**Podcast:** *The Consumer Finance Podcast*
**Episode:** AI in Collections: Balancing Automation, Compliance, and Human Judgment
**Host:** Stefanie Jackman
**Guests:** Aman Mender and Adrian Ferrante-Bannera
**Aired:** November 5, 2026

**Stefanie Jackman (00:05):**
Welcome to *The Consumer Finance Podcast*. I'm Stefanie Jackman, a partner in Troutman Pepper Locke's Consumer Financial Services practice, and I'm filling in for Chris Willis today, so feel free to send your feedback on how my hosting compares to his. Today we're going to be talking about use cases for AI, artificial intelligence, in servicing and collections. And I have two gentlemen from Corafone joining me, Aman Mender and Adrian Ferrante-Bannera. But before we jump into that and I ask them to introduce themselves and talk a little bit about Corafone, I want to remind you to visit our newly created Financial Services Hub, which is replacing our blogs that many of you may be familiar with, at troutmanfinancialservices.com or consumerfinancialserviceslawmonitor.com. And while you're at it, please be sure to check out some of our other great podcasts: The *FCRA Focus*, *The Crypto Exchange*, *Payments Pros*, and *Moving the Metal*. All of those are available on all popular podcast platforms. And speaking of those platforms, if you like our podcast or you like me as our temporary host, let us know. Leave us a review on your platform of choice and tell us how we're doing. As I said, today we're going to be talking about using AI in servicing and collections. And joining me are Aman and Adrian with Corafone. Thank you both for being here.

**Aman Mender (01:28):**
Thank you for having us.

**Adrian Ferrante-Bannera (01:29):**
Yeah, thank you.

**Stefanie Jackman (01:30):**
Before we jump into the questions, maybe just tell our listeners a little bit about what is Corafone? What are you all doing?

**Aman Mender (01:36):**
Basically, Corafone, we focus on building AI collectors for creditors and debt buyers and agencies and helping them run collections more efficiently. We have AI agents that essentially do the traditional job of calling consumers or taking their calls and helping them resolve their account or get current. In addition to that, we also have our own agency where we're looking to automate everything as much as possible with AI. So a little bit broader sets of use cases there.

**Stefanie Jackman (02:14):**
And it's interesting, I've been hearing a lot about AI and chatter. We were talking about that as we were preparing for our podcast. My sense is that AI adoption in collections has been moving faster in some instances than the governance conversation around it. You all shared with me that a 2025 industry survey showed only 7% of firms were reporting no AI plans, which is great because it's down from 27% a year earlier. But we're still working to understand how to oversee, monitor, and work collaboratively with our AI tools. And where I'm hearing a lot of public discussion is whether AI can hold a collection conversation and what that looks like and how you supervise using an AI collector in that capacity. Like, who sets policy? How do you stop it mid-campaign? Who's responsible for that? What the system can do on its own, what it can refuse to do, and then ultimately what record we have after. So let me put this to the two of you. Where does the line sit between what an AI collector handles and what a human handles, and who at the company who's deploying this technology decides where that line is?

**Adrian Ferrante-Bannera (03:23):**
First, if we rewind a bit, I can introduce ourselves so that we can share where we come from. Aman and I, we're both the co-founders of Corafone and we have experience building AI agents since it was a thing. And it started with voice AI very early on, before the models were good. And we struggled a lot with different languages, different use cases. Where we see AI currently is not where we started, let's put it that way. We're not really from the collections industry. We've worked with industry players to learn and build the moat around the actual collections experience. But at the end of the day, it is collections. What we're trying to optimize is collections with AI. We're not building a voice AI for collections. We're building AI collectors. And that use case where you started off, what could AI do that a human wouldn't be able to do, has moved a lot since we started. It started off with basic tasks like presenting the plans. Before AI, we had the IVRs, and then it slowly, as the technology improved and the models got better and also the experience within the field and the adoption rate of the consumers got higher, we started seeing a shift where AI is handling more and more complex calls. And then there's a line where it gets still in the uncharted territory, and that is a little bit up to the client where that risk appetite lies. Do we want to allow them to handle a bankruptcy case? Do we want to allow them to handle a debt settlement case? Or do we want to have that in a more controlled narrative with humans?

**Aman Mender (04:51):**
And typically when we work with someone, we start by understanding their consumers, understanding their operations, why people end up in collections, all those things help us actually build those first agents. And typically we often refer to those agents as sort of junior agents, and we typically have them do something a junior agent would do. So that's typically where we start. And so as Adrian said, when it comes to those more complex cases, it's really about how much the company in question wants to automate. And there's of course natural limits to where you need full human judgment involved. So

with AI, the way we do it is kind of you put some rules in place and you get that human judgment at the architecture level, and then you scale it up and have AI enact that. But if you want individual human judgment, then you have to get a human involved. Typically that's like a dispute or whenever someone wants a human or, yeah, as Adrian mentioned, beyond the bankruptcy, there's a fact-gathering part of it, but then there's also, okay, what's going to happen next?

**Stefanie Jackman (05:55):**
I hear you that where the line is drawn is really a matter of a different client's preferences, risk tolerance, et cetera. But it sounds like a lot of your clients that are using your technology are still keeping some humans in the loop. And I agree with that. It's advice I'm often giving is that these tools are not ready to just be set it and forget it. And regulators certainly don't think they are. But what does having a human in the loop actually mean day to day? Who's watching? What are they watching? Who's stopping the campaign? And Aman, maybe you'll start with this and then, Adrian, give us some thoughts if you have anything to add.

**Aman Mender (06:34):**
Typically there's a collections manager or some managerial role that will start a campaign. They will look at sample calls, they will look at aggregated data, they will look at automated QA. Basically, you start working at a more data-driven level where you get insight, but at a high level, it's kind of also like they get a sense for it. Like, what are they watching? Is the number of transfers going okay? The people actually on the floor are also a big indication. Like if they're getting transfers and they're saying that they're discovering something, that might be a flag. It might be the analysis of the calls running in. So typically there's a manager who oversees operations and this becomes another tool in their arsenal to run their collections operations.

**Adrian Ferrante-Bannera (07:19):**
There's also the fact of capability. It's not just about the risk tolerance. It's also about what have we done before that we can now offer, what have we learned? Because this is an iterative process where we wouldn't cover something when we started, but then we've learned, okay, this is a data point that we see a lot of people doing this. We build out the complexity of the agent and to handle those cases as well. Now this is an option to choose from, right? And one of those examples is really for example, the data that we're gathering. This is also one element with the human in the loop. What does the human do with this data? And right now, we're automating it. So for us, we have the data, we can do, depending on what happens in the call, we can choose which workflow we work. But the human element to this is like, okay, we understand this person might be willing to pay, might want to pay. Can we do some point calling? There's a human in the loop element there as well. What you decide, what's the next action basically for this account? And then of course, there's accepting the transfers, being there when we transfer.

**Stefanie Jackman (08:15):**
That's a great point, Adrian, that you're needing to have people know what to do with the data that's being gathered from these interactions. Another thing I was thinking of is what about when somebody, you're using an AI collector tool, like what Corafone has and others as well, and a consumer doesn't always announce like a cease and desist or a hardship in the words of the statute, right? They do not say, "I am refusing to pay this debt ever and please regard this as a full cease and desist" or something like that. So I'm curious, how does an AI collector recognize a cease and desist request or a stop contact request or a hardship request when they're talking about it perhaps in a more roundabout way? How does the AI collector recognize that? And how should a company approach this when the signal is ambiguous that's coming from the consumer?

**Adrian Ferrante-Bannera (09:06):**
There's no one way to do it, let's put it that way. But essentially there's the safe harbor approach where you flag everything suspicious and you kind of go on the safe side. For example, one of our clients would typically do some dialing and then no matter what happens to these accounts, they don't want to dial them again for seven days. Similarly, when there's an ambiguous do not call or cease and desist, one approach that we have is throughout the call, learn more. If there's a dispute or we suspect a dispute, the AI might ask, "Are you saying you're disputing this debt," to make it explicit. So this implied, ambiguous data point is very hard to act on. But making it explicit, if there's an implication there or, sorry, implied meaning there, that is one way to work around that, where the AI just states it and asks it. Another thing is also understanding the tricks and the hoops to actually keep the call going. So someone says "do not call" means, yes, we don't have permission to call again, but the AI might ask, "Can we continue this call?" And you still have that permission if they say yes. So there's nuance to what we do with the data point after, but there's also how we handle it during the call that make it more or less ambiguous.

**Aman Mender (10:12):**
So that's one part, engaging with the consumer, clarifying what they mean, especially if we're talking about putting AI into consumer contact. When it comes to things like cease and desist, it's also about the data you have. So having the accurate data labels and training around, okay, what is cease and desist and interpreting. So we have a benefit that LMs are actually pretty good at interpreting intent, but it needs the right context for what is something and what something else is not. So "don't call this number again" or "stop calling me" or "cease and desist," they're all in the same category, but how do you differentiate cease and desist from stop calling? And what we see a lot of people do and what all serious AI companies are doing is investing a lot of time and resources into building that data set that defines what something is and something is not. That's also what we spend a lot of time on, and that's what you really want to get right. And in that way, you can turn every call into an asset that makes your AI better all the time. So that's sort of the flywheel that you need to set if you're enacting AI and you want AI to

recognize different things, is to first do that foundational job and then have a flywheel where you turn data into more and more accurate AI.

**Stefanie Jackman (11:30):**
In your view, and I'll start with you, Aman, are AI collectors more compliant than human collectors?

**Aman Mender (11:37):**
In some ways, I think they are. They are more consistent. But a human collector can be more discerning and can be more flexible. It's more about pros and cons. So at some things, if you have 100 human collectors that you task to follow a certain procedure, you're going to have more mistakes than if you had an AI collector, or an AI for that matter, do the same job. But of course, the more you have to introduce nuance and judgment, then it becomes harder to have those right instructions and guardrails and guidelines. It's kind of like the more complex the conversation is and the more judgment you need, then a human is going to probably be more compliant because it could deal with all of the different ways it could go. But the most common reasons of compliance breaches, one could be just the data, having inaccurate data, but another is just human error. And human error comes from being tired, being exasperated, being frustrated, being under pressure, dealing with a lot of angry consumers. All of those things typically go out the window when we're talking about AI because it does not know one conversation to the next. It's just going down the same script that you gave it. So in that sense, it's more consistent. It's limited, but it's more consistent.

**Stefanie Jackman (13:05):**
It doesn't take things personally, basically. It doesn't have a bad day.

**Aman Mender (13:07):**
Oh, yeah.

**Stefanie Jackman (13:09):**
That makes sense. And that's what I'm seeing with a lot of clients is that those that have started playing with these technologies in this context, they're really helpful for things that are not discretionary, right? They're really helpful for handling calls where consumers want to pay or expressing a willingness, just kind of straightforward, easier calls. And as it gets more nuanced, that's when sometimes the models will be programmed, if you will, to say, "This might be something that's more helpful if I transfer you to a different team member," or something like that. And that gets to, Adrian, I'll direct this one to you. People have expressed concern about friction in using AI in the collections process. Do we tell consumers that they're dealing with AI? By the way, big fan of disclosure, whether the law requires it or not, in my view, for a variety of reasons. Do we make it easy, and how easy, to stop dealing with AI? Again, in my view, you make it as easy as possible. And frankly, some roundtables that I have been able to attend and just hearing what I predicted, just, I guess I got lucky, was that the more you

don't force people to engage with this technology, when they do and it goes well, it'll increase adoption. That's actually what a lot of companies that are doing disclosure, are making it easy to get to a live agent, are seeing. So that's encouraging. But in your experience, Adrian, how are your clients or the companies that you're seeing use your technology and others, how are they handling that disclosure? Are they making it? At what point are they making it? What is that disclosure that the consumer is speaking with AI?

**Adrian Ferrante-Bannera (14:46):**
It's a better question than maybe you realize because that's one of the things that we played around with the first months when we did outbound. How do we keep people on the phone? How do we increase the RPC rate? And we noticed moving the disclosure and making the disclosure different, we could increase our RPC rate by 100% almost. It was close to 100%. The trick is really, you don't want to make it a big deal. Don't make a big fuss that this is an AI. Disclose it, yes, but you can say, "Just so you know, I'm an AI agent. Look at me." Or you can say something like, "I'm an AI agent, but you still owe money. Let's keep this call going." I'm not saying that's exactly how we say it, but you get the point. The AI agent is kind of like a brushed disclosing part. Not all states require this at the moment, but we like to be on the transparent side, so our clients are as well. But that is one of the things where you can play around with where you put it. You can play around with how you say it. You can also play around with the words. "Digital assistant" works really well. We've had "digital advisor," "associate," "AI agent," and there's really no big difference between those. As soon as a human realizes it's an AI, they typically churn. So we want to keep the focus out of that and talk about something else.

**Stefanie Jackman (15:55):**
Aman, what are compliance officers that are using these types of technologies seeing that they couldn't see before in exercising their roles to oversee the compliance of the outbound and inbound call handling? And how is that changing the job of a compliance officer when you're starting to use AI, in your experience?

**Aman Mender (16:16):**
One thing that they are seeing is the ability to interpret these conversations at scale and to be able to investigate the things that would typically take a very long time to understand, to be able to highlight what they want to look at instead of going through a lot of different scenarios that are not that relevant. And also when you start using AI for collections and for calling, you're getting all this transcript and this information at scale that you can then interpret. So typically, I think the tools around this have become a lot better than they used to be. Before you had to rely more on word associations and interpreting what happened was a lot harder. But now you can have some very broad flags and that lets you see what you should look at. You can look at, give some kind of human judgment on that and then just focus on those calls or those interactions and scenarios that you need to. It brings up a lot of new challenges. Who's responsible?

What happens if it says something wrong? And also, what can we say? What does a disclosure need to be? I think this is definitely going to keep a lot of people preoccupied for a long time.

**Stefanie Jackman (17:35):**
I definitely agree there is a lot to be learned and done. But let me conclude by asking each of you one last question. I thought let's close the episode with some candor. What is an AI collector still worse at, including when on your platform, than a competent human collector? Aman, let's start with you. What would you tell a firm not to automate yet?

**Aman Mender (17:56):**
I would tell them not to automate the hardships, the people who are dealing with disputes or anything to sort of legal risk. Anyone who feels defrauded or badly treated, all those things should not be with AI. Anyone who's really, frankly, really upset also, I think, needs that human touch. We have lots of interactions where AI calms them down, but they're already at a point that you don't want any of that frustration that can occur if a conversation is not going as it should.

**Stefanie Jackman (18:31):**
Adrian, what is an AI collector still just not good at?

**Adrian Ferrante-Bannera (18:35):**
One thing that you as a consumer feel is authenticity, right? So genuine empathy, even though I can sound empathetic. And yes, it does help. We have a lot of cases where we've talked consumers down from being agitated to calm. We have a lot of cases where they've even thanked AI for listening and for hearing their complaint about the lenders. And we also have cases where we have for our creditors that we work with, for example, where they would say, "I want to talk to a human because I can't deal with this," or some reason. And then the AI keeps talking and five minutes later through the call, they actually disclose their dissatisfaction. So there are those cases. But what I would say is kind of the limit of AI right now is how do you deal with those dissatisfactions? I personally would want to have someone on the other side treat me as a human. And AI is great at being polite, great at following the instructions and doing what it's told. But there are some times where you just have to make exceptions. And it's an oxymoron as well. If you want to have a consistent AI that makes exceptions, you can't. You have to choose. So you want to have a consistent AI. When there's edge cases or exceptions, you don't want to have AI.

**Stefanie Jackman (19:43):**
Yeah, it's still a machine at the end of the day. And sometimes you do need that human touch. I want to thank both of you very much for taking the time to speak to our listeners today, talk about how you're seeing these types of technologies, including your own, being used, and really appreciate it. So thank you for joining us. This was a really

important discussion about where things stand and, frankly, where they're going with AI, because it's here to stay. And also thank you to our audience for listening as well. And as I said at the top of our podcast, don't forget to visit our new Financial Services Hub at troutmanfinancialservices.com as well as consumerfinancialserviceslawmonitor.com. And while you're at it, why not just visit us on the web at troutman.com and add yourself to our Consumer Financial Services email list? Then we can send you more announcements about podcasts with content like this and a broader array of topics, as well as copies of alerts and advisories that we release from time to time and invitations to occasional industry-only webinars. And of course, continue to watch your podcast feed on your favorite podcast provider every Thursday afternoon for a great new episode of this show. Thank you all so much for listening, and Chris will be back at your next regularly scheduled podcast.

---

*Copyright, Troutman Pepper Locke LLP. These recorded materials are designed for educational purposes only. This podcast is not legal advice and does not create an attorney-client relationship. The views and opinions expressed in this podcast are solely those of the individual participants. Troutman does not make any representations or warranties, express or implied, regarding the contents of this podcast. Information on previous case results does not guarantee a similar future result. Users of this podcast may save and use the podcast only for personal or other non-commercial, educational purposes. No other use, including, without limitation, reproduction, retransmission or editing of this podcast may be made without the prior written permission of Troutman Pepper Locke. If you have any questions, please contact us at troutman.com.*

*DISCLAIMER: This transcript was generated using artificial intelligence technology and may contain inaccuracies or errors. The transcript is provided "as is," with no warranty as to the accuracy or reliability. Please listen to the podcast for complete and accurate content. You may contact us to ask questions or to provide feedback if you believe that something is inaccurately transcribed.*