticket companies won't show you every seat that is actually available. They'll make it look like there aren't many seats left. And that tends to make people want to pay a higher price. And I've seen that because I'll go on one day and it'll look like, oh, there's only a handful of tickets left. If I want to see that show, I better jump on it. But then I might wait a few days, go on again. And all of a sudden, mysteriously, there are many, many tickets available and they're a lot cheaper.

Delicia Hand:

Yeah. I mean, I think if fundamentally the logic comes down to whether or not it's fair to use someone's personal data against them to charge them more. And so that's kind of what's at, I think, the heart of these practices that you're seeing, but also some of the legislative and policy responses as a result.

I had a similar experience the other day where I was on the train leaving the DC area trying to book some airline tickets. And I started the search in DC, prices were maybe like \$500. And then by the time I got to New York, but not Manhattan, so while I was still Manhattan, same price, but then went out to Yonkers, which is our headquarters. Refreshed the search. The prices were \$200 lower. So do I have to go to Yonkers every time I want lower airline flights? So I think it's time to look at the logic of dynamic pricing when you now have more data that is available and what is exactly fair for a consumer when personalization is involved?

Alan Kaplinsky:

Yeah. Well, and also related to it is, I mean, with dynamic pricing, the concept is supposed to be it's just a reflection of supply and demand. And that varies from day to day or hour to hour. But yet demand can be manipulated in a way by giving the consumer the impression that there aren't many seats left on the airplane or there aren't many seats left to see the NBA finals or the Super Bowl. I mean, there's just a lot of things and AI probably tells these companies the best way to maximize revenue. That is just my guess. I mean, I don't know. I don't understand the technology, but it's a black box, I think.

Delicia Hand:

Yeah. The black box is really what is in focus here, which is I think traditional consumer protection law tethers explainability to... So consumers now have rights to know, or not now, but have had rights, denied credit to get an explanation. With dynamic pricing and the use of more and more robust AI systems, but also wider and deeper data streams to determine pricing. Is explainability A, possible? Most people will say, "No." And B, so then what? What is the actual way to get to a solution of some degree of transparency, but also some degree of actionability for the consumer.

Alan Kaplinsky:

Let's turn to another subject. Oh, there's so much to talk about here, but in so little time. AI as a financial advisor. Increasingly, consumers use ChatGPT, Anthropic's Quad, Grok and other general purpose AI tools for getting financial advice. There might be particular stock that somebody's interested in purchasing and they want to find out what do you think about this? As if you're talking to a human being and they do give you financial advice. I mean, I've seen it. I've played with it myself. And sometimes what they'll do, they'll give a more balanced presentation of what you've asked them. But other times it's sort of the sycophancy. Have I got that right?

Delicia Hand:

Sickness.

Alan Kaplinsky:

Yeah. That gets triggered. But there's clearly financial advice. And last time I looked, these entities, AI systems were not licensed as financial advisors or brokers or dealers. How do they get around? How does that happen?

Delicia Hand:

I mean, I think this is going back to the asymmetrical nature of technology being out, deployed, in use, and regulation follows. So it happens because that's the cycle that we've been in, whether you're talking about a consumer using the very first check to ATM to now Claude or ChatGPT to ask for financial advice. And you've mentioned a lot of things.

So first, these systems are being used by everyday people. I ask my husband at least once a week, "Has Claude ever owned a credit card?" Why are you relying on something that has no real actual human experience? It is roughly leveraging all of human knowledge that's on the internet unless you give it more. But there is fundamentally an overconfidence problem when you're using these general purpose tools for financial advice. These systems can sound extremely authoritative while just being incomplete or wrong. And there's also a more subtle issue, which is these general purpose tools were not designed to contextualize financial information for your specific situation. The other thing that I often say to my-

Alan Kaplinsky:

In other words, you're saying they don't know enough about you.

Delicia Hand:

I mean, they may not even know what day of the week it is and the month and year. And at the same time, we are engaging these tools with our most sensitive and important information. And so they can surface information without knowing all of your situation, your tax situation, your risk tolerance, your liabilities, your timeline.

If I'm giving a general purpose tool, just a narrow context, it's going to be responsive to what it's given. And so the response may be technically accurate, but in the abstract, that can actually be harmful when applied to your specific circumstances. Have you given it enough information to give you an answer that then you as a consumer can confidently rely upon. When we make financial decisions I think we often don't just make them without the context that our situations bring. I buy something based on what my own set of affordability is that determines on when I get paid, the liabilities that I have, et cetera. And we may not be putting that context and additional information into the chat tools.

So that is, I think, the primary issue of people using these tools for financial advice. The other is accountability. Is ChatGPT or Claude, we're in changing spaces so frankly, there's a lot of partnerships and things like that being brokered now. But at the end of the day, I don't think any of these companies are registered financial advisors. They're not subject to the same suitability and other fiduciary duty standards that exist in financial services.

Alan Kaplinsky:

Yeah. By the way, it reminds me of another podcast I recently recorded. American Law Institute is developing a set of principles right now related to liability, tort liability of the various players involved in offering AI products. Is it the developer? Is it one of the middlemen that's involved? Is it the buyer of these AI system that made some changes to it? And it's right now a very murky area. There hasn't been a lot of case law yet, but it's surely going to. The courts are going to be absolutely inundated with liability issues.

Particularly, I'm thinking if these AI systems are not licensed as financial advisors and they're giving advice on buy this stock and it turns to be a very bad piece of advice and somebody loses a lot of money, what's the liability? Can you go after? Who do you go after? Who's responsible? How do you identify it?

But anyway, I just wanted to mention that ALI project. It's very much in its early stages. It's not yet anywhere near completion, but it's being. The reporter is Mark Geistfeld at NYU Law School. And he was on my program not so long ago. It's very interesting to hear about the work that they're doing. Let me turn to privacy and data aggregation. That is, and I guess what we're really getting into the area of open banking and CFPB rule 1033, when that gets finalized, if ever. The report highlights how AI systems are increasingly connected to real financial account data through tools like Plaid. How does that change the risk profile?

Delicia Hand:

Yeah. I mean, when you link your bank account or any account to an AI tool, you're often creating data flows that bypass the privacy and security protections that are built into your existing financial institution and their requirements to maintain this. And then one of the things our research showed was that just looking at the terms of service for many of these tools, they explicitly carve themselves out and reserve the right to use transaction data for model trading, sharing it with third parties, retaining some of that data indefinitely.

So there's a whole host of issues raised, be they privacy, security. And so consumers don't look at the terms of service every time they're using a lot of these tools and don't fully understand what they've consented to, partly because I think some of the consent mechanisms are being developed. And the example you just used, this is a space where many different sectors are now considering and looking at what are the specific implications for use of these tools?

So the financial institution holding your account is subject to GLBA, for example, CFPB oversight, examination, but the AI tool reading that same data is not. And so that creates a real vulnerability. And again, our current regulatory framework, there's a gap there.

Alan Kaplinsky:

Yeah, sure. And that sort of brings me to what I'm going to call the penultimate question I have for you. I'm going to have one more after this. And I just touched on this, is the issue of accountability and redress. You emphasize that consumers lack meaningful ways to challenge harmful AI decisions. In your view, what should an effective redress mechanism look like?

Delicia Hand:

I think at a minimum at this stage, a clear right to know when AI was involved in a decision affecting you as a consumer, a meaningful explanation of why, which seems simple, but really hard to achieve when AI is involved or Generative AI. Access to human review. Now that we're now in this world where models, model systems are being embedded in everything, we've moved from a system designed to keep the bots out to now a system that's going to be run, mediated by automation. So that makes access to a human, I think, an even more valuable proposition. So something like a right to human engagement and review. And then a real dispute mechanism with defined timelines. If money can move instantly, it can instantly be in the

wrong hands, then the dispute mechanisms and timelines matter even more with these robust systems that could accidentally facilitate fraud, scams, et cetera.

Now most of that doesn't exist. So I think it comes back to can the consumers identify that AI was involved? Can they get an explanation they can actually use? And then can they meaningfully challenge the outcome?

Alan Kaplinsky:

Right, right. Okay. My final question, this is going to be a hard choice for you, I'm sure, but what do you see as the biggest regulatory gaps today?

Delicia Hand:

Yeah, that is a hard choice because there are a lot, but I mean, I think there are three areas that stand out. I've just started to dig in today to yesterday's presidential executive order, which I think makes this conversation even more timely.

First, there's no coherent federal framework specifically for the use of AI in financial services. We have a lot of technology neutral statutes that were developed in a very different moment. And frankly, they're doing work that they weren't designed to do. We have statutes that literally refer to answering machines. We are way beyond that moment. So we do need strong federal frameworks. And then I'll go back to something I've touched upon, which is there's no meaningful requirement for explainability in consequential decisions. And I'll be transparent, which is I don't know that explainability is the right solve. I'm more and more inclined to say, "Well, who cares why? Most consumers want to know what to do next." So I'm denied. So then where do I go?

So increasingly, I think we should actually be focused more on outcomes rather than what's the reason. Everybody has elbows. I don't need to know why it was denied. I just need to know where to go next. But explainability I think is still important if it helps facilitate a different outcome or a different pathway for a consumer. So I think there's still value there.

And then I think we have to look at this whole new business model that has emerged the last several years and think about what accountability in the AI supply chain means. Right now, it's not really existent. So when an AI system causes harm, you said this yourself, who's accountable? Is it the model developer? Is it the cloud provider? The financial institutions? We thought we had issues with bank fintech partnerships. What's that going to look like with bank fintech AI developer partnerships?

So I think yesterday's executive order directing regulators to review and think about barriers for fintech firms, the framing is around removing burdensome oversight. But in this moment where I think even more nuance and complexity is being added to the system, I actually think we need the opposite. I'm very mindful that over-regulation doesn't necessarily produce better outcomes, but I don't think this is the moment to say, "If you're using technology, we need less guardrails." I think we need to define those things now.

Alan Kaplinsky:

Yeah. I said this was going to be the last question, but as you've been talking, there are two other things I want to briefly get your reaction to. And that is what should industry be doing right now? If you were advising financial, when I talk about industry, I'm talking about financial institutions. What practical steps should they be taking regarding AI governance?

Delicia Hand:

Yeah. So definitely first, know what your AI is actually doing. It is not a good idea. And I've sat in rooms where financial institution execs have talked about using off-the-shelf box models. I don't think that's going to be enough because accountability is going to be defined. So I think they need to conduct genuine internal audits and also frankly, stop treating AI integration as this PR signal or race. There are real consequences to using these models, and so you should know what your AI is doing.

Second, build real redress mechanisms now before you're required to. I'll stop there. Better build it now than later when you're doing it at cost to both the consumer and reputation risk. And then third, I think in the same way that we found it quite

valuable and necessary to think about, well, what does the use of AI mean for how we evaluate consumer outcomes and impacts?

You need to integrate separate AI performance metrics from consumer outcome metrics. They're not the same. And I think right now, most institutions that are integrating AI are looking just at AI performance metrics. And when consumers are using them, it seems like a missed opportunity that might have big consequences in the long run.

Alan Kaplinsky:

Right, right. So this truly is, next one, my final question. And that is, are you optimistic or are you pessimistic about where we're headed with AI? What's it going to look like in three years, five years? What's your feeling?

Delicia Hand:

Yeah. I mean, I think ultimately I have to be optimistic. Risk is all around us. Financial services is inherently risky. So I wouldn't be in this space if I was completely risk averse. And in the same way where there can be high risk, there can be high benefit. Fintech has made loads of different services and products available to consumers where they frankly weren't. I think there's been net positive to access in different ways of banking and engaging with your finances that had never existed before. I often talk about my auto mechanic cousin who lives in the Bronx, New York, who, because of apps like Stash and Robinhood, is now a very sophisticated investor and has a portfolio that I think really is beyond what anyone would have expected of someone in that place. So from that perspective, I'm optimistic. And I think we have an opportunity to really think about what's possible with these tools.

So not just a sort of gird against risk, that's important. Guardrails and the articulation of them are important. But how do we lean into this moment to tackle some of the things that the consumer protection laws that we have, have not fully achieved? So to the extent that we desire to have equal access to credit and all of these things that have not yet been fully realized, is there a way to achieve that through intelligence? And so from that perspective, I'm optimistic.

Alan Kaplinsky:

Yeah. So, Delicia, this has been an extraordinarily thoughtful and important discussion. A few major takeaways stand out for me. First, AI is no longer theoretical in consumer finance. It's already embedded in lending, fraud detection, pricing, customer service, and financial guidance. Second, your Consumer Reports report makes clear that many of the incentives that are driving AI adoption aren't necessarily aligned with consumer interests.

The market is prioritizing speed, automation, scalability, and competitive positioning while consumers are prioritizing fairness, security, affordability, transparency, and human accountability. Third, the regulatory framework remains incomplete. Important issues including explainability, redress mechanisms, privacy protections, and accountability across the AI supply chain remain unresolved at this point.

Fourth, the discussion highlights that AI risk is not confined to discrimination alone. It also includes manipulation, over-reliance, opacity, concentration of data, and the possibility that consumers increasingly lose visibility into how financial decisions affecting them are actually made.

And finally, your report provides an important roadmap for a consumer-oriented framework for AI governance. One that's focused not simply on technological capability, but on what AI systems actually owe consumers. This conversation underscores that AI in financial services is not just a technology issue. It is fundamentally a consumer protection issue, a market structure issue, and ultimately a trust issue.

So, Delicia, thank you again for joining us and for the tremendous work that you've done that is reflected in the report. And I do hope that when this net other report is coming out in a few weeks, that you send it to us because it may be worthwhile for us to revisit and to have a discussion about what's in there.

To our listeners who can't get enough content on AI and how it impacts consumer finance, we want you to know that during 2025 and 2026, our podcast series has continued to focus on the accelerating use of AI in consumer finance and the regulatory and policy responses that it's prompted.

Several episodes have explored in depth two important White House reports dealing with AI, including their treatment of AI risk management, safety and governance frameworks, as well as their implications for consumer protection, accountability, and responsible innovation in financial services.

In addition to the podcast that we're releasing today, we've examined how these federal policy developments are beginning to influence both regulatory expectations and industry practices, particularly around model oversight, transparency, and the use of automated decisioning systems that are used in credit, marketing, fraud prevention, and other core consumer finance functions. Taking these together, these discussions underscore both the transformative potential of AI in improving financial services, but also the growing need for clear, effective safeguards, or as you put it, guardrails to ensure these technologies are employed in a fair, transparent, and accountable manner.

To locate our AI resources that I've mentioned, simply go on our blog at consumerfinancemonitor.com and do a search under artificial intelligence. We want you to consider our blog, our podcast shows, our webinars to be the place where you can find the answers to your AI questions. You can view it as a place to do one-stop shopping for AI resources that relate to consumer financial use cases. So with that, Delicia, my thanks to you for coming on our program today and sharing with our listeners what's in that report that you put together. You're really doing good work. So my thanks to you.

Delicia Hand:

Thank you so much. Great to be here.

Alan Kaplinsky:

And I want to thank our listeners obviously for taking the time to download our show today and to spend their time with us. And with that, I hope everybody enjoys the remainder of their day.