

Consumer Finance Monitor Podcast (Season 9, Episode 9): Agentic AI in Consumer Financial Services: Opportunities, Risks, and Emerging Legal Frameworks

Speakers: Alan Kaplinsky, Joseph Schuster, Adam Maarec and Oren Bar-Gill

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 former practice group leader for 25 years and now senior counsel of the Consumer Financial Services Group at Ballard Spahr. And I'm very pleased to be moderating today's program.

For those of you who want even more information, either about the topic that we're going to be talking about today or anything else in the world of consumer financial services, please don't forget about our blog, consumerfinancemonitor.com. We've hosted our blog since 2011 when the CFPP became operational. 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 Apple Podcasts, YouTube, Spotify, or wherever you access your podcasts. And also, please let us know if you have ideas for any other topic that we should consider covering on a future podcast show, or speakers that we should consider as guests on our show.

And I'm very pleased to be hosting today's repurposed webinar on a subject that is very timely and it's going to focus on what at least I consider, many others consider to be one of the most consequential shifts in our industry. That is the emergence of artificial intelligence in consumer and financial services, and the ensuing legal compliance and regulatory issues that it raises.

As many of you know, AI has already transformed how financial institutions operate, from underwriting and credit decisions to fraud detection and customer engagement. But agentic AI, AI systems that can act autonomously, make decisions and interact directly with consumers, brings a new set of opportunities and a new set of risks that stretch beyond traditional automation and even generative AI.

Today's discussion will bring together three leaders who are deeply immersed in these issues across academia and in private practice. First of all, we are very honored and pleased to be joined by Professor Oren Bar-Gill, who is the John Edward Sexton Professor of Law and Economics at NYU School of Law.

Professor Bar-Gill is a highly respected scholar whose work explores algorithmic decision making, consumer harm, and the intersection of technology and the law. He will lead off our discussion by framing the discussion we'll have for the remainder of the webinar, and will help us understand what agentic AI is, how it differs from other forms of automation, and why it matters for legal doctrine and consumer protection.

He will draw on his scholarship and algorithmic differentiation and consumer harm, which includes a book that he published last year, co-authored with Cass Sunstein called *Algorithmic Harm: Protecting People in the Age of Artificial Intelligence*.

He will explore whether agentic AI represents fundamentally a new legal challenge or an intensification of existing doctrinal stresses that deal with such matters as consent, disclosure, accountability, and market power. So Oren, a very warm welcome to you. We're absolutely delighted that you've joined our program today.

Oren Bar-Gill:

Thank you very much, Alan. It's a real pleasure to be here with you.

Alan Kaplinsky:

Yeah. And I should add that not so long ago, but I guess about three years ago, I got to know Oren pretty well because he was one of the three reporters on the ALI project to create a restatement on consumer contracts. A 10-year project, I was on the board of advisors. It ultimately resulted in approval by ALI. And I know from having talked to Oren, it is now being cited regularly by the courts. So that, of course, is a major achievement. And knowing that Oren is now focusing some of his scholarship on AI, I think is terrific. It will be something that's going to help all of us, for sure.

Let me introduce my colleagues that are going to be joining us after Oren is finished framing the issues. We're joined by Joseph Schuster, who's a partner at our firm. Joseph, who worked for several years at Goldman Sachs and rejoined our firm about a year and a half ago, brings deep experience advising clients on AI, compliance, and regulatory strategy, particularly in the consumer financial services space. His practical perspective is invaluable in illuminating how institutions are thinking about risk and governance as they explore agentic AI use cases.

And finally, last but certainly not least, I'm very happy to be joined again by Adam Maarec, who rejoined us at Ballard Spahr a few months ago. Adam's work focuses on emerging technology and regulatory issues in financial services, mostly in the consumer finance area, and he'll help us ground today's conversation in the realities of the evolving legal expectations, from fair lending and discrimination risk to consumer disclosure and enforcement dynamics as regulators catch up with the technology. So let me turn the program over to Oren. The floor is yours.

Oren Bar-Gill:

Thank you very much, Alan. As Alan mentioned, I'm going to kind of focus my comments here on this recently published book with Cass Sunstein, and I'm going to try to link the analysis in the book more directly to the subject matter of this webinar.

So in particular, I want to discuss the potential harms to consumers from artificial intelligence, including agentic AI. And this links to, or ties with the subject matter of this webinar and specifically to the legal risks of using agentic AI in consumer finance, because these are the type of harms that create legal risks or potential legal risk for providers of financial services. I will also, at the end, share a few preliminary thoughts on kind of the question of ascent and agentic AI.

The framework of analysis that we use in the book to identify the potential harms from AI in consumer markets is a very simple one, which is captured in this little diagram. The first thing that we do is identify or distinguish between markets where most consumers are sophisticated and markets where most consumers are unsophisticated.

And when I use the word sophisticated, I'm referring to kind of information and rationality capabilities of these consumers. So sophisticated consumers are usually able to make optimal choices for themselves, and unsophisticated consumers often make mistakes.

For each one of these markets, for what we call the S markets, with sophisticated consumers and U markets with unsophisticated consumers, we compare a kind of pre-algorithmic, pre-AI benchmark with a lesser degree of differentiation, to a post-AI world where we think there is a greater degree of differentiation. So I'm talking about more personalized, individualized pricing and targeting of ads, marketing, and product offers.

And in order not to hide the ball, I will already kind of mention that one of our main bottom lines is that in sophisticated markets, in these S markets, the use of AI and the extra differentiation that it makes possible is usually beneficial. It's a good thing. On the other hand, in the U markets, it might actually be harmful.

Tying this to financial services, I want to emphasize that when I'm talking about sophisticated versus unsophisticated consumers, this is not necessarily about how smart the consumer is. It is mostly about the likelihood that this consumer will make a mistake in terms of what products to choose or how to use these products.

And this, of course, suggests that the complexity of the product or service is critical in identifying the market as an S market or a U market. And as we know, there are many financial services markets and products, from mortgages and credit cards and beyond, where the product is relatively complex, which can lead to potential mistakes.

I want to now focus on two types of harms. This is in the book, we talk about several others, but these, perhaps, are the main ones. We're going to talk about algorithmic price discrimination. This is when the algorithm and this AI powered algorithm is

actually used to set prices, and this can be done with various technologies, including with agentic AI. And when the algorithm is used in order to target specific advertisements or product offers to individual consumers.

I want to start with just a few examples, how this type of individualization powered by AI is increasingly happening. I'm sure many of you are aware of the Delta debacle from the recent summer, where the president of Delta was recorded on an investor call saying that Delta is making a lot of money by using AI to personalize prices. This created a big uproar, and then, of course, the PR folks from Delta deny that they ever do this.

I also want to mention this company called Fetcher.io, which its product is helping companies use generative AI in order to set prices. And this is just another indication that these are markets and practices that are increasingly prevalent. Also from a few months ago, the Instacart story was another example where research has found that different individuals are charged different prices for the exact same product at the exact same store.

Moving closer to financial services. In the insurance world, there was an important class settlement approved in 2024 in California, and the claim in this class action litigation was that the insurance company, in this case Allstate Insurance, was using AI in order to set premiums as close as possible to the willingness to pay of each different consumer.

Also in this space, we have another startup company called earnix.com. Again, their business model is to help other smaller companies do price optimization with the help of AI in the insurance and banking sectors. So the increasing use of AI in pricing has become concerning to many policymakers in the US and in other countries, and in the US at the federal level, although recently less so since the change of the administration.

But at the state level, we see a lot of interest in recent legislation about these issues, about the increasing use of AI in consumer markets. Normatively, I want to emphasize, as I've already suggested at the framework section, that using AI and creating increasing differentiation is not necessarily a bad thing.

And so, one of the main things that we want to do, and we try to do in our book, is to identify the conditions under which the AI-enabled differentiation is good for consumers and the conditions under which it is harmful. And of course, when it is harmful, it is more likely to lead to legal and regulatory risks for financial services providers.

And so in order to identify these two situations, when it's good and when it's bad, I want to go back to this distinction that I mentioned between the S markets and the U markets. In S markets where consumers are relatively sophisticated, they understand what's going on, adding differentiation and personalization with the help of AI is generally thought to be a good thing for consumers and for general efficiency.

In order to illustrate the point, I want you to think about a simple monopoly market. The analysis is not limited to monopoly, but let's start with that. In a simple monopoly market, if the monopolist has to set a single monopoly price, we know that this is problematic and inefficient because there are going to be a whole bunch of consumers at the bottom end of the demand curve who are going to be priced out of the market.

And these are transactions that efficiently should happen and do not happen. When AI allows for price differentiation and personalization, what's going to happen is that the seller firm is going to set different prices for different consumers, and those consumers who were priced out before now are able to enter the market and they will pay a lower price corresponding to their lower willingness to pay.

This creates efficiency gains in the market, and it's also good for these consumers at the bottom end of the demand curve, which are usually also poor consumers. So, this is also both efficient and distributionally attractive. The story is very different when we move on to the U markets. Recall in U markets, consumers are making mistakes, and I want to focus here on mistakes of the overestimation type.

So I'm looking at situations where consumers overestimate the value or the benefit that they are going to receive from a particular product or service. What will happen in these types of markets is that the AI pricing agent is going to set a price that corresponds to the overestimated valuation.

So for example, if I have a value, a true value of 100 from a certain product or service, but I mistakenly think that the value is 120, the AI will price perhaps at 115, and note that this means that I'm going to be losing money. I'm going to be paying 115 for something that's really worth 100 to me.

And so this is kind of an important category of cases in our minds, where the use of AI and the extra differentiation is going to be harmful to consumers. And the reason, again, is because each one of these consumers on the demand curve, with their own mistake and overestimation, are going to be paying more than what the product or service is actually worth to them.

Since our focus here is on financial services, I want to emphasize that what I've just said about overestimation of benefits carries over in exactly a similar fashion to underestimation of price. So when we're talking about financial products, like mortgages and credit cards that have often complex multidimensional pricing systems, okay, for these types of products, a main concern is that the consumer will underestimate the overall price of the product or the price that I will, as a consumer, pay over time. And this leads to the same type of issues when extra differentiation with AI happens.

I want to distinguish these potential risks that I've mentioned. When the AI is used to set prices just below the consumer's willingness to pay, I want to distinguish this from another very common situation, especially in financial product markets where the AI differentiation tracks the risk, the individualized risk of different consumers.

And I want to suggest that we should be much less worried about differentiation that ties price to individual risk as compared to differentiation that ties price to the willingness to pay. And as I've mentioned in the Allstate case, in that class action from California, it seemed completely plausible in an insurance market that there will be differential pricing based on risk. The problem was when AI started to be used in order to set prices, individual prices based on the different consumer's willingness to pay.

I want to end my presentation with some kind of preliminary thoughts on assent and agentic AI. And I just want to describe for you a relatively recent case in order to use as an expository tool for these thoughts. And so this is about chatbot contracting, and I think this is something that will increasingly be done by agentic AI. Perhaps it already is done by agentic AI.

And so let me give you the facts of this case for those who are unfamiliar. Jake's grandmother died. Jake is Canadian, so he visits Air Canada's website to book a flight from Vancouver to Toronto in order to attend the funeral. Jake is unsure how Air Canada's bereavement rates work so he asks the chatbot on Air Canada's site to explain this to him.

And the chatbot writes, "If you need to travel immediately or have already traveled and would like to submit your ticket for reduced bereavement rates, kindly do so within 90 days of the date your ticket was issued by completing our ticket refund application form."

Now, it turns out that the chatbot's description answer here was not reflective of Air Canada's actual bereavement policy, which explicitly stated that the airline will not provide refunds for bereavement travel after the flight is booked. You actually have to report and ask for the refund before the flight is booked.

In fact, an earlier response by the chatbot indeed linked to that webpage with the bereavement policy. So Jake, without reading the policy, just relying on the chatbot's response, he books a flight and later submits a ticket refund application form exactly as the chatbot suggested. Air Canada refuses to issue a refund, citing their bereavement travel policy.

Is Jack entitled to a refund? The court in Canada says, "Absolutely yes." And just citing briefly from the opinion, "Air Canada argues it cannot be held liable for information provided by one of its agents, including a chatbot. It does not explain why it believes that is the case," the court says, and of course the court, as you can see, is not persuaded by this explanation.

In my mind, the court here got it exactly right, and the same principle that we see here should probably be applied to agentic AI in this and other cases. As the court in Canada ruled, a principal is responsible for the action of its agents, and it seems likely that agents here will include not only individual human agents, but also agentic AI. With this, I will thank you for allowing me to participate here and return it back to Alan.

Alan Kaplinsky:

Okay. Well, Oren, thank you very much. And I think you did just an excellent job in framing these issues for now. And thank you for staying around and joining us for the remainder of the webinar, and we absolutely invite you to comment on what you're going to hear from Joseph and Adam.

And so, we're now going to turn to practical legal and compliance issues. We're going to be talking about use cases, both things that are happening right now in the consumer finance world, and those things that will likely happen in the future, that are on the drawing board, and things are moving at a very rapid pace.

And I'm sure if we were to do this webinar again in another six months, there will be things going on that we haven't even thought of. But anyway, I want to turn now to take the lead, and Joseph to then respond to what Adam's going to have to say. And we're going to talk about first marketing and underwriting. And Adam, you'll take the lead on that and then we'll go to fraud and servicing and that'll be you, Joseph, taking the lead. And then compliance monitoring and operations, then we'll go back to Adam. So Adam, floor is yours.

Adam Maarec:

All right, great. Thanks, Alan. Thanks, Oren. That was a great presentation. We wanted to highlight just a couple of the ways that financial services companies are using or thinking about using AI today. And these are some of the most prevalent that folks here have probably heard about already, but we'll go over them and talk about some of the legal implications and risks.

And so, first is with thinking about the customer life cycle, it's with marketing. AI is being used to help crunch data and figure out different segments of the population so that you can provide different personalized offers.

And if you've played with ChatGPT, you know how quickly these LLMs can spit out really good content that's focused and tailored to the younger consumer. With an advertisement for a new card, it might say, "Hey, bro, you could earn great rewards on this with this new credit card." And if you're reaching out to a different segment, maybe my mother, it might be, "Dear Mrs. Maarec, we have a new offer that you might be interested in."

And so that different kind of tone and personalization is able to be developed, and the volume of different varieties of marketing content that can be created is just becoming massive. And that creates some legal and compliance risks for us to think about as those different kinds of marketing content are being put out there.

Underwriting is another area we hear a lot about, underwriting and assessing risk, as Oren mentioned. This is using these tools to gauge actual risk and help determine pricing and availability for credit and credit lines.

Companies are trying to gather as much data as possible, and to analyze it, to find causal relationships between the data to beat out the competition and either bring more folks into the mix that can have credit or have differentiated parameters for that credit.

And developing new ways of thinking about credit underwriting than we have in the past, instead of just relying on credit scores and trade lines, looking at cashflow data for indications of how consumers' financial lives are actually operating, and being able to find that predictive data element.

I'll just mention that in terms of all of these use cases we're going to talk about, and I'll turn to Joseph in a second to keep going here, but all of these use cases rely on having lots of data. And there is an inherent challenge in terms of protecting consumers' privacy and only collecting from a data minimization perspective.

Well, collecting the data that you need and recognizing that there's some risk and holding more data than you need, because at some point it has declining value, but can still be a target for fraudsters. There's this tension between amassing lots of data in order to power these models, and also limiting the amount of data that you need to hold in order to deliver your products and services.

So, there's an undercurrent of privacy to all of these use cases that we're talking about, and legal and compliance professionals need to be thinking about. So Joseph, with that, I'll turn it to you to talk about fraud.

Joseph Schuster:

Thanks, Adam. And Oren, it was a great conversation. I'll say, I think that a lot of what Adam and I are going to talk about is how what we're seeing in financial services is that financial institutions are really using AI in a way that is beneficial for consumers here.

As regulators are asking these questions about how you're using AI and different things, which we'll get into, there's an expectation that AI is being used, because it is beneficial for the consumer, it is beneficial for the institution, and it's beneficial for the entire payments network and financial network that we have.

And a couple of those where there are strong areas, I know Adam mentioned underwriting. Oren had said that we should be less concerned about prices relative to individual risk. This is a great opportunity that a lot of institutions are looking at in terms of, "How do we offer credit or financing products to individuals who maybe we weren't able to before because we had the data, but we just weren't looking at the different elements of that data relating to that person?"

It was difficult to figure out whether or not somebody was credit worthy. That person may not have been offered credit. Now we're able to offer more people credit by doing these things. So, it's an amazing ability to expand on what can be done in financial services.

Let me touch on fraud quickly. Fraud is one where it continues to escalate. There is more fraud and payments in the financial industry today than yesterday, tomorrow there'll be even more. It continues to expand. AI is being used in a number of different ways, and it's incredible what it can do in terms of both number one, fraud detection, and two, evaluating claims of fraud.

In the fraud detection space, fraud has always been a pattern recognition type of activity that you're looking at, "Does the person live within this geography? Is this a high risk merchant?" If you're in a foreign country trying to pay outside of the pump, you might be declined, not because you should be or that you think that there's like an actual thought that, "Hey, what Adam is doing is fraud."

It's more just that this is a risky behavior. Now, you can look at things to say, "Is there a pattern of risk objectively?" So identifying more of those areas where it's like, "All right, this is a risky thing, we should look at that." But then is it outside of this consumer's behavior? So if there are a bunch of other purchases in this area, like you have the hotel, the rental car, all in this area, it makes sense that you're getting fuel at a pump at that point.

So you have less friction for consumers in that space, but you have more protection for financial institutions, and the consumers as well from preventing fraud transactions where it's like, this doesn't make any sense that this transaction is happening. And being able to then very clearly clear that with the person if there is a flag, from the standpoint of you can ping that person right away to say, "Hey, is this you doing this transaction?" They can then respond.

And this is where we get into the servicing collections, where you may have had to in the past wait for an agent to get on the phone. Now, with agentic AI, you can immediately talk to a chatbot to say, "This was me, this was not me." If it wasn't you, what other transactions did you not make to get that case all created right away?

And this is, again, in fraud where agentic AI can be used to then gather all of those materials, maybe help put together some of the evaluation for the materials that an agent will look at maybe. In that servicing collection space as well, I think that Oren's example about the chatbot concerns was a good one. When somebody calls into a financial institution and speaks with a person or chats in and receives information, they expect that they can rely on that information. Regulators expect them to be able to rely on that information. That makes sense. We'll get into responsibility later on.

In the collection space, this is where it's getting very interesting, that there are predictions about whether people might go ask you, might become delinquent. And being able to identify what are those stressors? Are there options and programs that can be offered to individuals before they become delinquent? As well as, and this is where it gets really interesting, outbound AI agents, where you will not know that you are talking to an AI. It will sound like a person. And there are different aspects of that that can happen, whether it's a male or a female voice, whether there's a particular accent.

And then there's data on this, too. If a certain gender owes money and a different gender calls that person money, the percentage more money you will get from the calls if you use specific voices in these types of calls. And this can all be done through AI. Now, there are concerns with that we'll get into, but there's a ton of use cases with it. And it's something that regulators are expecting to see and banks are using, and financial institutions, and it makes sense. But I'll turn it back over to you,

Adam Maarec:

Yeah. So, continuing on the same theme there of using lots of data about your customers and your customer experiences, and applying these tools to analyze that data for other backend purposes. So, from a compliance perspective, you no longer have to rely on sampling. These tools can really monitor all of your customer service agent activity at scale on 100% basis when applied properly.

And some of the tools, if you think about transcription services and sentiment monitoring, you can apply these tools to customer service, agent interactions on a full basis in order to detect issues, patterns, and potential problems that require compliance attention. So we're seeing some of the backend of these tools to operations, document processing, HR analytics. And then I do think we're seeing a lot of workflow automation, where things that humans are doing can now be done at scale using LLMs.

And there is a line between what is truly autonomous AI versus the use of technology that is workflow automation that we in the financial services industry have been using for decades. These are just very powerful tools to do a lot of the same things. And so, to keep that framework in mind as we think about the regulatory issues, that this is yet another tool in the toolbox to improve your operations. I think that's an important framework.

Alan Kaplinsky:

Thank you. I do have one question, and I guess Adam or Joseph could answer it. One of our listeners asked whether or not AI or a chatbot could actually make collection calls, where the consumer might not even know they were talking to a robot.

Joseph Schuster:

Yeah, I'll take that one. That is a use case that we are seeing, both there are financial institutions that are building their own internal outbound collections, AI programs, and there are a number of different entities that are creating these programs as well. They're fascinating programs. I've had the opportunity to interact with some, and you cannot tell the difference with respect to speaking to an agent.

A number of collections calls can all be happening at the same time using this technology, as opposed to trying to figure out agent staffing and how many calls can be made. I have a friend who's a data scientist and he works for a financial institution and put together like, "When is the most beneficial time to make outbound collections calls?" Surprise, surprise, it's like 5:00 PM, 5:00 to 6:00 PM, not great time that we want to receive calls then, but that's the time.

When you're working with agents, you're staffing an agent for six, eight hours. And there's a time period in there where calls are not that beneficial. Now, you can make all of those calls during the time period that is the most beneficial that are tailored to the individuals who are being called and people can't tell the difference.

The other thing that I'll say, Alan, is that it's also being used on the inbound. We have clients who are receiving a number of calls from people who are representing the customers of our clients from debt settlement companies and others, saying, "Hey, I can settle this debt. I want to make a payment for this individual." And the entities may not know that that is inbound AI calls either. And so, figuring out how to address all of those different things is a real consideration.

Alan Kaplinsky:

Okay. Well, boy, you told us about a lot of very impressive use cases, but let's now look at some of the legal risks that are involved. I'll call them the core legal risks. And we'll start with you, Joseph. Who's legally responsible when an agentic system acts? And I guess we heard the Air Canada, that wasn't a hypothetical, real case that Oren mentioned. Maybe that answers that question, but go ahead.

Joseph Schuster:

Yeah, I think that's right, Alan. I think that there are answers in existing law. And that's something, that as we think about AI and the use of AI, thinking about existing law is an important consideration here. We have agency law, we have UDAP laws, we have vicarious liabilities, safety and soundness principles. All of these things already exist and AI is being put into that framework. If an institution is responsible for doing something, that institution can use whatever tools it wants.

If it hires an employee, if it uses a third party contractor who's a human, it is responsible for those things. If it uses a computer program, AI, it is responsible for those. Now, there are different ways that you can look at this from the standpoint of contractual liability. All right, an institution's responsible for doing something. They have contracted with an entity that provides an AI service, maybe in that debt collection space like we were just talking about.

At the same time though, if those calls are being made on behalf of the institution, that institution is still liable for violations of law that happen. Not only that, they're probably also going to have some vendor management and their oversight responsibilities. They're going to get taken into account by regulators as well.

So the biggest thing is AI is a tool. It's a program. It's not a liability shield. We'll talk a little bit later on too about the different programs that institutions should be looking at in terms of managing these to reduce that legal responsibility that can happen from problems that allow you to use like the bonafide error defense. These are things that can be used when you run into a systemic issue with respect to technology. That's AI, but you have to have the right program in place. Without that, you're going to have legal responsibility.

Alan Kaplinsky:

Right. So, I'm going to turn to you, Adam, now. How do we think institutions should be thinking about, I'll call fair lending risk, or risk of there being violations under the Equal Credit Opportunity Act and Regulation B?

Adam Maarec:

Sure. So, the Equal Credit Opportunity Act and the requirement to not discriminate applies whether you're using traditional underwriting models or AI based underwriting models. There being a black box in terms of the criteria that are being used to make underwriting decisions is a risk, and not a defense to potential discrimination.

And so the baseline rule, but we are moving. I think there's an overarching shift in perspective from thinking about the explainability of models, and fully understanding how AI models are developing their outcomes, and shifting the regulatory perspective to focus on conceptual soundness and testing rather than complete explainability.

Now, that may not apply in underwriting and credit decisioning yet, but as you're using these models for other kinds of meaningful decisions for customers, the fact that you can't fully explain how an AI model reached its conclusion is an inherent challenge. And we have this balancing act where you want to adopt these new technologies because they are powerful, but at the same time, they are inherently difficult to explain.

Alan Kaplinsky:

Right. So Joseph, UDAP, a bullet point on -

Oren Bar-Gill:

Alan, can I chime in for a second?

Alan Kaplinsky:

Sure.

Oren Bar-Gill:

Would that be okay?

Alan Kaplinsky:

Absolutely, Oren.

Oren Bar-Gill:

This is just about the potential for discrimination with these black box algorithms. One of the main reasons why they're called black box is because explainability is a huge challenge. It's very hard to figure out what they're doing.

And just kind of following up on what has been said before, when we're thinking about discrimination issues, we will be, I think, facing significant challenges in litigation with respect to disparate impact, where the focus is going to be on outcomes,

whether the outcome is that certain members of protected classes received significantly different offers, terms, and so forth, with or without actually focusing on the inputs and the explainability of how exactly this came to be.

And there is a lot of recent research and also interest in policy circles about doing this kind of outcome-based disparate impact type analysis of what algorithms are coming out with. Thanks.

Alan Kaplinsky:

Yeah. Thank you, Oren. So Joseph, UDAP, what are the UDAP risks when AI communicates directly with consumers?

Joseph Schuster:

Yeah. And I'll get into those. I just want to touch on one item that Oren said about disparate impact. It's a great point. Disparate impact is something that has a lot of risk today in a number of different ways, both from just this specific piece of, is there disparate impact, but also concerns with respect to laws around whether or not there should be valuation of disparate impact. So that is one where, as you're looking at this, you're making sure that you're working with law firm, doing it under privilege, is very important and very valuable with respect to disparate impact.

With respect to UDAP, I'm going to touch on just a number of the high-level things that we see. Hallucinations is the term that's used now. We used to call this providing inaccurate or incorrect information. That's always been a UDAP. Now we just have a fancy word for it. Whether it's an individual who's providing something that's wrong or it's AI that's providing something that's wrong, that is seen as a UDAP, probably rightfully so. If somebody asks for a payoff amount and they don't receive the right value or for fees or for interest, how something works, and we saw that with the example that Oren provided, that's going to be a problem.

The second one is an overconfidence bias. The way that these work is, there's a probability of something being correct. And when it hits a certain probability, it takes the position that it is correct. This is probably happening, to some degree, with respect to humans today, but we don't see that. Now there is that underlying piece of, what is the probability for something being correct? But also, this is where you can put boxes around AI to say what the parameters are they can operate within, and they shouldn't have that overconfidence if it goes beyond something that is within their knowledge base.

Disclosure failures. This is one where depending on what they're doing with respect to servicing, collections, marketing, there may be specific items that are discussed that require disclosures. Those all have to be put in. Failing to provide a disclosure can be UDAP, as well as perhaps a violation of a different underlying law. Certainly UDAP, the FTC gets into a lot of this from the standpoint of how disclosures are used.

And then the fourth one that I had is kind of what we were just discussing, is that that personalization risk, and that hits on a number of different layers. And I think Oren talked about that with pricing, we talked about it with potential disparate impact. And there's also kind of the behavioral nudging that can happen with it as well, of suggesting that somebody should go towards a particular product, or telling them that it could be beneficial to go in this direction or that direction, just because the AI might be thinking that that's where the person is most likely to engage with an entity and end up with using one of their products. And so just there might be that nudging there, which maybe that's not the best product, but it's something that AI can be used for.

Alan Kaplinsky:

So, what about model risk management and governance, Adam?

Adam Maarec:

Yeah. So, model risk management is the primary framework I think that banks are using to apply as they think about the deployment of AI models. But I think maybe just to zoom out, the regulation of AI is a rapidly evolving area. There's been AI legislation introduced in all 50 states, Colorado, California, and a few other states have adopted laws that apply to AI.

At the same time, the current administration is opposed to fragmented state regulation of AI and wants a minimally burdensome national policy framework, and is thinking about ways to preempt these laws. And so, with that backdrop, we've

got federal banking regulators thinking about how to oversee and instruct their supervised regulated institutions as to how they should be managing these risks.

And so, model risk management is a good framework to think about how to manage these tools as just yet another instance of very complicated... But if I can, I know we're getting a little bit short on time, so I want to make sure we hit on a couple of points around governance and board oversight. The right level of management has visibility into the AI tools that are being considered and adopted is important.

And we've seen clients adopt some cross-functional governance committees to bring in different stakeholders from different parts of the institution to make sure that you have different points of view at the table to help evaluate new tools, new deployments before they are actually rolled out. And then setting up model inventories and risk tiering. Data mapping is a really important component of this. Understanding all the data that's coming in to power these models and making sure that the integrity of it is managed.

And third party risk management is another critical component of that overarching governance. As more likely than not, companies are going to need to rely on vendors for LLMs and other tools to deliver the value from these AI services. And so, as you onboard new vendors asking the right questions about the data that they're using and how they're making decisions, and how their tools work, is really important and challenging, given some of the novel use cases that are coming up.

Oren Bar-Gill:

There are particularly third party providers that do this type of kind of model risk assessment, and they will help you make sure that you are not, or you minimize the risk of disparate impact liability to make sure that there's fairness embedded into your models when you're designing them. Okay. So these are, again, tools that can minimize risk of liability and regulatory attention.

Alan Kaplinsky:

Thank you. That's very interesting to know that. Joseph, what are regulators and supervisors expecting when they examine companies' AI practices?

Joseph Schuster:

It's a great question, Alan, and it goes right to what Oren was saying about that regulatory attention. The better governance framework you have, the easier these conversations are. Examiners are not asking if you're using AI. They're asking how. So, you really want to have an AI task force or a committee or a working group, whatever name you use, set up that's evaluating the different AI use cases that you have. So you can show that governance and the different AI, where it's being used.

So when the examiners are asking that, you're not gathering that information, you're not creating it when they're there. It's something that you already have. And that reduces the additional questions that we see clients face from examiners, because examiners are then comfortable that organizations are being thoughtful about their use of AI with respect to laws that are specific to AI, but also those specific substantive laws that apply generally in how AI is used.

We work with a number of clients to create intake forms that ask, "Where is the AI being used?" Again, a question that regulators ask, "For what is it being used? What is the consumer impact? How would you compare it to what's happening for that consumer without AI? How does the explainability work?" All of those types of things, having it in that intake form.

And then you have all the right people in that task force, committee working group, from legal, compliance, your operations, your servicing teams. Everybody who can identify how it can impact different areas of the organization so those different laws that you're thinking about or that already exist are being covered.

I think that when we think about AI, the new laws at the state level, at the federal level that apply specifically to AI, those are the shiny new things. We're not going to miss those. Those are easy at this point. Where there's really the higher risk here, I think, is with respect to when an AI process is put in place without the appropriate governance in order to identify what the substantive other laws are that is potentially implicating.

Adam Maarec:

And we've talked a lot about the risks of adopting AI and how to manage those risks. The former Comptroller of the Currency laid out in the speech last year, a really interesting angle on this, that institutions should also be considering the risk of not adopting AI tools.

And that we are at an inflection point now where the tools are readily available, particularly for fraudsters, that you face a risk in not adopting these tools and not keeping up with fraudsters and your peers from a safety and soundness and competitive perspective. So, there is some urgency here to adopt these tools and figure out the right framework for your institution as to how to manage all the myriad risks that we've discussed today.

Alan Kaplinsky:

Yeah. So, let's narrow in as we're getting near the end of our webinar on agentic AI in the payments area. And going to turn to you, Adam, and what does autonomous initiation of payments look like in practice?

Adam Maarec:

Yeah, we could probably have a whole webinar on this, and I know we're about at time here, but really to get a flavor for what we're talking about here, agentic AI and agentic commerce generally refers to the concept of asking an agent to go out and make purchases for you as a consumer.

So, a good example is that I need to plan a birthday party for my 12-year-old kid and I go into an agentic platform and say, "Hey, please give me all the ideas I need for a baseball-themed party for 12 kids that are ages 10 to 12. Order me all the supplies and give me an agenda of things to do for two hours."

And instead of ordering supplies for 10 kids, it orders supplies for 1,000 kids. But the agent had my authority at that moment. I've given the agent my payment credentials, my credit card, my debit card, my bank account number, and the agent goes out and uses that authority, but makes a mistake.

And there's real consequences for everybody involved, for the merchant that's going to have to deal with that transaction, for the consumer that's going to have to deal with those erroneous payments, and then the agent, what is the agent's role?

And so this is an emerging area where a lot of thought and effort are being put into thinking about the legal and regulatory implications, but also the technological solutions to allow agents to interact with commerce platforms so that it's clear that they are an agent acting on behalf of the consumer. And so that will help inform merchant and payment processor risk protocols. So, I hope we can do another webinar on that topic at some point.

Alan Kaplinsky:

Yeah. I think we will probably be doing some follow-up webinars where we drill down in a lot more detail on some of these areas that we had to cover very quickly.

So, we're at the bewitching hour, I guess you could say. I'd like to really, all three of you, and if you don't mind, Oren, I'm putting you on the spot a little bit, but what do you think some of the priorities ought to be about companies that are beginning to consider how to use AI, either agentic AI or generative AI? What would be a takeaway that you might be able to provide to our audience?

Oren Bar-Gill:

So Alan, can you hear me because I'm a little bit frozen here?

Alan Kaplinsky:

Yes, I can.

Oren Bar-Gill:

Okay. All right. So my main takeaway is that we really need to think carefully about distinguishing the good side of AI from the bad side of AI. And it's been very helpful to hear from the panelists here in the webinar, all of the use cases that actually work in favor of consumers, but we should also be careful about various use cases that could harm consumers, even if not intended to do so.

And so I would encourage companies to be very much aware of that. And so just a very quick example, where we know that using AI to send collection calls, it could be very effective, but it can also be harmful to consumers if these calls are used maybe to threaten or to bully.

And then AI might learn to do that because that might be an effective strategy to get more money out of consumers. So I just think we should be very mindful and careful of the potential downsides and harms from these technologies that often can be very, very beneficial.

Alan Kaplinsky:

Yeah. Let me ask you, Oren, I'm just curious, it's a question that someone in our audience asked, is there a requirement... I think it was Joseph who mentioned, an ordinary person is unable to tell whether they're talking to a robot or a real human being. Is there some requirement, that at the outset, if a robot is involved, there needs to be an oral disclosure of that or a written disclosure? Is that, do you think, critical?

Oren Bar-Gill:

So, I think it is something that is critical. I will say there's a recent New York state law that requires a disclosure when pricing is being done by an algorithm. Okay. So this has to be disclosed to the consumer. This is a very specific rule about AI. I think that we have general law in place.

So I think we were talking before about UDAP and other unfair and deceptive, especially deceptive rules that we have at the federal and state level. It is at least, in my opinion, a plausible argument, that if you are pretending to be a human being, but you're not a human being, that might be considered to be deceptive. So I would be, I think it's kind of best practice to be very clear about who it is that the consumer's talking to.

Alan Kaplinsky:

Right. Okay. Thank you. Adam, give us one takeaway.

Adam Maarec:

I think governance is a key theme I want to leave folks with, making sure that your institutions have good policies and procedures, and controls established to evaluate new AI use cases, and that you're continuously applying those governance protocols and evolving them as the technology changes.

Alan Kaplinsky:

And Joseph, I'm going to give you the final say.

Joseph Schuster:

Sure. I'll say don't forget about existing substantive laws. How these new processes interact with all the existing federal state laws that you're already covering. Oren mentioned collections calls, threats or bullying tactics. Those are already prohibited by underlying substantive law.

Those apply to these types of interactions. And you need to make sure that the infrastructure that you're building, that governance that Adam was talking about, takes into account all those existing substantive laws that all the processes that you have today are following.

Adam Maarec:

Right. You can expect things to break, notwithstanding your best intentions with all your training for your agents and all your controls around your technology, expect things to break and have controls in place to catch it and quickly address it and remediate.

Alan Kaplinsky:

Yeah. Well, that's a really good point to end on, because while I think you mentioned, Adam, during one of your presentations earlier on, that there aren't too many AI-specific laws that apply to consumer finance transactions, Colorado being one of the leading AI-specific laws, and I think there may be exemptions in there for banks, I believe, but that doesn't mean that all the panoply and array of existing federal and state consumer financial services laws don't continue to apply.

They all continue to apply, every one of them. And so whenever you're implementing a new program, actually, you've got to think through all of the potential statutes and regulations a trip up on if you're not careful. My thanks again to our special guests, Professor Oren Bar-Gill, and of course, my thanks to my colleagues, Joseph Schuster and Adam Maarec.

To make sure you don't miss any of our future episodes, please subscribe to our show on your favorite podcast platform, Apple Podcasts, YouTube, Spotify, or wherever you listen. And don't forget to check out our blog, consumerfinancemonitor.com, for daily insights about the consumer finance industry. And if you have any questions or suggestions for the show, please email us at [podcast](mailto:podcast@ballardspahr.com), that's singular, [@ballardspahr.com](https://twitter.com/ballardspahr). Stay tuned each Thursday for a new episode of our show. Thank you for listening and have a good day.