

---

***THE CONSUMER FINANCE PODCAST: HOW VOICE AI AND MACHINE LEARNING REDEFINE CALL CENTER CONTROLS***

**HOST: CHRIS WILLIS**

**GUESTS: STEFANIE JACKMAN AND SCOTT HAMILTON**

**DATE AIRED: JULY 20, 2023**

**Chris Willis:**

Welcome to *The Consumer Finance Podcast*. I'm Chris Willis, the co-leader of Troutman Pepper's Consumer Financial Services Regulatory Practice, and I'm really glad you've joined us for today's episode, which is going to be about some interesting new machine learning voice analytics technology for the financial services industry. But before we jump into that topic, let me remind you to visit and subscribe to our blogs. We have [consumerfinancialserviceslawmonitor.com](http://consumerfinancialserviceslawmonitor.com) and [troutmanpepperfinancialservices.com](http://troutmanpepperfinancialservices.com). Both of our blogs devoted to the financial services industry.

Don't forget about our other podcasts. We have lots of them. We have the [FCRA Focus](#), all about credit reporting. We have [The Crypto Exchange](#), which is about everything crypto, and we have [Unauthorized Access](#), which is our privacy and data security podcast. All of those are available on all the popular podcast platforms. Speaking of those platforms, if you like this podcast, let us know. Leave us a review on the podcast platform of your choice and let us know how we're doing.

As I said, today we're going to be talking about a very interesting new technological development relating to call monitoring and voice analytics. I'm joined by two guests today. First, I have my partner Stefanie Jackman, who's a member of our Consumer Financial Services regulatory team here with me in Atlanta. But then we also have a very special guest with us today, Scott Hamilton with Prodigal. Scott has about 30 years of experience with top 10 banks here in the United States, and about 10 years of experience in a first line risk role in all kinds of call center environments, and he's going to be telling us a lot about this new technology today. Stefanie, Scott, welcome to the podcast.

**Stefanie Jackman:**

Thanks for having us, Chris.

**Scott Hamilton:**

Thank you.

**Chris Willis:**

Stefanie, let me start with you. You and I have grown up together in the financial services industry, and when we think about call monitoring, my mind first leaps to the old school way of doing things, which is when you have a call center in order to check for compliance with company policies, applicable legal requirements or whatever, you would pull randomly five or 10 calls for each customer service rep or collections rep or whoever the case might be, and have a human being listen to them and score them against a scorecard every month. That's the old way of doing things.

But of course, the old way has a significant limitation in that there's only so many calls that a human being can listen to. You end up listening to like five calls out of hundreds that a person may be on for a month. You question whether you're really getting a fair sample and catching everything you're supposed to catch. I noticed a few years ago there was a big uptick in the industry of adopting voice analytic or speech analytic technologies to help in call monitoring.

Stefanie, can you just tell the audience a little bit about those existing voice analytics technologies and what financial services companies are using them for?

**Stefanie Jackman:**

Yeah, I'd be happy to. As you said, Chris, we really could not historically speaking, and despite our best efforts, get much insight into what was happening on calls. Many of our clients have thousands and thousands of calls a week, very high volume, and sometimes you just have to get lucky. You and I had talked about at least having, instead of just random sampling, trying to have what we call targeted random sampling, like different accounts that were at certain statuses or had certain things happen.

But that still doesn't capture everything that you would need to capture in an effective call monitoring to capture potential violations, non-compliance with policy, training opportunities, and the whole panoply of things that can really add from a value proposition to your company. In addition to the opportunities that it can present to the business in just giving better client service, ensuring a more robust compliance management system, the CFPB really would like to see, and has for a number of years now, enhanced QA and call monitoring. When some of the existing tools came out, the CFPB was very excited about them. I remember you and I, Chris, we'd walk around and tell people you should invest in it.

But then what we started to hear maybe seven years ago or so and continuing to this day, well, do we have to have this? If we don't have this, will the CFPB or a regulator fault us? Well, why wouldn't you want to have it? Well, the answer that we heard is that it can be expensive, number one. Expense is always a consideration. But number two, there's lots of false positives. We're finding it challenging to work with. It's not giving us the insights we want. It's not allowing us to focus and make our actual human QA more efficient. What I really did hear were so many false positives and it just wasn't always a useful tool.

On the other end of the spectrum, you and I saw some clients who took these existing tools and implemented in-system enhancements and things on their own end to really customize them into something that served the business's needs and used them quite effectively, but that took a lot of effort and programming.

It's this uneasy world where we currently are in where some companies use these types of technologies, some companies don't. I do think, and Chris, I invite you to let me know if you disagree, I do see the CFPB, at least historically, it has asked, why don't you use call analytics? Or now we have text analytics and all sorts of things, and I'm seeing it raises a little bit of a flag. I think there is an expectation from regulators that a compliance management system should have these types of tools involved in it. That doesn't mean that they'll always fault you for not using them, but it gets attention and it's almost expected in this environment, even if there are shortcomings.

**Chris Willis:**

Got it. I think I agree with you that the speech analytics technologies that have been around for the past few years, although not perfect, certainly improved call monitoring beyond what we had with the old way of listening to five calls a month per person.

But Scott, let me turn to you now. Would you mind talking to the audience about basically how do the existing technologies work and what challenges are there with using those technologies?

**Scott Hamilton:**

Yeah, thanks Chris. There's a handful of companies that play in this space. They generally have been around for a decade. They generally are all working the same way. In the case of a call center, they take a voice conversation and transcribe it to text. They then empower the client to tell it what to look for. It introduces a tremendous value in that you can pretty easily find the needle in the haystack, where if somebody says bankruptcy, you can find all the calls that had the word bankruptcy in it, which is great. Before you had to get lucky and have it fall into your five a month example earlier, Chris. In that regard, they're very effective at finding a nice sampling of all your different types of call types. That's the good news.

The challenge back to how they work is the voice to text transcription tools generally are across the board with many different industries. They'll use the same transcription tool that they will for American Airlines as they will for ABC Bank. Those transcription tools are not super accurate. They're general tools that solve a general problem, and if you mumble or have an accent, the confidence levels won't be that high, so they'll just leave a blank. That's issue or challenge number one.

Challenge number two is the tools that they provide you, it requires for the most part the client to tell it what to look for. Most clients are really bright, but they can't come up with every way that customers can articulate a particular issue. So, you end up in this never-ending brainstorming cycle of different words or short phrases that identify the situation that you're trying to uncover.

To Stefanie's point, because you ask for, I'll use the bankruptcy reference, there are a lot of conversations where you might say bankruptcy, but you didn't mean it, like my neighbor filed bankruptcy. I don't want to file bankruptcy. But they'll catch those and they'll turn into false positives. I've heard ranges anywhere from 30% to 40% to 80% false positives.

Then finally, on the other side of that same coin, the false negatives. They might say, my brother passed the bar and he's going to be representing me today. I never said the magic word, but I meant it. While they're really good at finding calls that include certain elements, the generic transcriptions are a little bit flawed, and then they only are as good as what you tell them to go look for. They're a fair amount of effort, and the false positives and negatives oftentimes result in a net 30% or 40% accuracy rate, if you include negatives and positives.

It can be really challenging to then go to the regulator and say, we have an effective control environment or your product leaders or your efficiency leaders to really give a degree of confidence that you've got a very clean pulse on what's going on in your call center.

**Chris Willis:**

Understood. Thanks, Scott. Time marches on and technology improves. Now that we've got some years of experience with the transcription-based voice analytics technologies that you were just talking about, why don't you tell the audience about some opportunities that there may

be to leverage machine learning models to enhance the performance of these kinds of speech analytics?

**Scott Hamilton:**

That's actually exactly why I joined Prodigal. I won't talk much about Prodigal, but how it's built was really the issue. With that, those headwinds in Legacy Voice Analytics platforms, the only way that I've heard the industry uncover how to go from 30% to 40% accuracy to 80, 90, is to build smarter models or to build models that actually learn. The banks generally have had a fairly good set of experiences up in the front end of the business in originations or in fraud, and now they're starting to redirect that talent toward this problem.

The way these models will work in the voice environment is just like fraud or originations, you have models or algorithms that interrogate the dataset and identify patterns. They don't look for words, they look for patterns, and they take a tremendous amount of volume and a tremendous amount, and this is the hard part, a tremendous amount of training data. Not only do you need to build the algorithms, build the hardware, put lots of your calls in this case, lots of your calls into the algorithm, you need to review the output of those models. But when you do, you can tell them when they are right and when they are wrong, and over time they will learn, machine learning models will learn how to uncover the right patterns with a very high degree of accuracy.

In the case of Prodigal, we've built the models using natural language processing and machine learning and cloud technology for scale. We've ingested a quarter of a billion phone calls. There's maybe 10 banks in the US that have that sort of volume and the tech. But the most difficult part is to have the algorithms run, sorry, to get a little technical. Have it produce output, but have agents or QA analysts review the output and tell the model when it's wrong so that it can get smarter and smarter. With that training, you're able to yield accuracy rates that often exceed a human, exceed the analyst, exceed the judgment of the frontline agents, which opens up a tremendous amount of opportunity and efficiency. Hopefully that was clear on the technical design, Chris.

**Chris Willis:**

As listeners of this podcast know, we really enjoy getting into the technical stuff, particularly as it relates to machine learning. It's one of our favorite subjects, so don't worry about doing that.

Scott, let me do ask you a follow-up about that. You've talked about how you train the machine learning model by basically having people score the calls and then it learns from those scoring of calls. How does the model then work in production?

**Scott Hamilton:**

There are a couple of themes that are starting to emerge. Historically, legacy voice analytics as referenced earlier, typically has taken in batch calls. Every night, the various tools will go and we'll run the audio to text conversion and then run queries against that text, and that would happen at night. What is happening over the past four or five years is there's various technologies that are coming out that can not only do that with a much higher degree of accuracy, with much lower degree of effort, but it's also like many of these technologies shifting upstream to real time.

So, these same models can be run during the call as it's happening and can understand how the conversation is turning left, turning right, and can apply these learnings that the models have achieved to understand the context, can understand the intent of the evolving conversation, and

as an example, can recommend to the agent in the call what they should say next. It's about as close to a preventative control as I could imagine. You can't make an agent or stop an agent from saying something, but you can hold their hand and basically be that team leader listening into the same call and giving you some advice.

So, there are definitely real time solutions. There are others that, as an example, would listen to the whole call and would summarize it into a note, not a transcription, but a summarization of the conversation. So, from an efficiency standpoint, the agent doesn't have to take notes. From an accuracy standpoint, you get tremendous gains there. There's a variety of different use cases where this technology is showing up across the call center and non-call center, email, text, etc.

**Chris Willis:**

Well, first of all, I hadn't even thought about the real-time part that you were talking about, but one of the common practices in call centers is that you'll have a supervisor who's on the floor with the agents and he or she will have the ability to tap into a call at any time. Are you saying, Scott, that this tool would maybe flag them and say, hey, here's a call you need to tap into?

**Scott Hamilton:**

Both, yes is the answer. Yes, in both. On the screen of the agent having the primary conversation, it will prompt them with what to say next, and when they do address that topic or use that degree of empathy, it will check mark and give them some positive reinforcement. Then there's other user interfaces that, for example, as a team leader, if you are figuratively or literally on a hallway with 20 of your agents, you can look at a dashboard and realize that 17 of them are green and are going fine. Empathy's going well, procedures are being followed. Or there might be two that are going red and you can literally walk down the hallway or tap into the phone line and if needed, use your expertise in order to help get that call back on the rails. So yes, there's a variety of use cases that that can show up in.

**Chris Willis:**

Let me ask you about one other potential use case. One thing that we sometimes see clients struggle with is identifying and logging complaints that are received in a first line customer service phone line, because you either have to rely on the customer service agent to figure out it's a complaint and log it into the complaint monitoring system, or you may lose it. Is there any use case for this technology in basically pulling out complaints and helping to log them from a customer service phone line, for example?

**Scott Hamilton:**

It's a fantastic question. One of my favorite examples actually of something that you almost can't do without these types of models. As all of your listeners know, many banks are challenged with that exact issue, and it really starts with identifying the complaint in the moment and capturing it in order to be aggregated in the backend and validated and remediated, etc. But it's that in the moment complaint capture that is the most difficult.

To Stefanie's earlier point, many banks and firms are using voice analytics to expose certain words that equate to complaints, but oftentimes are finding that they're missing more than half of them because the word didn't exist. And these models are being redirected quite aggressively to exactly what you just mentioned, Chris. So, it's a very subjective question to answer.

However, there really isn't a choice in today's world. Banks and firms are leaning on their frontline agents to do as good a job as they can, and then they're QA'ing it like crazy on the backend.

The models are leveraging the learning that they've made because they understand lending or banking call types and conversations that they're then able to, with a lot more accuracy, identify very nuanced questions, complaints being a great one. You can load in, for example, identify calls with frustration or disagreement. Those aren't words that they're searching for. Those are intents that the models are searching for. So, yes is the answer. Identifying those types of intents do a much better job, much more efficiently, of identifying those complaints and bringing them into the back office for validation with a far lower degree of false positives. But most difficult is the false negatives. So the answer is yeah, it's a great example.

**Chris Willis:**

Stefanie, you and I have been sitting here listening to this, and you and I of course have to deal with regulators all the time. So, let me ask you to comment on, based on what we've heard, how do you think regulators might react to an implementation of this technology?

**Stefanie Jackman:**

I think that if it works well, and you can demonstrate that it works well, they would generally have a positive reaction, but they'd have questions first. Just generally, we're seeing this with the use of AI, which is a rapidly evolving area in machine learning, lots of attention to that. How are you teaching the machine? How do you unteach a machine? If it turns out that you've put in data or inputs that are resulting in the machine coming out with the wrong conclusion, you can't erase things from a person's head. How do you undo that in the AI context with some level of confidence?

How do you make sure the machine is not treating different populations differently based on protected characteristics? I think you'd need to be able to answer some of those. I don't know, Scott, if that's something you've thought about. It maybe way ahead of us, but I would expect questions. But then if you can show, here's how we control for that and here's how it has enhanced our abilities to monitor, be efficient, take corrective action, and better address consumer needs. I think they'd view it positively because they generally think that you should have some sort of voice analytics technology, whether you as an organization agree or not. That's what I was saying earlier on, that the failure to have one for whatever reason, always gets attention.

**Scott Hamilton:**

You bring up an excellent point that one, I have not seen or heard of any player in this space that runs the models, produces an output that the client acts on. There's always an agent or an associate interface. This is really producing situations or opinions to the agent that they ultimately have the up or down vote to release into subsequent activity.

I know from a machine learning or an AI space in the credit space or in the fraud space, what you bring up is hugely important. That said, I do think it's an excellent point that there is a validation set. There is a test and control activity to be able to prove and continue to prove that the models are accurate and not being over-trained a little bit left or a little bit right.

**Chris Willis:**

We've probably exhausted as much as we can do in a podcast episode on this, but there's probably a lot more to learn too. Scott, for those of our listeners who might be interested in learning more about this technology, where should they start?

**Scott Hamilton:**

There's various Google searches you can come up with or ChatGPT searches you can come up with, but I'm very passionate about this space. I joined Prodigal because they have the product that I wish I had for so many years. I'm more than happy to, without my Prodigal hat, connect with people offline and hear what they're wrestling with, share what some of their peers might be doing, what lessons learned they've had. So, I'll definitely make that offer to your audience and would be happy to follow up offline. Scott Hamilton, [scott.hamilton@prodigaltech.com](mailto:scott.hamilton@prodigaltech.com), or you can find me on LinkedIn as well.

**Chris Willis:**

Thanks a lot, Scott. Thanks for being on the show today. And, Stefanie, thanks for lending your expertise to our episode as well. Of course, thanks to our audience for tuning in to today's episode. Don't forget to visit and subscribe to our blogs [consumerfinancialserviceslawmonitor.com](http://consumerfinancialserviceslawmonitor.com) and [troutmanpepperfinancialservices.com](http://troutmanpepperfinancialservices.com).

While you're at it, why don't you go ahead and visit us at [troutman.com](http://troutman.com) and add yourself to our Consumer Financial Services email list. That way you'll get copies of the alerts that we send out as well as invitations to our industry only webinars. Of course, stay tuned for a great new episode of this podcast every Thursday afternoon. Thank you all for listening.

Copyright, Troutman Pepper Hamilton Sanders 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 Pepper 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. If you have any questions, please contact us at [troutman.com](http://troutman.com).