Consumer Finance Monitor Podcast (Season 9, Episode 43): AI in Debt Collection: Opportunities, Risks, and the Importance of Data Governance

Speakers: Alan Kaplinsky and John McNamara

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, founder, and now senior counsel of the Consumer Financial Services Group at Ballard Spahr.

And I'm delighted to be moderating and hosting today's program. For those of you who want more information, either about the topic that we're going to be discussing today, which I will describe to everyone in a moment, or anything else in the world of consumer finance, don't forget about our blog, which is called Consumer Finance Monitor. Yes, the same name as our podcast show. We hosted our blog. We have hosted it since actually the very same day that the CFPB became operational on July 21, 2011. So, there is a ton 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, you can visit us at ballardspahr.com.

And if you like our podcast, please let us know about it. You can leave us a review on whatever platform you use to access your podcast shows. And if you have any idea for any other topic that we should consider covering, or a speaker that we should consider inviting as a guest on our show, please let me know about that.

Before I introduce our speaker today, I'm going to talk a bit about the topic that we're going to be covering. We're going to be revisiting a topic that we've discussed on an increasing number of occasions, and that is rapidly transforming virtually every aspect of consumer financial services.

And of course, I'm referring to artificial intelligence, or AI. We've previously discussed on other podcast shows a number of common use cases of AI in the consumer finance world, such as consumer lending, deposit taking, payments, and activities such as marketing, origination, operations, complaint resolution, and much more.

Today, we're going to focus on AI use cases in first-party and third-party debt collection, and the legal compliance and regulatory issues that are raised by this technology. I, and I'm sure our listeners are going to be delighted to be enlightened about this subject by really one of the leading experts in the country about debt collection, and I'm referring to John McNamara.

Let me give you a little bit of John's background. John recently concluded a very distinguished career at the CFPB, where he served as principal assistant director of markets and assistant director for consumer credit payments and deposits markets.

In that role, he oversaw subject matter experts and analysts covering a wide range of markets, including debt collection, credit reporting, fintech, and the use of AI and machine learning in consumer financial services. Earlier in his CFPB career, John served as the bureau's debt collection program manager, and later as assistant director for consumer lending reporting and collections markets. Prior to joining the CFPB, John spent more than three decades in the debt collection industry, including senior executive positions with collection technology and servicing companies.

Today, John is Chief Growth Officer at Avtal. John is uniquely qualified to discuss both the operational realities of debt collection and the regulatory expectations surrounding the use of an emerging technology like AI.

So, John, a very warm welcome to our podcast.

John McNamara:

Thanks, Alan. I can't tell you how happy I am to be your guest. I feel like I've made it. I'm finally on the Alan Kaplinsky podcast, so happy to be here. Again, thanks for inviting me. And listen, with a buildup like that, Alan, if I ever get famous, you have a spot in my entourage.

Alan Kaplinsky:

Thank you. All right. First of all, let me give you a really easy question to begin with because I'm not sure how many of our listeners are familiar with the company that you joined after you left the CFPB, namely, I think I pronounced it right, Avtal.

But what does Avtal do?

John McNamara:

Yeah, so Avtal is a digital engagement service for third-party debt collectors. I'm here at Avtal because I'm pursuing my passion, which was - I'll go all the way back to when I was out of college - took a job as a collector. The first time I talked to a debtor, I thought, "Wow, debt collections is a marketing problem." And I've framed debt collections that way ever since. So, I followed the rise of players like True Accord and others that tried to reinvent debt collections around marketing engagement versus traditional models.

And at the CFPB, I was part of the team that wrote Regulation F, and we clearly tried to spell out how that could be used to engage consumers. I think it's a kinder, gentler way to do that. Anyway, roll the clock forward to 2025, Alan. I basically had three different offers to do three different things in the private sector. It was maybe time to retire and try something new. And I thought I was going to take one of those, but a friend of mine, Joe Gelbard, is the founder of this company, and they were trying to provide this digital engagement service and make it available to debt collectors so they too could have best-in-class compliant digital engagement with consumers. And the thing was working, and I loved the concept.

So, Joe called me at one point and said, "John, we got a nice round of initial funding. I can afford you if you'll take a pay cut." I didn't even ask him how much it was. I said, "You know what? I can afford to do what I want to do." So, I'm at Avtal now because it is a passion, and I think I'm pursuing a better way for responsible industry to engage with consumers in a way that limits the ability to rough them up or harm them in any way, shape, or form.

And quite frankly, it has been really successful.

Alan Kaplinsky:

Terrific. Sounds like an absolutely perfect job for you, John. I'm certainly happy that you didn't opt for retirement. You're much too young to do that.

John McNamara:

Thank you.

Alan Kaplinsky:

John, AI is one of the most talked-about topics in financial services today. Before we get into specific use cases, could you define AI in the debt collection context, and what tools are companies actually using?

John McNamara:

Yeah, no, it's a great question. I will say this upfront. The industry right now is almost abused by marketing puffery and hype, and everything is AI right now. Users and buyers trying to work through the smoke and the confusion, I think, have a rough time. And again, Alan, and you know this from your other podcasts and certainly your long and extensive career, that a lot of what's marketed as AI is just robotic process automation, and things like that. I would encourage anyone that's looking at AI in the debt collection space to make themselves familiar with the Gartner Hype Cycle, where literally you've got this very high peak of expectations followed by the trough of disappointment, and that's where we're at now.

But there are a number of uses of AI that are in play right now; they're in production, and it's sort of better to define the AI by what it's used for than what it actually might be.

But account scoring and segmentation, channel timing, like how and when to communicate with a consumer, almost treating the data set like a marketer would to find out what's going to cause a consumer to be interested in talking to a debt collector. And I think the last big one is payment plan settlement and optimization. Those are in production right now, in addition to a lot of back office things.

Of those, the one that I'm most excited about, and if you think about this, if you think about the data collected pre-charge-off by the big banks, even the fintechs, that is such a hugely rich data set of insight. And I think what I'm seeing that excites me are the smartest players are taking that data set, and they're combining it with post-charge-off data sets. And they're using the combined inferences from both of those very, very large complex sets to figure out a way to get the consumer to maybe think about taking care of their debt.

Now, that's what's exciting of those first three use cases. But in general, these use cases aren't too exciting to most people, and they're probably the least risky in terms of UDAP and wherever the law may take us. The other three use cases or four use cases are, three of them are actually kind of risky, but they get all the attention. One is consumer-facing conversations: AI chatbots, outbound voice AI, things like that. Agent assist is another one where you've got almost whisper coaching, or the collector talking to a consumer has literally... They can be omniscient, in a sense, in that they're whisper-coached on everything that's possibly knowable about that consumer in a way to engage them. That is a big one.

And then certainly skip tracing, that approach, and anything that's risky, I think, that may get the consumer possibly to be tricked or outwitted, like using unlimited computer power to take advantage of the consumer. And I think that gets the regulator's attention. And again, that's why I think the customer-facing thing is a lot slower. And again, in terms of practical deployment, for outbound AI voice bots, which everybody is excited about in the collection space, that's not nearly what you would think it is, largely because in the Supreme Court ruling of Facebook versus Duguid, they were really clear that outbound AI is considered a prerecorded message and requires prior consent.

So, anyone using that is still dealing with this existential threat of a TCPA class action lawsuit.

Oh, and the one back-office one, Alan, that also has me really excited, and that is compliance monitoring, where there was a time in my career - I'm not an attorney - I've got a JD, not an MBA, but I was dragged kicking and screaming to be the chief compliance officer for a large agency. And one of the things we did was we tried to look at every complaint and black box it to find out what went wrong at what point.

Sometimes that led to us penalizing or even firing trainers who seem to be instilling bad habits. The use of AI to draw a line through bad outcomes to the various inputs that led to it - that excites me too - about everybody getting better at operating in a more compliant fashion in the late-stage servicing and third-party debt collection space.

Alan Kaplinsky:

Yeah. So, most of the AI, I take it, that is being used today, it would be considered generative AI, right? Not so much agentic AI.

John McNamara:

Yeah, I think that's fair. I mean, I think that is absolutely fair because, again, there is more risk there. I mean, as an example, if you did have an AI. And there is a very active consumer plaintiff's bar, and I've always mentally, I've bifurcated that. Two groups: there is the noble consumer plaintiff's bar. They're going to sue any actor who they think is abusing a consumer to make the practice stop. And then the for-profit consumer plaintiff's bar is often looking for technical violations, and they're likely to sue a third-party debt collector.

I mean, does your AI understand that a consumer that says you've got a better chance of shoving a stick of butter up a wildcat's behind with a hot poker than me paying that bill? I mean, that's clearly a refusal to pay, but in the universe of things a consumer could say that any reasonable person would interpret as a refusal to pay, I mean, are you sure you're tight there?

Again, I think caution and a history of lawsuits, private rights of action on both TCPA, FCRA, and FDCPA have caused the debt collection industry to be, I think, very steely-eyed and mature in their adoption of technology. They're careful, and they have every good reason to be.

Alan Kaplinsky:

Yeah, sure. I assume that generative AI is used in drafting consumer communications. Am I right that that would be an area that would be suitable for use?

John McNamara:

Yes and no. And this is an area that at Avtal we spend a lot of time thinking about. Disclosures are disclosures; you know that better than anybody, and those are very prescriptive, and the content's pretty buttoned down. Things that aren't disclosure, there's more room to communicate with the consumer. But keep in mind the rails that these communications travel on in the digital world, which is the future, are the publicly switched telephone networks through SMS, MMS, and through email, which is way more effective than most folks realize. And there you're dealing with the ISPs, Google, MSN, and Yahoo. The latter want to cut down on spam. So, any content that your AI generates that looks like spam is likely to never reach the consumer, and therefore, not be effective.

In the case of the former, the networks absolutely do not want traffic on that they dislike. And so content that triggers the likes of Verizon or T-Mobile is likely to get screened out. I guess the rails that these communications move on have a vote. It would almost be like if the post office didn't like the look of your envelope, they would just throw it away. Or if the post office opened up and read your letters, God forbid, that they didn't like the way you worded something, they're going to interfere with that. It's a lot like that. So, I think people are using it. I know at Avtal, we basically craft them ourselves.

And it's not so much the writing them; it's using the big data to see how a consumer engages with that, whether that was effective or not. And again, that goes back to my conversation earlier about using these large, wonderful, rich data sets for servicing to see how a consumer may respond, infer what the consumer prefers in terms of channel.

I mean, we've moved, and this is a wonderful thing, I think, for both regulators, consumer advocates, and consumers, is we moved from this world where a debt collector is going to call, Lord knows how many times, just desperate to get somebody on the phone so they can talk to them until they get some kind of a payment. We've gone to this world where consumer preference is absolutely first and foremost, top-notch.

Because again, voice is dying. It's almost impossible to get a consumer on the phone, and what consumer really wants to talk to a debt collector when they could basically solve it. You can meet them at their place, they can solve it, or take a look at what solutions are available for that debt in their own time, at their own pace using asynchronous communication versus a highly intrusive phone call.

Yeah, I mean, I'm not a fan of AI sort of... And just because I say I'm not a fan of AI writing the content now doesn't mean that I won't down the line, but also let's face it, I mean, when you think about how consumers respond to anything, myself included, if it's a large, dense block of text and I'm not reading it as part of my job, I'm not reading it, period.

Alan Kaplinsky:

Right, right. So, a lot of the conversation about AI focuses on risk. What about AI can actually improve outcomes for consumers?

John McNamara:

Yeah. The layup answer I'm going to give you first, while I try to think of something more insightful, Alan, is there was a time where you might monitor 2% to 5% of calls that your collectors had with consumers, if you were lucky. Now, you have this powerful ability to monitor everything. I'm just talking about talking to the consumer, but we'll talk about how they interact with various emails and texts, but you can monitor that. You can monitor the tone; you can monitor the sentiment; you can look for over-talk, times where the consumer was trying to talk when the agent was, and vice versa. You've got all of that stuff where there's no reason not to do it.

It's funny, when I got into the CFPB, one of my favorite people at the CFPB was Chris Johnson, who ran the consumer response database. We would have conversations, and we would challenge each other, and I was like, and this is going to scare your listeners probably, and I'll apologize upfront, but I thought, "Gosh, you could turn the consumer complaint into almost like a minority report where instead of looking at each individual complaint after the fact, you're looking at it in a streaming fashion where you have streaming data of what's going on, and you're looking at everything in near real-time where you can spot a problem before it happens."

And of course, the CFPB data set is not rich enough to really do that, but a collector using all of their servicing data certainly could. You could probably infer that the text that you sent at 3:00 PM on a Tuesday, I think it irritated that consumer, not always. And of course, consumers respond to texts with lots of emojis, which again, are another rich opportunity to make inferences about how the consumer views that debt.

So, just that idea of taking all of that data and asking yourself the question, "Does that consumer want to engage with me more or less after that interaction? Am I building brand or building awareness with that consumer? Is that helping?" And then the one thing I mentioned that Avtal is the SMS and email, but all of those are to a purpose, and that is to get the consumer to look at a payment portal that was designed by people that think of shopping cart abandonment rate, not collections.

So, the consumer looks at something and says, "Wow, there are quite a few options for me to work this thing out." So, if there's a way for them to pay the debt, they can do it in their own terms and pay it. Or maybe they don't pay it, but you've created a little bit, maybe the budding of a relationship where the consumer distrusts you less, they distrust the communications less, and they're more willing to take a look at things.

Alan, I was talking to a litigation attorney friend of mine, and you know this, and I certainly know this because we're from the industry, but once an attorney has a judgment against a consumer, the clean little secret of this industry is that attorney will take almost anything. If they could meet them outside the courthouse or get them on the phone, once they have the judgment, because judgments are often hard to execute on and costly to execute on, they would take almost anything. Most collectors, trying to help a patient resolve a debt or take care of it, they're usually willing to work with them in a settlement, a settlement over time, payments over time, things like that. You want to make the consumer aware that they actually have some choices.

Alan Kaplinsky:

Right, right, right. Let's get into some of the legal strictures or requirements that apply here, talking about the Fair Debt Collection Practices Act and Regulation F, which you played a large part in drafting UDAP concerns. What exactly do debt collectors need to keep in mind when they're deploying an AI-enabled system? I know it's a very broad question, but I figured you'd have some ideas on that.

John McNamara:

Yeah, no, no, I've got some dislikes there. There are some providers, I think, that oversell AI, and they'll come into you, Alan, and say, "I want all of your data," and with more data sharing is more risk." And again, I firmly believe that when you listen to policymakers and regulators talk about AI, sometimes they're talking more data privacy. And sometimes when people are talking about what worries them about data privacy, they're really talking about AI. So, that give me everything requires some really good questions from the industry player, like why? Because each additional data field presents more risk.

I kicked this off at the CFPB. I initiated our efforts to monitor AI within the markets group of the CFPB because I knew that whether the CFPB was steeped in AI knowledge or not, industry was going to forge ahead. The opportunities were just too great. And so we began engaging with folks from Zest and some of the early providers of AI, the folks at Upstart that were using AI and alternative data. The one thing that stuck out from the ones that I talked to that I considered really, really bright, and they said early on, "It's data governance. It's absolutely data governance."

I heard that repeatedly, repeatedly, repeatedly. And I think that's probably where everyone needs to think. And along with data governance, okay, I've got the data. It's where it's supposed to be. I have a permissible reason to have the data. I worry about that, that at some point, lawsuits will bubble up over... So much of this comes from credit reports.

Do you have a permissible purpose to have the data? Are there disclosure requirements to the consumer when and how you're using the data? And again, like I said, do you know the provenance of that data? Did you go out there and just collect everything and pour it into your database?

In my personal life, I'm okay with that for my own uses, but I certainly want to know about where it is otherwise. It's in the early days of the fintechs where they were offering freemium-type pricing. And again, my group, Lori Udis and some of the excellent people, they're very knowledgeable about how payday lending lead generation works. So, our question for a lot of the early fintechs with the freemium model was: how are you monetizing what you're doing, and do the leads that you collect have a rich afterlife unbeknownst to the consumer?

Where is that data flowing to?

And even if there isn't a law that says you can't do that, should you let the consumer know?

And I see that it's a little bit like that with AI, where all of a sudden everybody wants to train their models, they want to ingest as much as possible, and there can be a little bit of this wild, wild west attitude about it. And again, you want to know what data is being used, how it's being used, where it came from. And at some point, could a state AG or somebody else challenge you that you have no reason to have that data in the first place? I think that probably worries me the most.

As far as actual cases, I think you have to use AI in a risky fashion, like we talked about the AI chatbot that doesn't understand that cantankerous debtor that makes a really colorful analogy of what it's like to get a payment out of them. This one, I think, is worrying because people are coming to AI late in the party. They're late to the party; they're trying to catch up. And again, their agencies sometimes have decades and decades' worth of data that they're tempted to use, or a vendor that really isn't thinking of the long-term best interest of their customers may encourage them to use that data. I'd be like, "Should you?"

Alan Kaplinsky:

So, where do you draw that line, John? I mean, AI needs data, right? That's what makes it work. It makes it more accurate. And you've properly expressed real concerns about some of these companies that will come in and ask you to just dump everything you have, dump it all, and let AI deal with all of it.

Where is that happy middle where you can feed the beast, but don't get in trouble?

John McNamara:

Yeah, no, no. Okay. So, great question. And I really equate it to the use of alternative data, AI, and lending. Rohit Chopra and I, we had a number of things we did not agree on. I did agree on his comment about black boxes in AI. If you're a lender, you can't just say that I initiated an adverse action because my AI told me to, right? You've got to explain it. And this was a question that I had for years for every provider is, tell me about the frontier of explainability, so it goes out this far. And then there is the frontier of what you can do, but you just can't explain how the hell you got it. That frontier is bigger and broader, and the area between those two frontiers is what you would have to give up to stay explainable.

I think AI should be thought of a lot like that. You have all this data lying around, and you've got a subset of the universe of that data that you know you're supposed to have. It's permissible that you have. You having a sense of what you're going to give up. And I think that's where you also don't fall prey to the marketing puffery. I was talking to one customer, and they were talking to a vendor that was trying to give them something to give them the best time to call the consumer, or the patient. But they wanted all of their data. And it's like, okay, just take a deep breath. They're not quantum-ly smarter than you are. And ask yourself, "Do they need their former address going back eight years to figure that out?"

No, no. They need a set of interactions that happen, and even then it's a little bit...

Look at the use case, look at the risk, look at the data necessary, and have enough confidence in yourself to... I refer to it and have, for years, of storing the BS flag and saying, "No, I'm going to give you a necessary subset." And you know what? If I'm going to lose a little bit... I mean, business is risk versus reward. Life is risk versus reward. Is the juice worth the squeeze to take that additional risk to use data? And I think a lot of times you're probably not giving up too much. I mean, I've had some folks in the lending space say that their frontier of explainability and their frontier of possibility - there isn't much room there. So, it's pretty easy. And again, if your frontier of explainability is here, given the development of AI and process improvement, it's going to be higher at some point.

You're going to push that thing out.

Alan Kaplinsky:

Right. One of the biggest concerns that you hear about all the time, and I've experienced it myself just fooling around with large language models like ChatGPT and Grok, is this hallucination problem. Hopefully, as AI, the technology becomes more sophisticated, that will become less of a problem, but it is a problem. People talk about it all the time.

What safeguards should companies put in place to keep AI systems from giving consumers inaccurate information?

John McNamara:

Yeah. Great question. I think hallucinations may be a smaller problem than most people realize. Quite frankly, I think the public loves the idea of AI hallucinating. I mean, it's just fun to think that it's making stuff up, and it certainly has made stuff up. I have firsthand experience with that. I'm going to give a book recommendation, an unasked-for book, and it's a book called Co-Intelligence by Ethan Mollick.

Mollick is a professor of entrepreneurship at the Wharton Business School. It's my favorite book. It was recommended to me by a compliance attorney in the debt collection space, Tim Collins. It's a quick read, but one of the principles: he has got four principles for AI, and I think everyone should adopt them.

One is always invite AI to the table. And that means use AI, get your hands dirty every day. It's cheap, it's omnipresent at this point, but let it do things for you. And you'll be surprised. I mean, to me, my journey with AI, Alan, is like asking my dog to make me breakfast. Each morning I come down, and I look at my beagle chopper, and I say, "Make me breakfast." And of course, with a beagle, it doesn't get any better. Maybe the dog brings you a bone, thinking that maybe this is what this crazy man wants.

But with chat, I use Claude. Claude is my favorite. It's as though each morning the dog gets closer to actually making your breakfast. And eventually, when you start asking the dog to do more, like I really want an omelet with some chives and mushrooms, it's getting there. While we're talking, I'm using DEVONthink, which I'm playing around with, but I'm using Claude to chew through all of my old unreadable files in DEVONthink. It's chewing through that right now to read old PDFs and things to make them indexable and searchable.

So, invite AI to the table, get your hands dirty, but number two of Ethan Mollick's principles: be the human in the loop. And that means that you always have to have a seasoned, experienced set of eyes spot-checking, or looking at things. And again, I think that maybe they spot-check periodically; maybe they spot-check one in a hundred based on the level of risk, and something that's less risky, it's one in a thousand. Something that's really unlikely to go wrong, it's one in 500. I can't tell you how important the human in the loop is.

I have a personal story for almost everything. I was trying to get a new iPhone. So, I'm at the Apple Store at the old Carnegie Library in Washington DC, and this thing is pristine. It has been in a gorilla case since the day I got it. There is not a mark on it, but that guy's AI and his phone is telling me it was Mar and they were going to knock down the resale price, $200.

I literally had to say, I had to be a little bit of a Karen. I said, "Can you bring your manager over because that's wrong?" I literally said, "Be the human in the loop. This thing, look at it. Can you see a defect? If you can see a defect, I'll go with your AI; otherwise I won't." And it wasn't because that AI was getting it wrong. So, that was a really long start, but to be the human in the loop, Alan, is the answer to that.

Since I'm halfway through, I'll give you Ethan Mollick's two other principles because I think they're valuable. One, treat AI like a person, but tell it what kind of person it is. And so if you give AI a sense of what you're trying to find, it's amazing.

This was Gemini AI. I asked it to tell me about a well-known CRM in the third-party declaration space without any kind of a hint of what I was looking for. It came back and said it's kind of a gold standard, and it does this and this and this. So then I said, "Hey, I've heard complaints about this CRM that data is hard to get out of it. The data is like a Hotel California for data where it checks in, but the customer can't get their data back." It was almost like I gave Gemini AI permission to be snarky, and it was like that gossipy friend we all have that just unloaded.

So, it's important to sort of play around with it. Don't treat it like a straight question every time, but give it a chance to add nuance. And then the last one is: assume the AI you're using now is the worst AI you will ever use.

Now, and you talked about this earlier about maybe the mitigation of hallucination problems, although I had a thought this morning while I was getting ready, I'm an inventor. I'm always worried about something, and I thought, "Wait a minute, what if the power consumption of data centers causes so much backlash, political and otherwise, that we basically begin choking off the power for AI so it can't develop?" And then that thought led to me, "What if AI kills Bitcoin?" Because it's like you can have power for AI, or you can have power to mine Bitcoin. Which do you want? I know which one I want.

Alan Kaplinsky:

Yeah, right. Me too. What's the name of that book again?

John McNamara:

Yeah, it's called Co-Intelligence by Ethan Mollick. It is a very quick read, and I'll tell you what it does, and this is what we aspire to do in the markets office of the CFPB is to make complex financial products accessible so people understood them. It makes AI very accessible to the average reader, and it puts it in perspective.

Alan Kaplinsky:

I'm sure it's available on Amazon, I would think.

John McNamara:

It is. Kindle, I think you can even get it on Audible.

Alan Kaplinsky:

Yeah. Okay. So many financial institutions already have model risk management frameworks that they've developed. How should debt collection organizations think about AI governance, and what should they look for when relying on a third-party vendor?

John McNamara:

Yeah. I think one is don't stray too much. They need to understand what they're asking for. A creditor may put very, very high expectations on an agency because they don't understand AI, and they may be limiting them. And quite frankly, they may be increasing risk.

Most notably, and I'm going to veer off AI for a personal rant, many banks and some fintechs don't let their third-party collectors use SMS and email. And it's an example where they're limiting that, but because they're limiting that, they're forcing their collectors to call the consumers more often, which is not a good outcome, and that's likely to lead to more complaints. And quite frankly, as someone that helped with the Reg F, we knew we were taking away the ability to call with the seven by seven rule, but we also tried to give back with a limited content message, and by clearly paving the way for the use of responsible digital.

So, understand what they're doing, think about it, understand how the agency is using the data, but I think they should be more principle-based and less prescriptive. A lot of times, I don't think they have the institutional knowledge to be prescriptive. And sometimes it's a middle manager at the bank asking for it that may not, again, have the universe of knowledge necessary to create that risk mitigation tool.

Alan Kaplinsky:

Okay.

John McNamara:

Since I already ranted about the banks, the other thing they do, I think that basically adds to risk is that a lot of them are on this gross remittance model, where the agencies have to remit 100% of the dollars collected, and then wait sometimes 30, 60 days to get their fee back.

So, banks put incredible capital stress on the agencies that they're putting through risk analysis. And because I'm no longer an agency, I could say this. I could say, "Hey, Acme Bank, don't even talk to me about the hypocrisy of risk analysis when you're literally almost putting me into bankruptcy."

Alan Kaplinsky:

Okay. Do you have anything else you want to... I've got a topic for you to rant about - the CFPB.

John McNamara:

It's one of my favorite topics.

Alan Kaplinsky:

Yeah, based on your years at the CFPB, what do you think regulators expect from companies that are using AI in debt collection?

John McNamara:

Yeah.

Alan Kaplinsky:

I mean, you were there primarily during that period of time. When did you join the CFPB? Remind me.

John McNamara:

Yeah, I started in the summer of 2014.

Alan Kaplinsky:

Okay.

John McNamara:

Largely to replace John Tonetti, who left. He was the bureau's first debt collection program manager. I came in to replace him because I really thought I was ideal. I had a weirdly varied experience in the debt collection space. I sort of walked all the way around the elephant, so to speak.

Alan Kaplinsky:

Yeah. I remember John Tonetti. In fact, he has been a guest on my program a long, long time ago. I mean, you've been there, and then you were there continuously during several presidential administrations.

In answering my question, maybe you have to focus on what it was like under one administration versus what you think it's like today.

John McNamara:

Yeah. Well, I'll tell you what. Richard Cordray was the director when I was hired in, and then we had Mick Mulvaney for a hot minute, and then with Kathy Kraninger, David Uejio for a while, who actually, in my mind, was a really good leader and was not bad. He's like the most underrated leader at the bureau in a lot of respects, and then Rohit Chopra. And then, of course, Russell Vought, and it's hard to compare what's happening now to anything.

I will say I often characterize how I look at something by reason to hope, reason to mope. I think Brian Johnson as being nominated, I am hopeful that as a fundamental shift in the White House's approach to the CFPB. I mean, Brian's a serious policy guy. He did work with Russell Vought, but he also worked with many other, I think, really good policy folks on the House Financial Services Committee. He was our deputy director. I found him to be very reasonable.

I hope, and I'm wrong in my life more than I'm right, but I've been saying for some time that the Trump administration at some point is going to wake up and say, "Wait a minute, the CFPB is not a bad vehicle for an affordability push, and therefore, we don't want to kill it anymore." And let's face it, most of your customers are what I would consider responsible financial services. They've watched, I think with horror, the attempt to kill the bureau has led the states just to go hog wild.

We've sort of accelerated this patchwork of financial regulation and expectations. I know when I was at the bureau, people would disparage me all the time with accusations of regulation through enforcement, but I think with the states, you're seeing regulation through supervision where an examiner sees something they don't understand or don't like, and they begin saying, "Well, you can't do that." I'm hoping the CFPB comes back under Brian Johnson to something that's not unlike the Kathy Kraninger years where the bureau followed the statute, it didn't push the envelope too far, and there are those that would say Richard Cordray pushed the envelope. And yeah, in some cases he did, but I found him to be eminently reasonable. I could actually have a conversation with Richard Cordray about why collectors have a right to collect time-barred debt, as awful as that sounds, but he understood the legalities of it.

I think Rohit was more headline than substance, and I think that caused a lot of problems, and maybe even the backlash the bureau saw when Russell Vought took over.

Alan Kaplinsky:

Yeah. I'm looking at it from an entirely different perspective, but I can't help but wonder why it is that the administration has largely left the FTC alone, although I'm sure they've had reductions in staff, like every other agency, but basically the FTC doesn't look that much different, at least in the consumer protection area, than it looked for the last five, six years.

I think it's because of Rohit rubbing a lot of people the wrong way, and there was a desire to get some revenge after he went to take everything apart.

John McNamara:

Yeah. I think actions have consequences. I think when Rohit announced the overdraft rule the day after he testified before the Senate Banking Committee, and when he refused to resign on the 21st of January, like the heads of other independent agencies, I think that triggered a reassessment of how the Trump administration would approach the bureau.

I think there's an alternate future that we were on that would've led to a Kraninger-like, Todd Zywicki-like, Brian Johnson-like director. And I'll tell you what: if that had played out - would have, should have, could have - but I think it would've been better because the states wouldn't have felt like they have to scramble, and there would've been this entity that was going to do good things. Kathy Kraninger did good things. As far as the FTC goes, I got to say, I mean, I know Chris Mufarrige; he worked with Brian Johnson and Kathy at the CFPB.

I mean, some of the cases they've done from the Bureau of Consumer Protection, I mean, it's hard to argue with them. I like what they're doing.

Alan Kaplinsky:

Oh, yeah. Yeah. No, I agree. Lets crystal ball gaze at this point about the future. What are you anticipating as you look out beyond the end of the Trump administration? Where are things headed, John? And in particular, where is AI headed in the debt collection area? What do you think things are going to be like?

John McNamara:

Yeah, so I am guilty, and there is probably a clip of me saying this somewhere, but maybe in response to all of the puffery around AI, I've said AI stands for almost infrastructure. I'm going to only speak for myself, not you, Alan, but I'm old enough to remember that when ATMs were considered fintech and there was gnashing of teeth and rending of clothes that, "Oh, my God, all the banks will fire everybody, and there'll be no more branch banking."

And of course, that played out. I was the chief marketing officer for a company called LiveOx when the cloud was scary, and people were terrified of the cloud. Now everything is cloud, and it's just infrastructure. Gary Stein, who you should have on your show if you haven't already, Gary Stein was my very able deputy and is an expert on payments. I used to joke that the 1033 rule was on its way to just being infrastructure, basically infrastructure like guardrails for normal business.

I think AI is well on its way to being normal business. I mean, it used to be special. Now, I mean, if you do a Google search, basically it's augmented by AI whether you like it or not.

Alan Kaplinsky:

There is no longer Google search. I doubt it. Whenever I put something in, I get an AI response.

John McNamara:

Yeah, no, no. It's interesting because AI let everybody else catch up to me. I'm a former bill collector, and I was a very good skip tracer, so I could use Google well better than almost anybody else. And I used to be frustrated because they'd say, "I can't find this." I'm like, "Let me see." AI has let everybody become expert at Google searches, so if they're going to find everything that you said in your speeches, and within five minutes, they're going to do it, and you'll have it, and then they'll synthesize it to get a gist of how you feel about any given issue. So, it's already in infrastructure. It's on the way. We won't talk about AI. As a matter of fact, I also predicted that any company that incorporated AI in their name is going to look back in a few years and go, "Were we drinking? What were we thinking? Why did we slap AI on our name? Because that just makes it sound silly."

That's going to have to come off very, very quickly. I see that.

In terms of financial services and lending, I'm hopeful that the debt collection, which has always had not exactly the best reputation, will become much more like consumer self-serve in a marketing approach where it's an extension of pre-charge-off servicing, and there are less complaints. I'm hoping Reg F leads to less lawsuits over technical violations that responsible collectors trying to do their jobs quit being... They quit paying this tax that they pay to the for-profit consumer plaintiff's bar for violations that probably had no material negative impact on the consumer. So, I think we're watching it right now play out.

I'm more pro-AI future than negative. I don't see doom and gloom, but it's a hell of a question, Alan. I don't even know what I'm going to have for supper, much less what AI is going to be in five years. But I think it's going to be one of those things where we think about it less because it's just there in the background always running.

Alan Kaplinsky:

Yeah. If you were giving advice to a son or a daughter who's interested in following in your footsteps to get involved in the debt collection industry, would you say to him or her, "Hey, you better think about that twice because it may not exist five, 10 years from now. Everything could be automated."

John McNamara:

Yeah, actually, I would. There is always going to be room for human agents, I think, for escalated cases, but will the need for vast seas of agents be the same in five years as it is now? No, absolutely not. AI is getting better and better and better at dealing with that.

Oh, the one tip for listeners in the collection space: some states require the announcement that this is an AI bot. I think most responsible agencies I know, they're doing that in all 50 states. "Hey, good morning, Alan. You're talking to Fred. I am an AI bot. Would you like to talk to a human?"

Alan Kaplinsky:

Yeah. I noticed something, not that it's relevant to debt collection, but the other day I was in Manhattan, I needed to get an Uber to take me quite a long distance. In any event, I went on Uber, and it gave me an outrageously high price. I mean, it was like this was to get from an apartment where I live in the Philadelphia area to Newark Airport, and it was something like $450, which it was crazy.

There was in big lettering on the page I was looking at that said, "In determining the price," in so many words, "In determining the price that we're charging you, we have looked at data pertaining to you, and you're consenting to what we have done." I guess it's a kind of surveillance pricing or algorithmic pricing. I was really surprised that there was that kind of a disclosure.

John McNamara:

Wow! Did you take a screenshot?

Alan Kaplinsky:

I should have, and I didn't. No, no. But what I did do is I then went on Lyft, and the price was 50% of what Uber wanted to charge.

John McNamara:

Okay.

Alan Kaplinsky:

Maybe because they didn't... I think it was the first time I ever used Lyft, so they probably had no data on me at all.

John McNamara:

Yeah. And maybe... That's interesting, but what I like is there was competition; you were able to find somebody else at a cheaper price, but that's the kind of thing that I absolutely think that algorithmic pricing or algorithmic pricing based on a surveillance capitalism, I mean, those are sort of triggering terms, I think, for New York City, for a lot of areas to go after because they certainly smack at hitting it at affordability.

If they chill competition and innovation, though, that's what I always worry about: is it's easy to go after them, but is Uber onto something? And were there other cheaper ways to get where you were going? Did you have availability of those? So, did Uber have an unfair advantage? It doesn't sound like they did.

Alan Kaplinsky:

Yeah. I kept thinking about what I could do to outsmart the system, and I was thinking maybe what I need to do is do a search on Uber every day, do at least one search with no intent of using it, keep doing it and doing and doing it. And then maybe I'll build up a different kind of database where they'll say, "Oh, he drives a hard bargain. I'm going to give him a good deal and stuff."

John McNamara:

And that sounds exactly like a job for an AI bot to do that for you. I think that's the other thing is I think it'll spur creativity because how many ideas do you have like that, but you just don't have the time or the capability to pursue them? It's almost like AI... I mean, AI should be taught in first and second grade; everybody has the ability to follow up every thread and every hunch.

What if everybody that heard a conspiracy theory had the ability to go chew through all the data to say, "What's the likelihood that this is complete and utter baller dash?" That's where I tend to be more optimistic. I worry about consumer advocates sometimes. They seem to be genetically predisposed to hate fintech and hate anything new. So, part of the market's job, I think, was to try to translate as much as possible so that they would realize things like earned wage access actually have a benefit in some cases; buy now, pay later is not necessarily evil.

That is absolutely fascinating. And again, back to Uber, this is probably apocryphal, but I thought I heard somewhere that in the early days, the app grabbed your battery life in your phone when you ordered, and so it knew that if your battery was almost dead, it could charge you more. I mean, if that's true, I mean, it's interesting, but also evil.

Alan Kaplinsky:

Yeah, yeah. Agreed. Well, John, we've come to the end of our program, and boy, time flies.

John McNamara:

It sure does.

Alan Kaplinsky:

Yeah. I want to thank you again very much for joining us today and sharing your thoughts about AI, how it's being used now, and where you see it headed in the future. This was very, very illuminating for me, and I'm sure for our listeners. Again, thank you.

John McNamara:

Alan, thanks so much for having me. That hour flew by so fast. This was fun. If you want a second bite or a third bite of the apple, you know how to get me.

Alan Kaplinsky:

I know how to get you, for sure. And I want to thank our listeners also for taking the time to listen to our program today, and I hope everybody enjoys the remainder of their day.