

Consumer Finance Monitor (Season 6, Episode 34): A Look at the Growing Use of Generative Artificial Intelligence and Chatbots in Consumer Financial Services, with Special Guests Ron Shevlin, Chief Research Officer, Cornerstone Advisors, and Reggie Young, Product Counsel, Lithic

Speakers: Alan Kaplinsky, Ron Shevlin and Reggie Young

Alan Kaplinsky:

Welcome to our 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. I'm your host, Alan Kaplinsky, I'm the former practice group leader for 25 years, and now senior counsel of the Consumer Financial Services Group at Ballard Spahr. And as usual, I'm very pleased to be moderating today's program.

For those of you who want even more information, don't forget about our blog, consumerfinancemonitor.com. Yes, it goes by the same name as our podcast show. We've hosted our blog since 2011 and there's a lot of relevant industry content there. We also regularly host webinars on subjects of interest, to those in the industry. To subscribe to our blog or to get on the list for our webinars, please visit us at ballardspahr.com. And if you like our podcast show, please let us know about it. You can leave us a review on Apple Podcasts, Google, Spotify, or whatever platform you use to obtain your podcast. Also, please let us know if you have any ideas for other topics that we should consider covering or speakers that we should consider as guests on our show.

Well, let me give you a little bit of background about the topic we're going to talk about today, and then I'm going to introduce our two guests who are very, very deep into this subject that we are going to explore. And the subject, once again, is artificial intelligence, generative artificial intelligence, or we'll frequently refer to it as AI. And in particular, today, we're going to talk about chatbots and what their use case is in the banking and consumer finance industry. And we're going to talk about the benefits of chatbots, we're going to talk about the legal risks involved in using them, and what the future holds in store.

Now, this is not the first time that we've had a program on AI, and even the first time we've had a program on generative AI. Those of you who follow our show regularly will remember that, on April 20th of this year, I had, as my special guest, Alex Johnson, the founder and author of a newsletter called Fintech Takes. And the title of that podcast was, A close look at generative AI and what it means for the consumer finance industry.

And one other thing I want to mention before I introduce our guests, and that is, on June 28th of this year, our firm launched an AI task force to advise our clients on opportunities and risk mitigation. This is a multidisciplinary AI team that cuts across, literally, every practice group in our firm, including consumer finance. And if you're looking for more information about what our task force is up to, you can either contact me at kaplinsky@ballardspahr.com. Or you can go to the Ballard Spahr website and you can look under news and press releases, and you'll find the June 28th press release.

Now, let me introduce our two speakers who are truly experts in this area we're going to be talking about. First of all, I want to introduce you to Ron Shevlin. Ron is the chief research officer at Cornerstone Advisors where his research on banking and fintech trends helps shape the direction of financial institutions and fintechs. He's the author of the FinTech Snark Tank on Forbes and host of the What's Going On in Banking podcast. Ron is ranked among the top fintech influencers globally, and he is a frequent keynote speaker at banking and fintech industry events.

Ron, a pleasure to have you on our program.

Ron Shevlin:

Thanks, Alan. It's great to be here. Appreciate this opportunity.

Alan Kaplinsky:

Great. And now, let me introduce to you, Reggie. Reggie Young is product council at Lithic, which is a card issuing platform that powers car programs for fintechs like Mercury and Novo. Before Lithic, he worked for the online business lender called BlueVine, and before that, he worked at law firms. He's the host of Lithic's podcast called Fintech Layer Cake. I love that name. And writes a Fintech regulatory newsletter called Fintech Law TL;DR.

Reggie, thank you very much for being on our program today.

Reggie Young:

Thanks for inviting me on. I'm a big fan of Ballard's Consumer Finance Monitor. And your email updates, I think I read them more regularly than the CFPBs. So, big fan of everything you're doing.

Alan Kaplinsky:

Well, thank you very much. Appreciate it a lot.

So, let's dive into the subject. And as I said, we're going to do some drilling down. But before we do the drilling down, we got to do a little bit of level setting. And even though I'm sure a lot of our listeners know, by this time, what AI is, and may even know what generative AI is, let's lay that foundation. We're going to start with Ron.

And Ron, the basic question is, what is generative AI? How is it different from other AI tech and how can it be used?

Ron Shevlin:

Great place to start, Alan. Thanks.

I think the easiest way to describe what generative AI is in context of other forms of AI like machine learning. Basically, generative AI is a type of artificial intelligence technology, and in fact, maybe even more specifically, a form of machine learning. The big difference between generative AI, and say, other forms of machine learning is, really, in terms of the output, the type of output that generative AI can produce. Typically, machine learning is producing data or content, but generative AI produces a wide range of output including text, images, audio, can produce synthetic data, can even produce programming code. It's the output that differentiates the generative AI from other forms of machine learning or other forms of artificial intelligence. And so, therefore, it opens up a really wide range of use cases. I have seen banks that are now exploring ways of using generative AI to create programming code, marketing copy, generate legal contracts, analyze fraud data and fraud trends, automate email responses, automate customer interactions to do things like proofreading.

So, the list of use cases for generative AI is pretty wide. And I think most importantly, from a financial institution or fintech perspective is that, it cuts across a lot of the various functions within the organization. Whereas, other forms of AI tend to be a little more specific to various aspects of the institution or the fintech. For example, RPA, which might be more about mortgage or mortgage processing or things like that, or automation or chatbots and conversational AI, which might be more specific to customer support departments, things like that. I know we'll talk a little bit more about that. So, I think the easiest way to think about the differences of generative AI from others is in terms of the forms of content, and again the super wide range of use cases that can be applied here.

Alan Kaplinsky:

Yeah. Reggie, you're the lawyer here. Ron's not a lawyer. What are the legal risks and the benefits of generative AI?

Reggie Young:

Yeah. I think a lot of the benefits are just around operational efficiency and things you can unlock with fewer resources, less like legal specific but more what can it power legal teams to do. So, there are some pretty amazing word plugins now for contract drafting and analyzing that.

I remember, when I was at firms, there was this dream that we'd have a word plugin where you just press a button and analyze the contract. But it wasn't quite there. And now, the ones I've seen recently, they're very, very much there and they can plug in a clause very easily, if you need. So, there's some efficiencies there. I know some lawyers will use it to use like a ChatGPT tool to draft a first response and negotiations because it's always easier to edit something than to write it from scratch. So, I think there's some interesting benefits there. You can generally do more with less. There's a lot of folks who are interested in executive assistants that are effectively AIs.

With all of that, comes risks, I think. There's some kind of non-legal risks if you're using, for example, written tests in hiring folks and screening them. I know some engineers are having troubles now because engineering candidates can just ask ChatGPT to write the code for their interview, for their take home tasks. So, they're kind of having to change how they're analyzing prospective employees. There's also potential for better fraudulently written emails for fraudsters that want to phish and whatnot. The emails can get better because they can ask ChatGPT to write them, and there's going to be fewer typos, fewer grammar errors.

Focusing more on legal risks, I think there are four big ones. The first one is, just the fact that these AI tools can spit out inaccurate and misleading outputs. So, potential youth apps, unfair, deceptive acts and practices. They will sometimes hallucinate the phrase that's used to refer to, just making things up or creating things that don't make sense. I think, about a month ago, there's a lawyer in the headlines who submitted a draft to a court that cited non-existent cases because they asked ChatGPT to do some case research. So, there's some definite problems there.

I think I've had an interesting experience recently where a colleague asked about the Durbin amendment, which I'm going to guess, most listeners are familiar with, but in case they're not, it's the law that caps the fees that larger banks can charge on debit cards. And a colleague asked, out of curiosity, whether it applies to both consumer and commercial cases. And it applies to both, but there's a common misconception that it only applies to consumer. So, I was curious and asked ChatGPT. And sure enough, it told me it only applies to consumer, which isn't correct, but it is a common misconception, that AI tools are often based on scraping what's out there. So, it makes sense that these AI tools might regurgitate what folks have commonly misconceived online.

So, I think, again, first big risk is kind of inaccurate, misleading outputs. Second, big area of risk is privacy and confidentiality. Some companies may want to upload their Slack history to ChatGPT, to generate FAQs about the company's history. Tricky thing there is, you're going to be uploading sensitive information that may be subject to confidentiality provisions, maybe subject to privacy laws. I think a good example is California's privacy law applies to employee information. And so, if you're uploading your Slack history, you better have given your employees a policy that says like, "You can do that", because they're their names, which are going to be in that Slack history, are covered by California's privacy law. So, I think some risks there, making sure that companies are doing correct privacy standards.

I think the third bucket is IP issues. There's been a lot of headlines recently about Getty Images, and other artists suing open AI and other AI companies for scraping the web to build their largest language learning models. I actually think, as an in-house lawyer, I think the more interesting IP problem from my point of view is that, there have been a few cases saying that, under US patent and copyright law, if something is solely created by AI, you don't get a patent, you can't copyright that, you can't protect it. And so, maybe you can imagine an engineering team working on some interesting problem and asking ChatGPT how to solve it, and if it comes up with an interesting solution that's potentially patentable, that would be potentially patentable if the human did it, you can't undercurrent US case law, you can't protect that. So, I think that's a risk to that I always warn our engineers and other folks about.

The last big risk I'll mention is just bias and discrimination. This is big, especially in any kind of credit fintech. I think about it from a credit card perspective, just because I'm in the card space. But obviously, there are laws applied to housing and in other areas, generally, that prohibit discrimination. I look at Regulation B, the Equal Credit Opportunity Act that prohibits discrimination in extending credit. There was a big CFPB release... I actually wouldn't call it big release. It was slipped out in

one of their blogs, I believe, saying that, "Under Regulation B, you have to be able to explain why somebody is rejected for credit." And if you have a black box AI and you can't actually explain that, CFPPB has said, "That's not a defense." For compliance, you can't throw your hands up and say, "We didn't know, it was too complicated. It was a black box." So, I think that's a big risk.

But to kind of circle back to the benefits, I think there's a really interesting benefit, potential benefit, along that bias and discrimination line which is that, these models are ultimately based on massive amounts of data they've scraped. And so, if you have a bias issue, you can feed your AI's better data to learn from. And while you have to be able to identify the bias issues, it's probably easier to get an AI to update their bias and their beliefs. It's probably a lot easier to do that than a human, like a human psychology, to be able to identify bias and be able to change. That's a very long process if it can be done at all. And so, I think that's actually a potential benefit that doesn't get enough attention.

But yeah, there's plenty of other risks and benefits, but I think those are some of the big ones.

Alan Kaplinsky:

I guess, the other thing I would add, and I'm glad you identify fair lending because in our world, I think if you were to ask Rohit Chopra, the director of the CFPB, he would say, that's the biggest concern that he has.

But the other concern that I have is that, generative AI is probably never going to be a hundred percent accurate, and what it tells you or what documents that it might generate. And a lot of things in our consumer finance world require pretty much a hundred percent accuracy. Sometimes, there might be a tolerance on the annual percentage rate that you disclose. But say, you were asking the computer to generate truth in lending documents and if the APR was off, was actually turned out to be higher than what was disclosed and it wasn't covered by that tolerance, you've got a violation.

And I don't think it would be any defense to say, "The computer did that" or "Generative AI did that, I didn't intend it." That wouldn't help you.

The other thing is, really, more of a question I have for you, Reggie, and that is, you gave the example of, you asked ChatGPT something about the Durbin amendment. What is it? And it got it wrong. Now, how do you correct that? Can you input information into that database saying, "You're wrong. Durbin amendment applies to consumer and commercial credit card transactions or debit card"?

Reggie Young:

Yeah. Definitely, in OpenAI, ChatGPT's case, they have a kind of, you can flag answers that need improvements. I haven't looked too closely, but there's a pretty easy button where you can submit. Whether that gets addressed, my guess is, correcting Durbin amendment is probably not the top of the list for open AI right now, but they definitely have avenues to do that.

Alan Kaplinsky:

Okay. So, I have a question for both of you, but let me go to you, Ron, first, and then have you respond, Reggie. And that is, how should banks and fintechs go about deploying generative AI tools today? What should they be doing?

Ron Shevlin:

Yeah. This is going to be very dependent on the type of the company, type of bank, size wise, and from a fintech perspective, too, because so many banks below a certain asset size are reliant so much on their vendors, whereas some of the largest banks have large in-house teams to do it, to build a lot of their technology. But I think it's important for both the fintechs and the banks to recognize that, while there are generative AI tools out there like Bard from Google and ChatGPT from OpenAI, from a pure banking perspective, pure financial services perspective even, there aren't a lot of vendors out there providing specific tools. And when they do, like Bloomberg has their own version of a generative AI capability. Kasisto hasn't come out with KAI-GPT. These still have limited applicability in a lot of banks and credit unions and fintechs as well, so I think the way to start is by really identifying the opportunities in the institution for productivity improvement.

I think you kind of nailed it, Alan, when you honed in on the transparency aspect of this, as well as the case that things get, are necessarily wrong. So, what it really means is that, this is not a technology that is going to necessarily automate and replace people. I think the big opportunity is, for this technology to accelerate and enhance the productivity of a lot of people in the organization, and people who typically have not been impacted by process improvement and process automation before. So, it's the creative people, it's the marketing copy people, it's the legal people. Not that you necessarily think of yourselves as creative, but yeah, you write contracts and there's a creative aspect to it.

So, to me, I think the place to start is by looking at and basically tasking the various leaders of the functions in the organization to experiment and explore with ways that they can use the technology to improve the productivity of the people in their organizations, which is a very different approach than we've usually had with technology, which is telling the CIO and the IT organization, "Go fix something, go create something, go build something." Very different, this is more about individual creativity and productivity. I think it's a very kind of different way to approach the tools and technology than other types of systems and technologies that the banks and fintechs use.

Reggie, feel free to argue with me too on this.

Reggie Young:

Yeah. I think in terms of how to actually roll out and deploy, and how banks and fintechs should think about it, I think there's a big question of whether you even need it and whether you should apply AI in the first place. And that's, I think, going to depend a lot on the product and use case. Like, we were talking about bias and fair lending. You probably don't want to use AI in underwriting. And if you think that's a good idea, you may want to go talk to a lawyer and have them explain why you shouldn't do that. There's a reason the industry is kind of focused around FICO and cashflow for underwriting AI, is generally going to get you in some pretty dicey territory. So, I think there's that initial question of like, what do you actually want to use it for?

And then, I think, tactically, you're going to want to start out checking the terms of use of what the AI tool you're going to want to set, establish policies and kind of set rules for employees, like don't feed it confidential, sensitive information, don't feed it customer PII. And I think the thing I haven't heard talked about a ton is just, you almost need an audit function.

To your earlier point, Alan, these things will never be a hundred percent accurate. And so, in a sense, if you want to use AI, you almost need an AI compliance function within the company that looks and reviews the outputs regularly, audits them annually or quarterly, whatever you need to make sure they're still accurate. I think AI is great. Putting my lawyer head on, you probably still need a human in the loop, at some point. Instead of having humans do a hundred percent of that, whatever, manage those support tickets, maybe it's only 2% auditing regularly, but you still want that oversight.

Alan Kaplinsky:

Right. Okay.

Well, I think we've laid a pretty good foundation now, so that we can dive into the subject of chatbots. And I'm going to ask you, Ron, what's the business use case for a bank or fintech that wants to add a chatbot today? And in other words, why do it?

Ron Shevlin:

I would venture to guess, and then I have done some research around this, so from a quantitative perspective, I can tell you that many of the... Both banks and credit unions that I have surveyed and talked to, get into chatbots and see it basically as a customer support productivity enhancement tool, that they're using it to kind of deal with the easy questions that consumers or their customers or members have, to ease the load of questions on their customer support staff. I think that's a really, really narrow way of viewing the opportunity for chatbots, because many of the more sophisticated companies that I've talked to, that are using chatbots... And we should get into, also, the definition of chatbots, why, by the way, in a minute, and I'll kind of throw some thoughts out on that, is well beyond the frontline peer one support of customer service. Meaning that, I see that there's a bigger opportunity for chatbots, or worldly speaking, conversational AI to support employees in the bank or credit union or even fintech as much as supporting the customers.

And the reason I kind of hesitate on using the term chatbots and conversational AI interchangeably is that, there is a big difference. There are a lot of deployments of chatbots that are really nothing more than rules-based engines. You ask a question, you get an answer, and that's really the extent of the conversation. The next question also produces an answer, but the system or the tool does not really know or recognize that there was a question that came before that, so there's no history.

Conversational AI is different. It's a more sophisticated approach from a technology perspective that is much more conversational in nature, which is why they call it conversational AI, which says, "Wait, we know that there was another question asked before that, and now we see this next question, so we start to see a trend and can anticipate where this conversation is going." So, there's a lot more than just sort of a front and tier one easy question support capability. In fact, I would argue that much of the technology that's being developed today is actually more sophisticated than what the human support people can provide.

And so, therefore, the opportunity to use the chatbots, the conversational AI driven chatbots, internally, to better inform the people on the front line, to better inform the people in the back offices about what's going on in the business and with customers and with support and with the need for advice and all of that, is actually a whole lot more sophisticated and opens up opportunities for productivity improvement.

So, part of the challenge we've got here is definitional. There's a wide range of, what's being called, chatbots that range from simple rules-based engines to more conversational AI machine learning type approaches, and a wide range of deployment that ranges from banks and credit unions who are really doing nothing more than just trying to skim off the top of some easy questions, and those institutions that are, in their terms, making chatbots a member of the team, which is a weird kind of concept to think of having a teammate that is a chatbot.

But I have to tell you, from companies like Boeing Employee Credit Union to Royal Credit Union, there's a whole bunch that really believe and buy in that they're making chatbots members of the team.

Alan Kaplinsky:

It's interesting, Ron, that you mentioned credit unions. And after we first talked on the phone a while ago, I watched a webinar in which you were involved, and there was somebody from a credit union on your panel. And I was astounded to find out that credit unions, as opposed to banks, seemed to be in the vanguard here, with the respect to chatbots. I mean, they're generally very, very small. Most credit unions are well below a billion in assets. I would not view them as being particularly technologically savvy, and yet it seems like they're the ones that are trying out chatbots. Why is that? I'd love your observation on that, if you have one.

Ron Shevlin:

Well, Alan, you kidding? I'm Mr. Fintech Snark Tank, I got a opinion on everything. I think there's, seriously though, two reasons why this is. First of all, the big secret in the industry is that, credit unions are actually more aggressive in their technology investment than most community banks are. And Cornerstone Advisors where I work, we've actually been quantifying that for 15 years. As a percentage of assets, credit unions actually spend more on technology than community banks do.

So, that's one, but it's very closely related to the second reason which is, because compared to a lot of community banks that often see themselves as really nothing more than commercial lenders, credit unions are more consumer-focused, have more of a customer support need requirement versus many community banks, again, who see themselves more as commercial lenders. And so, therefore, don't invest as much either in a technology or in consumer facing technology as other community banks, especially the credit unions. So, I think that's really why you saw that and see that difference, is a couple of those reasons.

Alan Kaplinsky:

Yeah. I mean, I can't remember the name of the credit union, it was on your program. But the person from that credit union said that their customers or members like it so much and they've given the chatbot a name. It may be called, Alice, I don't remember. But they'd rather speak to the chatbot than they would to a live person.

Ron Shevlin:

It's not unusual.

Alan Kaplinsky:

Yeah. But even the big banks don't seem to be, as far as I can tell, using chat box. I mean, I'm talking about the mega, mega banks.

Ron Shevlin:

Oh, no, no, no, no, no.

Alan Kaplinsky:

Am I wrong?

Ron Shevlin:

Number one, Bank of America's been out there for a number of years with their chat back called Erica. They boast that it gets millions, I don't know, probably even each month, at this point, of interactions.

Wells Fargo just introduced theirs. I'm not sure if Chase has one, but yeah, the large banks have been out there. US Bank has something. I don't know what they call it. But yeah, the large banks, Alan, have been very aggressive about implementing chatbots.

Alan Kaplinsky:

Interesting. Yeah, I didn't realize that. Well, it shows you how infrequently I call customer service.

How does it differ from a voice response unit, which is the least... Up until the use of AI in chatbots, that's the thing I'm used to using. And as a customer, it's frustrating as hell. I mean, it's a terrible experience.

Ron Shevlin:

Well, very simply, the best way to think about the difference is that, the voice response unit is much more akin to a rules-based engine. You're asking a question, you're getting an answer, there is a path and a rule. Visualize it as a decision tree that the customer press one, two or three. If they press one, go to this question, ask, do they do one, two, or three until you get to some end of the tree. The chatbot, especially one that is powered or driven more by machine learning and conversational AI tools and technology is much more conversational and relying on what is, hopefully, a deep set of data, both in terms of input and output, which is important.

So, it's learning about what questions and answers, what answers satisfied the question, what answers produce these sets of questions. It's much more sophisticated than a voice response unit, which is really just simply a decision tree.

Alan Kaplinsky:

Yeah. In other words, it thinks before it gives the answer, whereas a VRU, it's just completely programmed, it doesn't do any thinking.

Reggie, I assume there've got to be some legal risks that are involved in fintechs and banks and other financial institutions using an AI chatbot. So, what kind of legal advice would you give on what's doable or not?

Reggie Young:

Yeah. I think a lot of the risks we generally talked about with AI apply accuracy, misrepresenting these AI tools, the chatbots potentially hallucinating their answers.

Alan Kaplinsky:

How in the world does that happen? I mean, I'm just curious, the hallucination part of it, there's got to be a technical reason how that can happen. Like, coming up with a fake legal case.

Reggie Young:

Yep. I think one interesting example that I find very intuitive is like, imagine a bank, updates its fee structure on its cards. This AI chatbot is going to be, potentially, built by scraping historical data on the web, related to this bank's fees. And so, if this fee has only been in place for a week, it might still cite the old fee, and so it can provide something that's factually inaccurate. I think that that's a pretty intuitive example, for me.

So, I think one of the most compelling AI legal or AI chatbot legal use cases to me is in handling basic kind of legal requests, complaints, those sorts of things. I think the two that come to mind are under many states privacy laws, you have a right to opt out of sharing data or a right to have your data deleted. And normally, that goes to a support team, and they have to identify, and it gets handled, sometimes escalated to legal to help deal with, is this actually a deletion request?

Another good example is Regulation E, regulate debit cards, prepaid cards, which requires certain handling of disputes and claims of error on... If somebody sees an unauthorized transaction, you have certain obligations for consumer cards to respond and handle those complaints.

Alan Kaplinsky:

And I guess, also, truth in lending or credit card dispute resolution, right?

Reggie Young:

Yep, exactly. Regulation Z with credit cards as well. And I think, in all of these, what you have, I actually think is a big, big potential benefit of these AI chatbots which is, as a lawyer working with a support team, you will put together templates. And these are pretty standard because you need to respond within X days. Your support team needs to acknowledge they got a dispute about a charge, for example. While we talk about how AI can get things wrong, humans can also get things wrong, regularly.

And so, I think with AI, you have a more standardized response, right? Your team is going to respond to privacy opt out requests in a more standard way. I know that it's a frequent headache for lawyers to deal with, the small things that humans can get wrong, that can potentially create a lot of regulatory exposure. The risk, of course, is that, the AI chatbot doesn't accurately flag, "Hey, this is a privacy deletion request." I think that gets to my earlier point of, you really want a human in the loop doing some kind of audit and oversight, to make sure that those requests are being picked up correctly.

Ron Shevlin:

Hey, Alan, if I can jump in for a second on Reggie's point, because I think it's really, super important. Look, you guys are the lawyers here, but Chopra from the CFPB has also expressed concern about bank's use of chatbots as ways of... I think in his perspective of kind of offloading human interactions to the chatbot and is concerned with that. And I think to your point, Reggie, it just kind of goes to show the short side of this. This is that, number one, the chatbot's responses could often be actually be proven to be more accurate and better than the humans. And number two is the, I'll use the term paper trail, although that's not obviously the right term. But the virtual paper trail, and that being able to show, demonstrate and prove that the chatbot's responses to a series of questions is a consistent and based on some form of input that drove that, versus humans' responses which are... Who knows what's driving some of that. Not everything in the call center is overly scripted.

I often like to chat a lot of banks who tell me how great their branch experiences by saying, "You got to be kidding me. Post pandemic, you've seen so many of your good, long tenured people leave." Too many of the people in the branches are folks with five, 10 minutes worth of banking experience. When my wife goes in, she knows more than they do, not just about finances and financial services, but about their own products.

So, the chatbot is actually a way of improving the consistency, the quality of answers and the paper trail of this. I think Chopra is way, way off base with his concern of chatbots.

Alan Kaplinsky:

Yeah. Well, I'll give you a simple example, I guess, of a case where it would've been nice if I were dealing with a chatbot. I called up a bank where I have a CD account and it came up for maturity, and I think it was paying me 0.01% for the last 12 months, and I figured they must be offering some kind of special rate, now with interest rates being much, much higher than the last time I renewed that cd. And I talked to some customer service rep and she was absolutely convincing that, no, the best rate I could get... It was no longer 0.01, it was 0.1% for a one-year CD. And I said, "You sure? You must be losing a lot of deposits to your competitors." Nobody is paying that today. They're paying for one year CDs closer to four or 5%.

So then, when I got off the phone, I went on their website, and sure enough, there was a special seven-month CD paying 4.5%. She didn't know about it. She's in customer service. The chatbot wouldn't have forgotten that, I don't think, right?

Anyway. So, when all is said and done, Reggie, what's really doable today, in terms of the chatbot world? If you're a fintech company, you're a bank, and you're interested in using chatbots because you think it would be a more efficient and productive way to relate to either your employees or to your customers.

Reggie Young:

Yeah. I think, generally, it's going to be less of a legal risk for simpler uses, these basic fill-in, basic support questions, potentially no legal questions about exercising privacy rights or I want to dispute a charge.

Like I was saying, I think a key thing is making sure that there's a human in the loop that kind of has some oversight, some kind of audit function for AI. And I think it also probably toggles a bit on product. Like, you probably don't want AI involved too much in credit products, given the fair lending hot water you can get yourself into, versus kind of more like, "Oh, what's my bank account balance?" or "What fees was I charged this past month by using my bank account?" That kind of basic question, I think, is more palatable.

Alan Kaplinsky:

To give my example, if I had gotten a chatbot instead of a human in the call I made to renew my CD, chatbot probably could handle all of that. Without any kind of human intervention.

Reggie Young:

Yep, agree with that.

Alan Kaplinsky:

Yeah. But the credit world is probably... We're not quite ready for that yet, I guess. Is that the point that you're making?

Reggie Young:

I think maybe it's a question of how much... Maybe it's that audit function I've been talking about. Like, if you want to do it in credit, you're probably going to have to have an outsized audit and oversight of your AI tools and their responses and outputs, compared to, say, a bank account type of product.

Alan Kaplinsky:

Yeah. What about complaints? The complaint function? I mean, anybody who operates a financial institution knows that you get a ton of complaints. Some of them are lodged online, some of them are lodged over the telephone, some of them, unfortunately, get lodged with the CFPB or the Federal Trade Commission or state attorney general. Is that another area that you would be very wary of using a chatbot?

Reggie Young:

I think it depends how you implement it. I think if there're going to be basic like, "Hey, your product sucks" type of comments, there's going to be, "Hey, I have this complaint about fees that were charged, even if they're correctly in the legal terms, whatever." I think that the risk is having AI that correctly identifies when a complaint is a complaint, and making sure it gets escalated as necessary. If somebody's threatening to sue, you want to make sure that gets escalated to human, probably. And maybe that ends up being the hack that customers do to get around these AI chatbots to get actual human services, just kind of make threats like that. Who knows?

I think there's definitely some use case for it, but I think you almost want a chatbot that is more sensitive to flagging complaints and escalations than maybe warranted, just to be safe.

Alan Kaplinsky:

Ron, what are the credit unions that you're familiar with? How are they using chatbots right now? Are they staying away from credit, focusing on deposit accounts? What about complaint handling? What's going on there?

Ron Shevlin:

They, and a good number of banks, large and small, I think, are kind of all over the map with this. From a customer service perspective, I think what's important, too, is, we have to go back to distinguishing between the rules-based engine chatbot and the AI conversational machine learning driven, because those who are simply doing front-end, rules-based engine stuff, I don't even want to call those chatbots anymore. And in fact, some of the vendors don't, they like to distinguish them by calling the AI driven ones intelligent agents, or there's a word in the middle I'm blanking out on.

But the important point is that, it's more than just sort of the response. There's the benefit. And one of the benefits and opportunities with the more sophisticated use of the chatbot is the backend data analysis that identifies the trends of what's going on. And so, number one, it's not enough to just say, "Hey, somebody says your product sucks. How do I respond to them?" Well, a sophisticated chatbot, intelligent digital assistant... That was the word, digital assistant. I was looking for... Response would say, "Number one, who's lodging this response?" If it's Ron Shevlin, my age, my history of interaction, it may pay to not respond digitally. But to say, "Let me call you. Let me respond."

My kids, the millennials, the Gen Zs, they don't want to talk to anybody, anyway. They're more happy to deal with the automated response. Number two, how many responses or inputs are we getting that says, "The product sucks"? Is it a trend? What triggered that? Is there something about it that says, "Wait, this is not a one-off person who has a problem, this is a deeper trend we have. This is a capability that, I think, most contact centers, customer support groups, really don't have and don't do, that they look at the back as a big opportunity for that.

So, there's a lot of opportunities to use the chatbot, digital assistant capabilities to do more deeper analysis, trend analysis, and also sentiment analysis. It always drives me nuts when I'm dealing with a chatbot and I'm telling them, "You're stupid. You don't know what you're talking about". There's no..." Hey, okay, this guy is really on the verge of having a nervous breakdown, let me get him a real person right off the bat". There's a sentiment analysis perspective that, generally, we, as humans, are pretty good at. We can usually tell by talking to somebody, is this person really off about, off the rails? And should I escalate this up to a manager or whatever it might be?

Again, the big differences in terms of the capabilities, my big fear and what I keep preaching is that, you've got to look at this more as a tool that supports a wide range of your organization in terms of your ability to gather and use data, and not just reduce the number of hours that your customer support people are providing support to. Big difference in terms of their use. And I think we're early in this evolution. The technology's been around for a while, but I think in terms of the percentage of firms that have adopted this, it's certainly increasing exponentially these days, but I still think we're early in its impact on productivity and customer satisfaction.

Alan Kaplinsky:

So, I'm just curious, Ron, can a chatbot or digital assistant detect when the person they're talking to is angry? If somebody starts shouting or cursing, can it tell?

Ron Shevlin:

Some, can. It's not necessarily a universal capability of a lot of chatbots, because it is dependent on the quality and the capabilities that underlie it in terms of the system that they're using. But yes, some absolutely can, it's sentiment analysis, they look at the words that are used. If it's typed, they can sense in terms of the tonality and volume, if it's more verbal driven, oral driven.

So, yeah, there are those capabilities in some of the systems that are out there, and certainly something that can be taught if it was a self-developed chatbot system.

Alan Kaplinsky:

Yeah. I mean, we've been talking, mostly, today, about voice, the use of voice chatbots, but I take it, it can also be used online. A lot of banks and other companies have chat boxes where they encourage you, if you've got questions, type it in. And I assume, a chatbot can respond to that.

Ron Shevlin:

Yeah, absolutely. I will tell you that, all along, in this discussion, in my head, I've been thinking more the written interaction than the verbal interaction. So, interesting that we were coming at this from different perspectives, because I think a lot of those interactions now are happening, especially on mobile devices, which is funny, because they're great for voice, but it's funny how nobody thinks about using the phone as a phone anymore.

Alan Kaplinsky:

Yeah.

So, let's talk in just the remaining couple of minutes, and we're going to wrap up our show. What do we see down the road? What do we see in five years from now, 10 years from now? Can you use your crystal ball and figure out where chatbots are going?

Ron Shevlin:

So, I will put the stake in the ground as I have written in on my blog, that... I want to use the term conversational AI, not chatbot. That conversational AI is becoming, and will be a foundational technology in financial institutions and fintechs. And by foundational, I mean everybody has to have one or more. It has to be there. It's not an optional technology. And not just for customer service support, but for a wide range of functions in the organization, as much to support internal employees as it does external customers and members. The opportunities to deploy conversational AI and various functions, again, it ranges just as much for conversational AI across functions as it did the examples that we gave earlier around generative AI, where the real differences are more in terms of the output than the actual input.

And someone's probably going to get on my case for that one and argue with it, but I think that the five to 10 years out, we're looking at conversational AI as a technology that every financial institution has, with multiple types of deployments, meaning that there might actually be more than one chatbot in the organization. And Erica is great for Bank of America to support customers, but I don't think it supports the internal employees as much, which means there's opportunities for Bank of America to develop a whole different type of chatbot to support employees. And given the range of work that employees do, one might be for the retail group, one for commercial, one for investment banking. I think it becomes a foundational technology with multiple types of deployments within the institution.

Reggie Young:

I agree with that. I really like Ron's point about, we're probably going to see some specialized uses when we were talking about whether Chase is an AI. I was briefly googling that on the side and saw some headlines around how they want to roll out an AI investment advisor that they're been working on the past few months, apparently. I think you'll see uses like that. But I can

also imagine the compliance and legal team is going to want a chatbot that's specialized in handling complaints and privacy requests and those sorts of things. So, I think that that specialization will definitely, increasingly happen.

And also, I agree with Ron's view that this is going to be kind of standard. These are going to be expected. I think, right now, it's like, the right now AI and financial services is kind of dealing with the problem of that self-driving cars we're dealing with, where you look at the data and they're safer than human drivers, but that's not the perception that a lot of folks have, and that it's going to take time to change that. I think there's similar, to Ron's earlier point, but CFPB not totally getting it. I think there's a lot of good arguments that AI is going to be more accurate and better to deal with, and more compliant than human beings, but there's a lot of education and enlightenment that needs to happen between now and then. But we'll progressively cross that chasm and just see more and more ubiquity of these chatbots.

Alan Kaplinsky:

Well, the thing I worry about the most is what you just mentioned, and that is that, the regulators, some of them have a knee-jerk negative reaction to new technological developments. And I think Rohit Chopra is a perfect example of that. Probably, Lina Khan, the chairwoman of the FTC. The banking regulators, I don't think, are on top of it. Unfortunately, they don't have the best experts on their staff. They're not paying sufficient money in government. Why would somebody who's an expert in this area want to go to work for the CFPB if they can get a private job and they can make 10 times as much money?

So, they're lagging way, way behind, and all they think about are how consumers can get hurt and how chatbots might discriminate against consumers. And it just worries me that that could stifle the growth of technology.

Either of you have a reaction to that before we conclude things?

Ron Shevlin:

Yeah, I do. Well, I have to tell you that I tend to not agree with a lot of the things that the CFP does. I have to say that, almost to a person I've ever dealt with at the CFPB, they've been super sharp, what really experienced, and know their stuff. So, I don't diminish their ability, capability, credibility or any of that, I just kind of disagree with them more philosophically.

But I do share your concern that, in general, the regulatory environment cannot keep up with technology change. I don't think that that's what they're set up to do, and I think that's the biggest problem. And I'll just tell you, just anecdotally, a couple of years ago, I was on a panel at the Money20/20 conference with a former CFPB person, and the question posed by the moderator was, can the CFPB and other regulatory bodies keep up with technology change? I argued, "No." The CFPB person, ex-CFPB person argued, "Yes." They put the vote out to the folks who were attending. I won, and that was the last and only time I ever beat a lawyer in a debate.

So, I'm going to leave that one at that one. Leave it at that, Alan.

Alan Kaplinsky:

Okay. Do you have anything more you want to add, Reg, on this last question of the regulators being able to support the development of the technology, rather than to do everything they can that put an obstacle on the way?

Reggie Young:

Yeah, I analogize it a little bit to what's happened with fair lending where, if you want to do something novel and unique, you need a bunch of data to back up, that it's not going to be discriminatory. You probably need to go get a no action letter in a way that's just like, there's no interesting different underwriting methods that are being done besides FICO or cash flow, as a result. Like, it's kind of stifling, and I worry that that would happen with the use of AI. It's kind of chicken and egg problem, where to be able to go and convince a regulator that you're not a bad actor doing or that you're AI use isn't irresponsible, you're going to need data. But in order to get that data, you need to have the AI bot fielding complaints and questions to be able to get results. And so, it's a little bit of a chicken and egg problem, I think.

Alan Kaplinsky:

Yeah. I mean, the other thing I worry about, I read the other day, Elizabeth Warren, and... I'm trying to think of the Republican that is involved with her. They want to create a new agency devoted to AI, just to AI, and they're bragging about the fact that they've got bipartisan support for this. That really scares me.

Anyway, we have to wrap up our show. I want to thank both of you, Ron and Reggie, for joining us today. This was extremely informative to me. I'm sure our listeners feel the same way. And I'll definitely want to invite you back on other occasions when we, once again, are going to be focused on generative AI.

So, thank you, Ron. Thank you, Reggie.

Ron Shevlin:

Thanks, Alan.

Reggie Young:

Yeah. Thanks, Alan. Like I said, a big fan of Ballard and everything you all are doing, the Consumer Finance Monitor, so honored to be a guest on the show.

Alan Kaplinsky:

Thank you.

So, to make sure you don't miss any of our future episodes, subscribe to our show on your favorite podcast platform, be it Apple Podcast, Google, Spotify, or wherever you listen. And don't forget to check out our blog, also, going by the name of Consumer Finance Monitor. We have daily insights there on the consumer finance industry. We launched the blog on July 21, 2011, a day that we live in infamy, when the CFPB got stood up and became operational. We just had our 12th anniversary, like the CFPB.

And if you have any questions or suggestions for the show, let me remind you again, you can email me, kaplinsky@ballardspahr.com, or you can direct it generally to our email address, podcast@ballardspahr.com. Stay tuned each Thursday for a new episode of our show. Thank you for listening today and enjoy the rest of your day.