Consumer Finance Monitor Podcast (Season 9, Episode 42): The End of Shopping? Agentic AI and the Future of Consumer Financial Services Introduction

Speakers: Alan Kaplinsky and Mark Bartholomew

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 business, consumers, and 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. I'm the founder, 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 are looking for more information, either about the topic that we're going to be discussing today, which I will introduce to you in a moment, or anything else in the world of consumer financial services, don't forget to consult our blog, which also goes by the same name as our podcast show. It's called Consumer Finance Monitor. We've been doing the blog since July 21, 2011, the very same day that the Consumer Financial Protection Bureau became operational. And there is a ton of industry content there, and including a lot of content about the topic that we'll be discussing today. And if you like what you hear today, please leave us a review on whatever platform you may have used to access the podcast show today. And if you have any ideas for any other topics that we should consider covering or speakers that we should consider as guests for our show, please let me know about it.

Today's episode examines a development that could fundamentally change consumer commerce, and it's called technically agentic AI or agentic artificial intelligence. For decades now, technology has made shopping much easier for consumers. We move from brick and mortar stores to online shopping, from desktop computers to smartphones, and from manually comparing products to receiving personalized recommendations generated by sophisticated algorithms. But according to my guest today, we're about to enter an entirely different era. Instead of consumers making purchasing decisions themselves, AI will increasingly make those decisions on your behalf.

Can you imagine telling your AI assistant, whoever that may be, things like, "Replace my running shoes," when they're worn out, "Find me the lowest cost auto insurance policy. Switch my checking account to another bank that offers a better deal than the bank that I'm using now." And I've signed up for a lot of subscriptions in the last couple of years. "Cancel any subscriptions that you see that I'm not really using, and book my family's vacation within my budget," which you're already familiar with. The AI, not the consumer, conducts the search, compares the products, negotiates the prices, and completes the transaction. Major technology companies are already developing versions of these AI shopping agents. As these systems become more capable, they may transform how consumers purchase virtually everything, from groceries to obtaining mortgages to obtaining credit cards.

The importance of agentic AI is underscored by the number of podcast shows we've already devoted to this topic this year. On July 16th of this year, the guests on our show were Bridget McCormack, President and CEO of the American Arbitration Association, and David Hoffman, who is a professor of law at University of Pennsylvania Law School. They are the authors of a very thought-provoking short paper, I wouldn't call it a law review article, that's entitled Agentic Commerce Needs a Legal Infrastructure, and the Courts are Coming. And it's published on the AAA website, and I would commend that to your reading, or you can listen to our podcast show. It's in our archives. That episode that we did built on an earlier May 21, 2026 podcast with Professor Mark Geistfeld of the NYU Law School who discussed the ALI's principles of the law project entitled Civil Liability for Artificial Intelligence. Together, these two programs examine both the tort law and the contract law challenges that are posed by agentic AI systems.

Our guest today argues that this represents far more than another technological innovation. In his recent article, which he co-authored with someone else who I'll mention in a moment, the name of the article is called The End of Shopping. And I'll tell

you exactly how to get it in a moment. He contends that autonomous shopping agents raise fundamental questions about consumer autonomy, competition, privacy, behavioral economics, and the adequacy of our existing consumer protection laws.

Joining us today is the co-author of that article, Mark Bartholomew. Mark is the vice dean for research and faculty development at the University of Buffalo School of Law. He writes and teaches in the areas of intellectual property and consumer protection law. He received his BA from Cornell University and his JD from Yale Law School. After clerking on the Ninth Circuit Court of Appeals, he practiced law both as a litigator for a San Francisco law firm and in county government. His latest book, called Intellectual Property and the Brain: How Neuroscience Will Reshape Legal Protection for Creations of the Mind, it was published by Cambridge University. Together with Professor Samuel Becher, I hope I pronounced his name correctly, or Becher, he has written The End of Shopping, which I found to be a fascinating and a provocative article exploring what happens when shopping itself becomes fully automated. It's available on SSRN. Don't recall what the initials stand for, but if you go on ssrn.com, that will take you to the appropriate website. It will soon be published in volume 68 of the William & Mary Law Review. Mark, a very warm welcome to our podcast show.

Mark Bartholomew:

Thanks for having me on the show, Alan. I'm really looking forward to the conversation.

Alan Kaplinsky:

Okay. Well, I have a lot of questions for you about your article, but let's set the stage first. It's got this intriguing title that's called The End of Shopping. What exactly did you mean by that type of title? I know it gets the attention of a lot of people, but... Hold on a second. Christina, I just got a thing that said installing an update. I hope that doesn't interfere with anything. Yeah, okay. Installing an update of Adobe in the middle of this. Okay, let me start over.

Mark, your article has got this intriguing title called The End of Shopping. It naturally draws the attention of somebody who's wondering what in the heck are you talking about in this article? What exactly did you mean when you came up with the title, The End of Shopping? Mark?

Mark Bartholomew:

With the title, we meant to get across that shopping is a task that we perform. We research the prices, we go to the store, we compare things online, we do market research, we see what best meets our needs. And if shopping is suddenly something that we no longer perform but something that's performed for us, as you alluded in the intro here, that it becomes automated, then it's not really shopping anymore. That's why we wanted to call it The End of Shopping. Of course, we'll still need consumer goods and different things like financial services, but what if AI suddenly does that for us instead of us doing the task ourselves?

Alan Kaplinsky:

Yeah, yeah. Many of our listeners are familiar with what are called AI assistance like ChatGPT or Claude, which is the creation of Anthropic, or Grok, which is related to SpaceX or Gemini. And I could go on and on. There are many, many assistants. And how is an AI shopping agent different from these types of AI assistants that I mentioned?

Mark Bartholomew:

The critical difference here is that instead of merely recommending things... And people are using the LLM models you just talked about, things like Claude and things like Gemini to help them with their shopping decisions. "Oh, can you recommend me a mid-price family sedan?" And they can offer some thoughts for you. But in the article, we're more focused on the next step forward, which is these AI shopping agents, which is not just presenting recommendations, you go to sleep and they do the transactions for you. They go ahead and buy that subscription or they cancel your internet access from one business and switch you to another that's a better deal. Or, like you mentioned in the intro, they say, "Mark has been running for about nine months now. Let me go ahead and buy him that new pair of shoes that matches his budget and his running style." It's the

doing it for you as opposed to just offering a recommendation that a consumer can choose whether they want to follow or not.

Alan Kaplinsky:

Yeah. How close are we to the future? I mean, I know that there are systems already in place, but I wouldn't say that the usage is very widespread yet. Are these systems still in the speculative stage or in the very early stages, or is this something that is going to become very widespread in the near future?

Mark Bartholomew:

I think the average consumer is not going to be using this in the short term, but that I would say is more an issue with comfort than technological capability. A couple things about the technological capability. One is that these agents are being introduced all over the place now by major players. Google, Walmart, Amazon, they've all unveiled their own agentic tools that will do the shopping for you, and they've done things like autonomously ordering your shopping list and recipes for the week. Or Anthropic had an agent that actually managed the store, took care of inventory, ordered the items, managed the store. JP Morgan had an agent that was executing stock trades.

I think this is more and more feasible. And we can talk maybe a bit about why it makes business sense. I think the biggest hurdle is just making people comfortable with this. A lot of people tell me, "Oh, I don't know if I'd be willing to turn it over to AI. I don't trust it yet to do it, but shopping for me." I think the bigger hurdle here is going to be making people comfortable enough so they're willing to sign over their shopping needs to these agents. But the technology's getting closer and closer.

Alan Kaplinsky:

Yeah, and I would guess it's not like all of a sudden you're going to download some software and immediately be able to start using it. Don't these things have to be trained over some period of time? They've got to get to know you as an individual. You can't expect them to come in as a stranger in your house and start giving them all kinds of commands to do the things that you mentioned or that they mentioned.

Mark Bartholomew:

Sure. Sure. And there have been some trial runs where some publicized mistakes were made, ordering 100 of an item that should only require one purchase or going rogue and spending more than the person really wanted. There are mistakes, and probably these mistakes get a lot of attention. Like with autonomous driving vehicles, they're more safe than when I drive or someone else drives in general, but we pay attention to the mistakes. And so I think that's happened here. But you're right, it will take a bit of tailoring.

I would just say, though, that once an AI shopping agent is ready and they can access both a wide array of retail options, surveying the options for retailers, and they have access to my full data profile... Think of everything that Axiom has about me or everything that Google has about me. Once they access both those pools of data, I think they're ready to execute. I think at that point they will know us enough to be able to recommend and execute things that will satisfy our needs, that won't be something different from what I really would purchase if I was making the choices.

Alan Kaplinsky:

Yeah. Yeah, yeah. Let's talk about why shopping matters. One of the most interesting aspects of your paper is your argument that shopping serves purposes beyond simply acquiring products. Why do lawyers and policymakers often overlook that point?

Mark Bartholomew:

Yeah, that's a great question. I think in general, maybe even beyond the lawyers and policymakers, shopping is just so much in the background. It's so ubiquitous. It's part of our daily lives that we almost don't pay attention to it. Sometimes you can't see the things that are in plain sight because you see them all the time.

I think beyond that, a couple theories about why the non-economic components of shopping haven't garnered a lot of attention with policymakers and lawyers. One theory I might have is that when you look at the political economy of this, who has lawmakers ear, it's businesses. And businesses are less concerned, I think, with some of the things we'll talk about, the sociological, psychological benefits of shopping. They're more concerned with the bottom line, "How many of these units can I sell?"

I also think there's something, at least that I've noticed when I talk about this article with people, particularly in the university setting, but sometimes with lawyers too, they don't like to admit that they like shopping. It's like, "Oh, that's something that's tedious." Sometimes it's even coded as feminine, and sometimes a guy doesn't want to say that they like shopping. And so they're not really in touch, I think, sometimes with some of the things we'll talk about, the psychological benefits, the spiritual benefits, the social benefits of shopping.

Alan Kaplinsky:

Yeah, yeah, yeah. Some people might respond to you, "If AI saves me time, why should I care?" What's missing from that perspective?

Mark Bartholomew:

Absolutely. And I don't want to come on your show and be Debbie Downer. I think there's lots of good things about AI shopping agents, just the article is meant to point out both the pros and the cons of them and also how we can regulate them so they do what we need.

But in response to your question, "Well, if it makes it faster, if it saves me time, who cares?" I'd say that one thing that shopping does for us is beyond just getting me the things I want, satisfying my rational preferences for items, it's a key component of our mental health in a way. One thing that psychologists talk about when I've done the research on the psychology of shopping is they talk about how people need to feel in control of their stories. And there's so many things we don't feel in control of, like what's happening in politics today or even sometimes what's happening in other aspects of my life, my job. But when I'm in the grocery store and I'm picking out what I'm going to get for dinner that night, I feel like I'm in control. It's something where I'm exercising my autonomy in a successful way.

And I can express my identity through how I shop. If I'm an environmentalist, maybe I buy Patagonia. If I code more on the right side of things, maybe I buy Black Rifle Coffee. In general, maybe I want to buy a Harley Davidson motorcycle because I want to say I'm a Harley rider kind of guy. Shopping is a way to communicate identity to people.

It also has other benefits. Just quickly, I'll name a couple more psychologically. Anticipation. It turns out that some of the things we buy, the main value we get out of it isn't actually consuming it or using it, it's the sense of anticipation up to we use it. For vacations, turns out that often vacations aren't that enjoyable, but the time as we're looking forward to that trip to Florida or that trip to Disney World with our family, maybe we fight when we get there, but in the time up to it, everybody's looking forward to it, and we get more utiles of pleasure from that than the actual trip itself. Well, if an AI agent just does this while I'm sleeping and delivers the product to my door, I'm losing out in that anticipation component.

And I'll just mention one other thing that doesn't have anything to do with efficiency, it doesn't have anything to do with economics, but there's also the sociality of shopping, and that's another key part of what shopping does for us in a positive way that could be lost if we outsource it. A lot of shopping that we do occurs with other people. We might shop with friends. I might shop for gifts. I might just encounter other people in stores. Even if I'm doing online shopping, I'm paying attention to comments. I might be going on social media and getting thoughts about the things I'm going to buy. And so we might lose some of the social benefits of shopping if this thing we perform ourselves for so long gets turned over to a machine. I think those things should be part of the mix as we think about how much we want to entrust these agents to do the shopping for us.

Alan Kaplinsky:

Yeah. I get what you're talking about, I really do. I mean, I've heard people say... It's mostly women, and I don't want to sound like a misogynist here, but, "I think I need to do some shopping therapy today. Not to see my therapist, but just go out and

shop because it's fun." It's something that's enjoyable, and it's distracts you from other problems that you may be dealing with, either personal problems or problems, things going on in the world.

Mark Bartholomew:

One interesting thing I learned is that retail therapy... You probably heard that phrase, retail therapy before.

Alan Kaplinsky:

Yeah, retail therapy. Yeah.

Mark Bartholomew:

That's a real thing. Health professionals actually talk about the benefits of retail therapy. It can go too far. Don't spend your college savings on retail therapy, but there is something to it as a therapeutic tool. And of course, if you let a machine do it, it wouldn't have the same kind of benefits.

Alan Kaplinsky:

Yeah. Let's turn for just a few minutes to consumer financial services, which is the area that our podcast show concentrates on. Our audience consists largely, though there are certainly retailers that are part of our audience, but a lot of them, it's mostly banks, fintech companies, mortgage lenders, credit card issuers, other kinds of consumer financial services providers. How might AI shopping agents change the way that consumers select financial products?

Mark Bartholomew:

Yeah, I think financial products, and we didn't dwell on them in the article, but maybe we should have, is a great example because I think the average consumer finds shopping for financial products more intimidating than the other shopping they do. It's different than going to the grocery store. It's different than buying a lawnmower. I think there's more concern at the consumer level about being outmatched. How do I scrutinize these things? It's sometimes hard to make comparisons. The terms of service are going to be more complicated than looking at that warranty for other more physically tangible items.

To me, these things like shopping for a credit card, shopping for an auto loan is something where consumers would be pretty likely to embrace it. And the way it would work is you task the agent, "I'm thinking about a new auto loan," or, "I'm thinking about a new credit card. You know things about my financial status and what I'm looking for and what I have now. Go out and comb through every listing available and tell me what my options are and what I should get." And the agent, unlike me, who gets bored or doesn't have the tolerance to look at more than five or six, let's say, credit card profiles, could look at hundreds, thousands of them, potentially. Unlike me, who even though I'm a trained lawyer, I tend to tune out the terms of service, the AI could potentially scrutinize every credit card company's terms of service and say, "You know what? Don't go with this company. They really have just a Draconian arbitration provision. Go with this one instead."

And then the other thing I'll add is why this might be a valuable area for agentic shopping is switching costs. I know for myself, I know I should think more strongly about doing comparison shopping for, say, homeowner's insurance or locking in a new mortgage rate when interest rates change, but, oh, it's a hassle and it's not as easy as just doing things like buying something at the grocery store. But we envision a world where the AI could do that for us. It calls them up, it goes to the website, it switches you out of your old mortgage loan and switches you to another. We could see something here where this would be a great candidate for value added for consumers.

Alan Kaplinsky:

Yeah. And you mentioned a bunch of things that it could be used for, new checking accounts, you want to buy purchases, certificate of deposit. Get me the highest rate in the country on a one-year CD. It's insured by the FDIC, so I don't care what bank it's with. Investment products, I guess agentic AI could trade stocks for you. I mean, I think I've seen some things that they advertise on CNBC that sound like you can do that today, like you can get software that will take over your investments in stocks and bonds and et cetera, et cetera. It exists, I think.

I know very recently, a brokerage firm that I was using, the individual who I had the relationship with went over to another firm, and, oh, the headaches involved in switching all the accounts and making sure that the apps on my phone were capturing all my accounts. And it was the biggest headache in the world. I would imagine that that's something that agentic AI could handle very rapidly and probably a lot more accurately than human beings.

If an AI shopping agent opens up a new account, it could be a deposit account or switches lenders, who is legally making the decision? Is it the consumer? Is it AI, just some software on my computer? Or maybe I've got a robot, a humanoid. Or is it the platform that's operating the AI?

And now we're starting to get into some of the legal relationships. And what is this thing that's making all the decisions for me? If the AI agent, if I tell him to buy one pair of shoes and he buys 100 that get delivered to me, can I tell the seller, "You got to take the other 99 back and pay for the shipping. I didn't do this. It was some robot that did it."

Mark Bartholomew:

Right, right. Yeah. Great question. And I think what we'd say is this is an example of delegation, of course. And we delegate responsibility in all sorts of areas, so the individual user would be delegating the authority to make these choices to the AI assistant. And so we would say that the assistant is making the choice, but then the question is, well, who's going to be responsible when the assistant messes up? Like you said, buys 100 pair of shoes instead of one. And I think the platform in that case would have some responsibility. We've already seen some situations. There's a fun case out of Canada where an airline had a customer service chatbot that promised certain reduced fares, discounts, and then the airline refused to honor them. And a tribunal said, "No, no, you have to comply with that. It's your agent."

Now, the contours of all this responsibility is still an open question, and maybe this is something you've talked on the show before. Can we officially say, "Well, this is a question of agency law"? The AI is the agent and the principle is the platform that supplies the agent. And I think there's some reasons to be nervous about that or not think that it's a one-to-one fit. Agency law is developed for humans. Both the principal and the agent are humans, and there's certain things about incentivizing better behavior through agency law and sensitizing the agent through liability for bad decisions. But of course, the agent that mistakenly orders you 50 Ferraris or 100 pairs of shoes, they don't care about their personal financial liability.

I think that's a question that we punt on the paper, say, "I don't think this is a panacea." What we're more concerned about in the paper is not the ex-post rules about who's responsible when something goes wrong but the ex-ante framework to try to make sure these things work as well as possible from the start. We offer some reasons why, if these things aren't set up the right way from the beginning, these agents won't deliver on the promise that we've been talking about so far in the podcast. We're focused on what's the best way to build these things so they do what we want them to do.

Alan Kaplinsky:

Right. Let's turn to consumer protection law for a minute. Much of the consumer protection law that we have both at the federal level and at the state level assumes that there's a human being involved that's able to read disclosures, to compare products, to look at the specs of different products and make decisions. What happens when those assumptions aren't true? Is the whole foundation of consumer protection law basically irrelevant?

Mark Bartholomew:

Right. Well, this would definitely be a paradigm shift in a couple ways. Let me just give you a couple things that would change greatly if these assistants are doing the shopping for us. One is I think we'd have to have a different approach to data. We haven't had to worry that much, I think, about privileging consumer availability to large amounts of data. And price, availability, quality, we haven't said, "Oh, this is guaranteed for the consumer." Because after all, if you're like me, a human consumer, I can only internalize so much information. We have consumer law that's focused on this bilateral exchange. It's me, the consumer, versus the seller and regulating that relationship.

But to realize the promise of AI shopping, the agents are going to need access to vast stores of publicly available data. They need to be able to scrape online postings. And currently, the law and the legality of this scraping is ambiguous. Retailers and

others are allowed to put up barriers in terms of service and those CAPTCHA things they force you to fill out. Maybe this needs to be changed and ensure this public availability of data if we're going to have agents doing the shopping.

Another thing that I think would be different when we no longer have a human being in the loop is thinking more about how we need to shift away from worrying about human cognitive biases. I work on trademark law a lot, and trademark law is about protecting the consumer from search costs. Don't waste their time by looking at something that's confusing to them. And also, certain areas of consumer protection law say this is unfair emotionally. You're going after a vulnerable population or this is an unfair way to tug on the heartstrings. We won't worry about that with AI agents. They're not emotional. They don't have cognitive biases. They're not influenced by prices that end with 99 cents.

Instead, we're going to have to worry more about what I would guess I'd call algorithmic vulnerabilities, an agent that self-deals, but you can't tell because the AI is a black box or retailer efforts to somehow mislead the AI. The AI isn't getting the right sense of the available prices and availability out there. Those are different, I think, than the traditional things that consumer protection law worries about because they're thinking about the human brain, not the algorithm.

Alan Kaplinsky:

Right, right, right. Does current law, either at the federal or state level, adequately address AI agents making purchasing decisions? Or do we need new laws?

Mark Bartholomew:

I think we need new laws. It's not that clear. I'll just mention the Ninth Circuit just decided last week, this case between Amazon and Perplexity and where Amazon users could authorize Perplexity to comb through the Amazon listings and offer comparison shopping for them, and Amazon sued. And the court said... I'll just quote here because the decision just came out. The court said, "There is little to no existing case law directly dealing with how to ascribe responsibility for AI agents like the assistants," the assistant issue in that case, so we don't really know.

And some have called for this common law of agency like we talked about before to impose duties on the firms for the actions of their AI agents. But like I said, I don't think this is a no-brainer. And even if agency law applied, even though we think there's a bit of a mismatch here, we still need these things on the front end to make sure that AI agents are working on behalf of the consumer and not fooling them or somehow engaging in self-dealing for the platform or hidden retailers.

Alan Kaplinsky:

Right, right. I guess once again, you're making your point that your article focuses more on getting things right at the beginning rather than dealing with problems that occur or mistakes that happen and figuring out who's ultimately responsible for that.

Mark Bartholomew:

We want to make sure that these things are actually tools, and maybe I can talk about in a minute that some of the great things about this, like how it could really help consumers, we've talked about a little in financial services already, but for AI agents to do that, there needs to be guardrails in the first instance.

Alan Kaplinsky:

Yeah. Well, let's talk a little bit about some of the guardrails that are necessary.

Mark Bartholomew:

Yeah. Yeah. One thing we think that's important is that it's going to be really hard for a user of an AI shopping agent to tell when it's doing a good job. An AI agent would be something that e-commerce refer to as a credence good. A credence good is one where even after you look for the product, even after you buy the product, you're still not sure if it's delivering on its promises. Think about something that says it's environmentally efficient or says it's organic. You don't really know, you're

taking the business' word for that. And AI agents would be the same thing. It's going to promise you, "Yes, I'm going to comb through 10,000 listings and I'm going to look for the best choice for you, mark the user." This is what my promise is. But did they really or did they steer me towards someone who paid for their rankings to appear higher in the AI agent's choices?

One guardrail we need, and this is something that happens to other credence goods, is we need some way to give consumers something to grab onto here. And so we think there ought to be some system of audits that allow for some disclosures to someone who's choosing a shopping agent. Think of the Energy Star program would be one example. I can't evaluate whether my washing machine is going to be energy efficient, but, oh, if I see it has the Energy Star label and it gets a high rating, I know that's good for me. How would this translate? Well, we might have something where we'd have these audits, and the audit would take a series of goods and shop for them and compare them to, say, other AI agents and say, "Oh, well, this agent gave the best prices within the 90th percentile." Or we could have an audit saying, "Well, how many listings does your agent actually look at? I'm shopping for that new lawnmower, new riding lawnmower. Did you look at 50 listings? Did you look at 1,000 listings? Did you look at 10,000 listings?" The more, the better. If there's a way we can score that and report to consumers so they can pick, they won't be in the dark. That would be one thing we recommend.

I'll mention one other thing here. We have a whole suite of recommendations. It gets wonky towards the end of the article. But one other thing I think is important is because I think this brings up antitrust issues, and so one way to combat that... We're worried about one or two players being the only game in town for these shopping agents. We much prefer a competitive playing field when it comes to the shopping agents, so one thing we would want is data portability. If you are using, let's say, Google's shopping agent, Google already knows a lot about me.

Let's say I use Google in all sorts of ways, and it also has great, great listings, but then I'm unhappy. I see that it's not combing through listings or I get word it's prioritizing some businesses that it has a relationship with already. I want to be able to switch, but maybe that's a problem. How do I bring all my taste preferences, my choices over the years, my data to the new shopping agent that I think will perform better? We'd like to ensure some sort of way that Google can't say, "Oh, I'm sorry, you can leave, but you can't take any of that data with you. You have to start from scratch." We'd like to make it portable.

Alan Kaplinsky:

Yeah. Well, let me ask you a question. I have assumed, and I guess this is an improper assumption, that not only will companies have their own agentic systems that they're hopeful that consumers will buy into, but won't consumers be able to just get an agentic system that has no relationship with any company at all that's working just for me and I don't have to worry about somebody being preferred by Amazon or some other company?

Mark Bartholomew:

Yeah.

Alan Kaplinsky:

Or will there be both?

Mark Bartholomew:

Well, that's what we would hope, just the scenario you described. Oh, well, here's a good agentic shopping agent put out by a third party with good reviews; I'll use that. But the concern is that there's this lifeblood of data that will be needed for this. And so if we just have the major players, Google, OpenAI who have a choke hold on this, that'll be hard for these other upstarts to also provide their own shopping agents. That's the thing, we need to make sure that the data is portable and there isn't these lock-in effects from the beginning.

Alan Kaplinsky:

Right, right, right. You mentioned antitrust issues and competition issues. I want to go back to that because your article talks about market concentration as being a concern. Why might AI shopping agents actually increase the market power of the largest technology companies?

Mark Bartholomew:

Yeah, I mean, a shopping agent we've been talking about in terms of, well, it's a consumer product, but it's also a kind of market infrastructure. The shopping agent would become your universe, your ecosystem for what's available for you to purchase. And so the concern there is that already dominant retail platforms might already have a somewhat receptive consumer base that could introduce their AI shopping agents. And if there aren't some guardrails like preventing self-dealing, allowing the data to be portable, well, the firms with the high market share might make it hard to switch.

We already know about how much of say online advertising is controlled just by Google and Meta. When we already have this concentration, we really need to make sure there are guardrails in place so shopping agents work in a way that best benefits consumers. We think antitrust law is going to be important in here, essential for ensuring that there is not just dominant platforms infiltrating this new technology and not opening up to other parties.

Alan Kaplinsky:

Yeah. Privacy, you talked already about massive amounts of data, personal information about consumers about their shopping habits and other things that are already online. What privacy risks concern you the most? And what can be done at the outset, which is what you're focused on principally, to guard against someone's data being improperly used?

Mark Bartholomew:

Right. I'd be worried probably the way most people are about these rich personal profiles being sold off or being heisted thanks to lax cybersecurity or just being used, sold off to steer users in ways that they're unaware. And in some ways, this is different than other things because if the shopping agents really take hold, they'd really be comprehensive. We can see it being embedded into daily life like, "Oh, well, okay, every time I turn to OpenAI or open up my phone or turn to Claude, the shopping agent will be there." It'll be persistent, and it's a little harder then to turn it off, to give meaningful consent to these things. Yeah, it's a great point about privacy concerns.

I don't have a magic cybersecurity idea for this. There's probably tech people who might have more thoughts about it. I will say that these fears, even though I think they're real, we think about this with our personal data and lots of context today. "Oh, they know a lot about me. I hope they don't get my profile. Oh, here's another letter that's saying I get a year's free worth of credit reporting." I guess what I would just say is we'd want these businesses providing the agents to maintain rigorous cybersecurity protections the way I still would with a bank account or even my social media or things like that. And on the other hand, a big point of the paper is we want the publicly available data out there, all the listings, every publicly available website to be declared free for the consumer-facing shopping agents to use because these things won't work to their potential if they can't ingest all that information to get the most information about the consumer possibilities out there.

Alan Kaplinsky:

Yeah. Yeah. We talked a little bit about regulation earlier, but your paper refers to something called an algorithmic nutrition label; what is that? You're not talking about food, are you?

Mark Bartholomew:

No, although it's designed to make us think of food. We're all familiar with nutrition labels on food packaging that tell us nutritional content per serving size. And the point of that is that it helps consumers evaluate. Lots of food, some of these things can be credence goods. I don't know how many trans fats are on this thing unless I see it in the nutrition label. We think having these algorithmic nutrition labels on shopping agents would provide the same sort of service for me when I look at the bag of potato chips and weigh just how bad this is going to be for me. They create some mandatory disclosures that people could use to try to decide if this shopping agent is right for me, if it's going to deliver the way it should, and also to help build up that competitive field, one shopping agent competing against another to gain consumers' business. And so things like I mentioned before, these nutrition labels could talk about score as far as how they delivered on benchmark prices for consumers, number of listings they canvas, things like that so the consumer just isn't in the dark when they're picking one of these things.

Alan Kaplinsky:

You also recommend preserving consumer self-determination. How can regulators encourage convenience without eliminating meaningful consumer choice?

Mark Bartholomew:

Some of our recommendations are simple, like insisting that platforms provide an off button. Consumers should be able to take control when they want the agent to no longer have the reins. One example would be gift giving. I might want the agent to do most of my shopping. When it's time to shop for my wife's anniversary present, I want to be able to hit the off button there so the gift means something. I don't get in trouble because I didn't sacrifice any time or put any thought to it, I just outsource this to the agent. I want to be able to hit the off button. Making sure that's available.

And then short of an off button, a way to provide some of these intangible or less tangible benefits from human-oriented shopping while still allowing for some automation is calling for just structural features that maybe these are supplied by market pressure from competing shopping agent providers. Maybe they're mandated, but doing things like allowing me to switch to, say, explore mode or surprise mode so I'm not just getting things that are based on Mark's 25-year shopping history because maybe Mark needs to update his wardrobe or maybe Mark needs to listen to new music every once in a while. Doing things like that so I'm not just treading in the same past purchases.

Now, maybe businesses don't always want that. Past purchases are predictable. I can probably get me to buy more things when I rely on the past. But maybe I can toggle that surprise mode to at least get some of the more psychological benefits of shopping for me that I used to have when I did all this stuff by myself.

Alan Kaplinsky:

Right, right. You talk in the article about a sunset provision for AI delegation. What do you have in mind there?

Mark Bartholomew:

Yeah, our thought there is that we definitely have opinions, and overall we think that the benefits should outweigh the costs of agentic shopping if we put some of these guardrails in place, but we don't know if we're going to get this balance exactly right, whether with our recommendations, the suite of recommendations like algorithmic nutrition labels and portability mandates and voluntary audits, and also just how effective this technology going to be. If we don't get the balance quite right, we'd like a mechanism to prevent inertia from just taking hold and cementing... the guardrails we think are important, cementing them in place as is years down the road.

The thought would be is that we would pass some of these regulations, but there'd be a mandatory revisiting in three, five years. And what we draw on is some things that I think have been seen as successes. The Patriot Act, a controversial law when it came down, but its surveillance features are an example of one thing where Congress said, "This is a big deal. We need legislation. We can't wait. But we'll impose expiration dates to make sure that we thoughtfully think about this and not just rubber stamp it in the future."

Another example is the Telecommunications Act of 1996. I think that's seen as a smart legislative success. Now juxtapose that with the Communications Decency Act of 1996 and Section 230, something I think probably some of the listeners will be familiar with, which is giving this giant free pass to user-generated content for social media platforms. And I think in 1996, we thought the internet's new, we don't want crushing tort liability. This is a good idea. Let's let the internet flourish. But 30 years later, I don't find many people who think it's a good idea. I think a lot of people think that there are some harms from social media that need to be reigned in, but it's just been really hard to dislodge Section 230. There's no sunset provision there.

We recommend something here. Let's have some full bore rules for AI. You have to have portability of the data. You have to submit to certain mandatory audits, things like that, but sunset them if we get the balance wrong so we can create a new suite of policy tools three, five years later.

Alan Kaplinsky:

Let's talk just for a second about the need for actually having new agencies or existing agencies of the federal government or state governments involved in policing what's going on. I assume you think there is a need for that. Am I right? And then the next question or the related question is are the existing agencies capable of doing that or would you create a brand new agency that would just deal with AI, AI issues in general? And it would try to draw upon leading experts in the country. I don't know how they'd be able to pay them enough money when they're being offered much, much more money to work for some of the high-tech companies or the Amazons and the Googles of the world.

Mark Bartholomew:

As far as the FTC, I mean, I think they definitely have a role to play here. They definitely have the authority to police some potentially egregious, we might call it AI snake oil in this arena, false promises, really obvious self-dealing or betraying customer trust. I think that's at least part of that unfair deception practices authority. And they've done some of that already with setting judgments against some AI companies for misrepresenting what their systems can do.

I think there's a role for that whenever there's false claims involved. And so we see a role for the FTC, but I think there's clearly some areas where the FTC can't step into the void. One is this discussion we hinted at earlier about what are the rules about agents as fiduciaries? Do they owe fiduciary duties to their users? When an AI shopping agent is nominally acting on the consumer's behalf but it has commercial incentives, referral fees, or preferred merchant placements that's it's inserting in its decisions for the user, that's closer to an agent fiduciary duty problem than straight up false advertising. Maybe the FTC catches that, maybe they don't. And then of course there's the question about, well, how is the FTC going to deploy its resources? They seem more inclined to go after obvious scoff laws than blanket rules.

Another thing I'd say is that some of this is going to involve complicated technical interoperability requirements that I do think you need your experts, you need some standard setting to be done. And I don't know if the FTC is set up to do that now or not, so maybe we need another agency. Just-

Alan Kaplinsky:

And then of course you have the other problem, Mark, of what's been going on in the Supreme Court this year with the relationship between the president and what a lot of people thought were independent agencies, and we have found out they're not so independent anymore than except maybe the Fed. You have situations like the FTC where there are no Democrat representatives on what was supposed to be... Congress wanted to be a bipartisan commission. Same thing with a lot of the federal agencies. Basically whoever is president can make sure his people... It's really an extension of the executive. And if you have a Republican administration, it's more laissez-faire. They don't like to have heavy regulation. But this is a pretty perilous time not to have any regulation.

And I speak as someone who represents the industry. I get very concerned when I hear about various systems running amuck or taking over and making or trying to hack on its own other websites and stuff like that. It would be, I guess, nice if we could create an agency, and this probably require a constitutional amendment given this Supreme Court, that was insulated from all kinds of political pressure and would be able to just do the right thing. Do you have a reaction to that?

Mark Bartholomew:

No, that's a great point, Alan. Standard agency rulemaking, the FTC under the current Supreme Court, even if we thought they had the authority, I don't know anymore. And then you also bring up the point about a political football, and I don't think we want AI regulation to be this kind of political football switching between administration to administration. Yeah, I think I see some rule for legislation from Congress setting out some of these things that we're asking for.

And then maybe you're right, maybe we need an agency created from scratch and specifically say this is modeled on the Federal Reserve or something like that so it has a little bit of insulation from these political pressures. Because the AI companies themselves don't want these rules to change every two, three, four years, they want something they can depend on. Even if they might want a light footprint, they at least want to know what the rules of the game are so they can plan.

Alan Kaplinsky:

Right, right. That's for sure. We just have a couple of minutes remaining, Mark, we're drawing toward the end of our show, but I'd like you to put... Let's do a little crystal ball gazing. And if you could look into your crystal ball and we were having this conversation let's say five years from now what changes do you think we will have already seen? And what developments would most surprise today's consumers? And what worries you the most? And conversely, what makes you optimistic, assuming that you are? There's a lot there, but what do you think? Where are we headed?

Mark Bartholomew:

Yeah, I think we'll see more and more of these AI shopping agents in everyday use. It won't be that unusual for someone to turn over the keys to their shopping, at least for some shopping tasks, to an agent. We already do things like purchase subscriptions on Amazon, and that's just like these companies getting us used to, like the frog and the boiling water, to more and more of these things being planned, automated for us. I think we'll see more of this.

I think something that might surprise consumers is that they'll see, A, some benefits from this. If we get a critical mass of enrollees, they'll do things like switch 700,000 people at once to a better subscription and it'll save them money or a better terms of service. I also think they might see some benefits in that, "Oh, you know the things that stressed me out from shopping or the things that were a cognitive load barrier, so I didn't bother to switch my homeowner's insurance? I'm freed up now because the agent did it for me." That's my rosy scenario that I think actually is realistic.

The thing that worries me the most is this potential for self-dealing. If this isn't regulated, if there aren't some guardrails, then instead of doing what these things can do, which is efficiently satisfy my preferences even better than I can do it and much faster than I can do it, then this will just be another tool that I'm not quite certain of. And you talked in the introduction about online shopping, and I think when online shopping started, people said, "Oh, great, this is going to empower the consumer." And in some ways it does. I can take up my phone and look up product reviews in the store. But I think most people don't think, oh, wow, the consumer has all the power in the world of online shopping. There's these dark patterns, there's choice overload, there's people trying to scam me.

Online shopping is still... There's information asymmetry. It's not always a fair fight. What I'm worried about here is if we don't get this right, then AI shopping agents, instead of doing what they could do, are just going to be another thing where we don't really trust it. Maybe it adds some convenience, but it doesn't really help us as much as it could.

Alan Kaplinsky:

Yeah, yeah. Well, Mark, we've come to the end of our program. And want to thank you very much for joining us today. And your paper challenges one of the most basic assumptions underlying modern consumer protection law. The consumers themselves are making the purchases. And if AI agents become commonplace, which is where I think we're headed, regulators and businesses and courts may need to rethink a lot of legal doctrines that have governed consumer transactions for decades that may not be a good fit for this new marketplace phenomenon that we're talking about.

Again, my thanks to you. Good luck with your publication. Again, remind people the name of it's The End of Shopping. It's available on ssrn.com. And it very soon will appear in the William & Mary Law Review. And I hope everybody enjoys the remainder of their day.