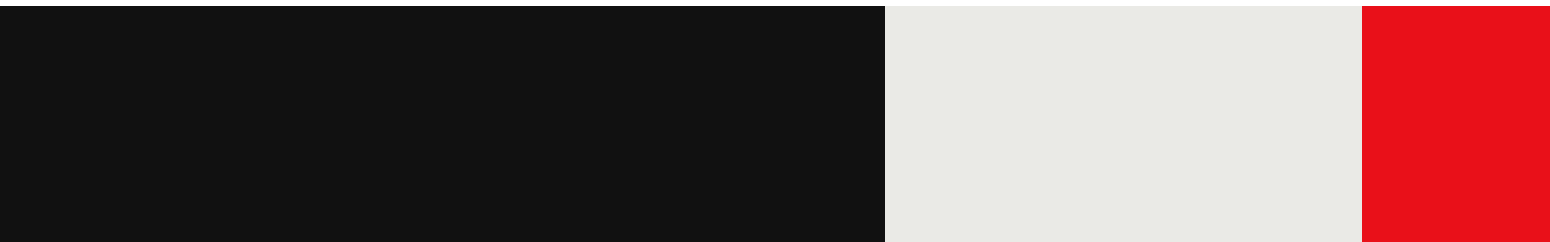



The Banking & Payments Show: AI risks to banks

Audio



In today's episode of The Banking & Payments Show podcast, we will be discussing the potential risks that AI poses to financial institutions. In the 'Headlines' segment, we will examine an article from BBC.com titled "Could AI Trading Bots Transform the World of Investing," which discusses risk-related issues such as AI bots making financial decisions autonomously. In the 'Rankings' segment, we will rank the 5 AI risk categories that financial institutions must address in terms of importance. Join the conversation as host Rob Rubin chats with analysts Jacob Bourne and Grace Broadbent.

Subscribe to the “Behind the Numbers” podcast on [Apple Podcasts](#), [Spotify](#), [Pandora](#), [Stitcher](#), [YouTube](#), Podbean or wherever you listen to podcasts. [Follow us on Instagram](#)



Episode Transcript:

Rob Rubin (00:00):

This episode is brought to you by TikTok for Business. Brat Summer is over, but TikTok for Business is just getting started. Whether your marketing goals are brand performance or full funnel related, leverage TikTok for Business' insights to drive business growth. Learn more at TikTok.com/business.

(00:24):

Hello everybody, and welcome to the Banking and Payment Show, a Behind the Numbers Podcast from eMarketer Made Possible by TikTok. Today is October 8th, 2024. I'm Rob Rubin, Head of Business Development at eMarketer and your host. Today I'm joined by two great analysts at eMarketer, Jacob Bourne who follows developments in AI, and Grace Broadbent who follows the payment industry. Hey guys, how are you doing?

Jacob Bourne **(00:50):**

Pretty good, Rob. Thanks for having me today.

Rob Rubin **(00:52):**

Thanks, Jacob. Hi, Grace.

Grace Broadbent **(00:54):**

Hello. I'm excited to be here.

Rob Rubin **(00:56):**

I'm excited to have you guys here. Let me ask you guys an icebreaker question. How much of your day is spent interacting with an LLM like OpenAI, ChatGPT or Claude?

Jacob Bourne **(01:11):**

So really good question. I haven't timed it yet, but it's probably become more over time just because of the nature of my work as a tech analyst, have to be using the technologies that we research. So it kind of goes with the territory.

Rob Rubin **(01:28):**

So you're doing it for work, but do you do it for any personal stuff?

Jacob Bourne **(01:32):**

Yeah, occasionally. It comes up with good answers to casual questions about things that just would be difficult to search and find using traditional search methods.

Rob Rubin (01:43):

I find that I'm looking at Gemini a lot now, when I'm Googling, and then I see that and I read that response.

Jacob Bourne (01:49):

Yes.

Rob Rubin (01:50):

Grace, what about you?

Grace Broadbent (01:51):

I agree. I feel like I've been looking at Gemini a lot when I do just any kind of Google search. I don't use it too often though. I'm probably more limited than Jacob. I've been playing around with it some, but I probably only use it once, twice a week, and mainly for work purposes. Nothing personal yet.

Rob Rubin (02:08):

All right. Just trying to take everyone's temperature. Today's subject is the risks of AI to financial institutions. Jacob, the last time you were on this podcast, we were talking about the opportunities financial institutions have for AI.

Jacob Bourne (02:23):

Yeah.

Rob Rubin (02:23):

So before we talk about all the different categories of risk, let's start with the headlines to get us into the topic.

(02:33):

In the headlines, I pick an article that's related, and there's a link in the show notes for the article itself. For today's headline, I chose an article on BBC.com, Could AI Trading Bots

Transform the World of Investing? And the reason I chose it is it gets to the heart of a lot of these risk related issues, and at the heart of it is, is that AI bots are ultimately making decisions autonomously. So they are listening to what you say and making a decision by themselves. So Grace, what could go wrong?

Grace Broadbent (03:08):

Oh, goodness. So much. So much could go wrong, Rob. I mean, AI is subject to hallucinations, biases, inaccuracies. You name it, there's all sorts of problems. And this is scary for investments because a lot of times people's money, future savings, are the things at stake and that can go so wrong so badly.

Rob Rubin (03:32):

Right. And all of the due diligence that's necessary behind the scenes to get the right data and if they're actually giving you the wrong data, how do you even know where to look?

Jacob Bourne (03:42):

Yeah. I mean, I think a central issue here is just accountability. I mean, if a person breaks a banking law or someone working in the industry makes an error, well, they either face legal consequences or lose their job. But if AI is a lawbreaker or making the error, well then who's responsible for that? And I think it could be is it the developer of the AI model itself? Is it the company deploying the AI model? Is it the person using the interface? And I don't think that there's great answers to any of those questions, at least not yet.

Rob Rubin (04:15):

It seems like it should be the company that is providing the incorrect information. They're responsible.

Jacob Bourne (04:21):

I mean, it could be. I think it depends on the particular instance and what happened. But I think that ultimately the stakes are high here because without accountability and a good system for established accountability, then it's really hard to have public trust in institutions that are using generative AI.

Rob Rubin (04:42):

Because talking about investing here, one of the challenges that comes up a lot with AI is the bias. So they're making investment decisions and there's an inherent bias to the decisions that they might choose to make. How does that get overcome?

Jacob Bourne (04:59):

Yeah.

Rob Rubin (04:59):

Does it?

Jacob Bourne (05:00):

I mean a lot of it goes back to having quality training data, that you're not training AI models using data that's biased. For example, not using data that's really representative of the full demographic scope of your consumer base, for example. But I think it's a deeper technical issue than that too. And I think it's something like with hallucinations, tech companies are working on it, but can you get to 100% resolution of the problem? And if so, when does that happen?

Rob Rubin (05:35):

But can they be smarter than their developers? How does AI overcome its developers weaknesses?

Jacob Bourne (05:41):

Well, I think it can and it can't. I mean in the sense that it can is a sense that generative AI is a type of AI that can learn, and it has these so-called emergent capabilities that allows it to act outside of its programming, at least for the most advanced models. And it's also trained on data that's so vast, that's beyond the awareness, I mean, the people, the developers making these models aren't fully aware of every data point that's going into training the models.

Rob Rubin (05:42):

Right.

Jacob Bourne (06:14):

So in that sense, yeah, the models are, they're capable in certain respects beyond the people who train them. And I think we have to remember here when we're talking developers, we're talking about a team of people.

Rob Rubin (06:24):

Right.

Jacob Bourne (06:25):

And so if one person on the team has a limitation, probably not a big deal, but if the entire team has a systemic bias or has a blind spot in terms of how they're building the model, well then that could be a big issue and could ultimately, I would say in that sense then no, the model can't probably overcome the weaknesses there.

Rob Rubin (06:47):

So it's going to really be a situation where it's going to be hard for us to know where the weaknesses are until they're discovered.

Jacob Bourne (06:56):

Absolutely.

Grace Broadbent (06:57):

Yeah, and I think that's a big issue is when people are looking through how AI makes decisions is we don't even know exactly how they are. Like the developers working on it can't tell you exactly how the AI came to this decision so how do you figure out what the bias is if you don't know how they made the decision?

Jacob Bourne (07:13):

Yeah, that's a really crucial point, Grace, is that we can't see the chain of thought between the input and the output. And when you have someone's credit application being decided by an AI bot, well, you kind of really want to know how the decision is being made, which you can't really at this point.

Rob Rubin (07:32):

But to give somebody information to execute a trade that's wrong or to make a recommendation on a series of stocks that there's a bias to why the AI is making those recommendations, all that could actually just stop the industry in its tracks. No?

Jacob Bourne (07:53):

I mean it's here though, and it seems like we're using it, so the industry is using it.

Rob Rubin (07:59):

But are they using it to make, they are using it to make investment decisions right now, but are they actually executing?

Grace Broadbent (08:06):

I think something interesting to think about, pivoting that a little bit, is we said can AI overcome developers issues? But can AI even predict the future, predict investment decisions better than a human, a trader can? AI doesn't have any knowledge that humans don't have, so can they be smarter than the human? That's really the question at stake in terms of investment.

Rob Rubin (08:31):

They didn't predict the pandemic, right?

Jacob Bourne (08:34):

Yeah.

Rob Rubin (08:34):

Like those sort of things which create sea change or just a tremendous change, it wasn't able to predict that.

Jacob Bourne (08:42):

I mean, I think we also have to make a distinction between predictive AI and generative AI, and I think we're increasingly seeing them being used in tandem with each other.

Rob Rubin (08:42):

Good point.

Jacob Bourne (08:51):

The models are getting more advanced and I think the predictive capabilities will increase. But I think it's also important to note that, I mean, one human being also can't think with the same computational speed, with the same amount of data, like instant recall to the same kind of data that AI model can with the same speed. That's not to say that AI is not going to make a ton of mistakes. It does. But in terms of healthcare, like diagnosing conditions, AI has proven to be very accurate in terms of predictions. So I think the potential is there. But I think also, just to be fair, people also sometimes can't give a logical explanation for their decisions as well. So the human brain can kind of be a black box in a way as well. We make a lot of our decisions based on intuition.

Rob Rubin (09:41):

Right.

Jacob Bourne (09:41):

So there's that aspect as well.

Rob Rubin (09:43):

This is an excellent time to transition because I wanted to use the headlines to get us sort of warmed up into the topic and I think what we've learned is that there's a lot to cover here in terms of all the risks that a bank could make. So in our final segment, we're going to do something we haven't done in a while, which is something we call the rankings.

(10:07):

Prior to recording this show, I shared a list of five AI risk categories that financial institutions must address, and we each ranked them in terms of importance. And I'm going to read the list and then we can go around and see what we ranked as one, two, et cetera, and then we'll discuss. So let me just read the list and then, Grace, I'll get to yours first.

(10:30):

The first on the list is cybersecurity. So things like an AI-powered cyber attack. Number two is frauds, so deepfakes, automated money laundering, those kinds of things. The third is regulatory compliance risks bias, which we've talked about, fairness, opaqueness, sort of covering all that. Fourth are sort of operational risks like systemic errors. And the fifth area is

data privacy or security, so things like breaches, handling PII, and that sort of thing. So, Grace, what was your number one?

Grace Broadbent (11:02):

I said fraud.

Rob Rubin (11:03):

All right. Jacob, what did you do number one?

Jacob Bourne (11:06):

I did cybersecurity as number one.

Rob Rubin (11:07):

Ah, okay. I did cybersecurity as number one as well.

Grace Broadbent (11:11):

I'm the outlier.

Rob Rubin (11:12):

Yeah. Grace, what about number two?

Grace Broadbent (11:14):

I put regulatory compliance.

Rob Rubin (11:16):

Ah.

Jacob Bourne (11:17):

I put fraud for my number two.

Rob Rubin (11:20):

Me too. Great. Now what did you put number three, Grace?

Grace Broadbent (11:24):

I have data privacy and security.

Jacob Bourne (11:26):

And I have the same for my number three.

Rob Rubin (11:28):

Okay, so we're all the same for three.

Jacob Bourne (11:31):

Okay.

Rob Rubin (11:32):

So we each have that. Grace, what did you choose as your fourth?

Grace Broadbent (11:36):

I have four, cybersecurity.

Rob Rubin (11:37):

All right.

Jacob Bourne (11:37):

And I have regulatory and compliance as my fourth.

Rob Rubin (11:41):

And my fourth I had regulatory and compliance challenges. So that was the same. It seems like so far, well it seems like now I know for sure, Jacob and I were the same.

Grace Broadbent (11:53):

Yes. I'm the-

Rob Rubin (11:54):

And Grace we all had operational as the last one.

Jacob Bourne (11:57):

Yes.

Grace Broadbent (11:58):

I did, yes.

Rob Rubin (12:00):

Okay, so let's talk about, to me, this is really interesting actually. So Jacob and I put cyber security as first and Grace, you had it fourth.

Grace Broadbent (12:00):

I did.

Rob Rubin (12:10):

You got to tell us why is cyber security, obviously all these things are super important, but why is it not as important as these other things?

Grace Broadbent (12:18):

I don't think it's a matter that cybersecurity isn't important. I think fraud is just so extremely important and so top of mind, especially as I come from the payments world. And it is just I think, I mean obviously all five are extremely important and we should be concerned about them all, but to me, fraud is top of mind. And I think to show how top of mind it is for financial institutions, particularly payment providers, you can just look at the headlines. The past two, three weeks, both Visa and Mastercard made billion dollar or multi-billion dollar fraud acquisitions specifically to fight AI fraud. And that's not a coincidence that they both happened so recently. So Visa bought Featurespace for almost 1 billion, Mastercard bought Recorded Future for 2.65 billion, and that was both in the past month.

Jacob Bourne (13:09):

Yeah. And Grace, I think that's a fair analysis, just because fraud is something that's almost non-stop, it's a non-stop issue. Cybersecurity-

Grace Broadbent (13:18):

Yeah, I agree. I think it's an everyday issue for consumers.

Jacob Bourne (13:22):

I put cybersecurity even though it's not necessarily going to have a major breach every day. When they do happen and I think generative AI's ability to generate malicious code and also pinpoint weaknesses in security systems, means that we have an existing problem with cybersecurity that caught so massively damaging financially and into organizations reputations, that I think that that could just escalate even more from generative AI and become a really catastrophic problem.

Rob Rubin (13:49):

I agree that it's like fraud is the everyday problem, but cybersecurity is the nightmare scenario, the thing that shuts everybody down. And I think we've seen tastes of it where an intern by accident pushes code to production and takes down, like didn't that happen with Microsoft? Like a whole bunch of stuff got taken down for Outlook?

Jacob Bourne (14:10):

Yep.

Rob Rubin (14:11):

So that wasn't even malicious. So that's why I put cybersecurity number one. And I think we kind of agree with you, Grace, in that fraud for both Jacob and I is number two. Like the idea of deepfakes, the ability for AI to figure out how to launder money better than the money launderers.

Jacob Bourne (14:30):

Yeah. And I think with the fraud and the AI deepfakes, it's getting worse almost by the day as well because now you have these AI clones that are clones of people that are getting better, more realistic.

Rob Rubin (14:44):

I know.

Jacob Bourne (14:45):

So a clone that can sit on a Zoom meeting with you with co-workers and it's not quite convincing but it's getting there. And the audio deepfakes, you have instances where someone's voice gets cloned and then the bad actor calls a family member, says, "I'm in

trouble. I need money." These are really nightmare situations for individuals. So generative AI just kind of enhances that risk.

Rob Rubin (15:08):

Yeah. Now, Grace, I'm going to pick on you again because here's the other big difference.

Grace Broadbent (15:12):

Go for it.

Rob Rubin (15:13):

Regulatory. So, again, obviously it's important, but Jacob and I put it at number four and you said it was the second most important. So we'd love to get your rationale.

Grace Broadbent (15:23):

I'm again coming from the payments angle, consumer facing angle, because that is my bread and butter of what I look into every single day. And in terms of financial inclusion specifically, I just think the issues around bias and fairnesses is such a large concern. Biased AI systems result in unfair and discriminatory outcomes, whether it's denying a credit card application, a mortgage application, whatever the case may be, it can really have big impacts on everyday consumer lives if they get it wrong.

Jacob Bourne (15:56):

Yeah. And I think we can all be in agreement that all of these five are really serious issues and this one is no different. But I think the reason why I put it as my four is just because maybe the consequences aren't immediately catastrophic in the way that a major cybersecurity incident is.

Grace Broadbent (16:14):

Yeah, that's true. I think I'm looking at everything from more short-term lens and you guys are looking at the long-term, wider lens for sure.

Rob Rubin (16:22):

I feel like with regulatory, it sort of slows things down. It's like puts a spanner in the works in that it'll slow development down. Maybe for a good thing, or on the other side, something bad

happens and then everybody gets locked down for a little while. But I just don't see it as, if I'm ranking risk, I just don't put it up in the same area as a cybersecurity catastrophe.

Grace Broadbent (16:47):

Right.

Rob Rubin (16:48):

That could really blow something up.

Grace Broadbent (16:49):

Right.

Rob Rubin (16:49):

Or a major data breach.

Jacob Bourne (16:53):

Yeah, it's kind of more of a slow-moving insidious problem than something that's immediately nightmarish.

Rob Rubin (16:59):

Right. And then the last one, which we all said the last one, probably the people at the banks are laughing that we put operational last because none of the others work without it.

Jacob Bourne (17:08):

Right.

Rob Rubin (17:10):

But obviously systemic errors in your data sets or in your systems cause systemic problems and that would be the biggest problem that we can occur. So I just thought maybe I could quickly just review where we were at and maybe we can come up with sort of the ranking.

Grace Broadbent (17:28):

Well, it's two against one right now.

Rob Rubin (17:30):

I know.

Grace Broadbent (17:30):

That's not fair.

Rob Rubin (17:32):

Well, I=it's more than two against one because I'm the host.

Jacob Bourne (17:34):

Right.

Grace Broadbent (17:36):

Yeah, this is really not working.

Rob Rubin (17:37):

I'm going to say fraud and cybersecurity are sort of a one and a half. I'm going to go there. But I do think that privacy and security are, I'm going to put above, I'm going to say-

Jacob Bourne (17:48):

And the three of those are closely related too, is the other thing.

Rob Rubin (17:51):

Yeah, so it's like one and a half, one and a half, three, and then regulatory operational is four and five. So cybersecurity slash fraud, then privacy security, then regulatory, and operational. So I think we've solved some problems for people today, huh?

Jacob Bourne (18:12):

Yeah, I think we came up with a good threat mitigation prioritization list or something.

Rob Rubin (18:17):

Exactly. I think so. I want to thank you guys for coming on today. It's been a lot of fun.

Jacob Bourne (18:24):

Yeah, thanks for having me.

Grace Broadbent (18:25):

Yeah, thank you guys.

Rob Rubin (18:26):

Always fun. And I want to thank everyone for listening to the Banking and Payment Show, an eMarketer podcast made possible by TikTok. Also, thank you to our editor, Victoria. Our next episode is on November 12th, so be sure to check it out. See you then. Bye guys.

Grace Broadbent (18:41):

Bye.

Jacob Bourne (18:42):

Bye.