

# Barefoot Innovation Podcast with Chris Calabia, Senior Advisor, AIR; Head of Central Bank Digital Currency Programs, Massachusetts Institute of Technology Digital Currency Initiative

**\*Note that transcripts may sometimes contain errors and that transcript timing notations do not match the posted podcast**

Jo Ann Barefoot: I have been looking forward for a long time to today's show because my guest is Chris Calabia, and Chris and I have known each other for a long time, and we are going to be talking about AI. Chris, welcome to the show.

Chris Calabia: Jo Ann, thank you so much for having me. It's great to see you again.

Jo Ann Barefoot: It's fantastic to see you. I'm really excited about our conversation. You and I, I haven't counted when we first met, but I remember that you were at the Federal Reserve Bank of New York and invited me to come in and speak to a group there. And ever since then, we've been collaborating as you've worn some different hats really, I have too, on how technology is changing the financial world. So I'm going to start by asking you to just share with us your background and then we're going to talk about the paper that you have written for AIR, but give us the overview of your background.

Chris Calabia: Sure. Well, Jo Ann, like you and like many of your colleagues at AIR, I'm a former regulator. As you said, I was at the Federal Reserve in New York as a supervisor and banking regulator for almost 24 years or so, mostly working on regulatory policy, but I was also a line level regulator as well, dealing with firms directly. And that included also spending time in Basel at the Bank for International Settlements, working with the Basel Committee on Banking Supervision, the global standard setting body for bank regulators. And I subsequently worked with the IMF on some assignments, helping to provide advice to regulators in a couple of different countries.

And when I left the Fed, I went to the Bill and Melinda Gates Foundation, and there I worked in regulatory policy reform and technology innovation to promote access to financial services among people who are underserved or unbanked, and especially Sub-Saharan, African and South Asia. And from there I went on to work as a regulator with the senior leadership team at the Dubai Financial Services Authority, so one of the premier global hubs for banking, certainly in the Middle East and Northern Africa. And now I'm at MIT where I work with the Digital Currency Initiative and we do research on the future of digital currency.

Jo Ann Barefoot:

So this is such a fascinating background and I know people just heard it, but I'm still going to reinforce it. You've been a US bank regulator, you've been a leading financial regulator outside the US and you've worked for the Gates Foundation with its amazing work in financial inclusion around the world, including, I know you were involved in their big project on the future of the Central Bank globally, and you have the academic background as well. So you have looked at these challenges from a lot of perspectives and just bring a wealth of insight to everything.

And I will also say, I remember feeling when I first did meet you that you were way ahead of your time in terms of regulators who were really thinking about technology change. Regulators always think about technology change of course. But as we entered into the digital age, you were really one of the outstanding thinkers who was seeing that this was going to be bringing a tremendous amount of profound change to what regulators need to do and what the industry is doing. So we are thrilled to have you as an advisor to AIR, and we did ask you some months back if you would take the lead in writing our seminal paper in a series of work that we're doing on artificial intelligence and especially generative AI as it has captured everyone's imaginations. And so let me ask you to talk about the paper.

Chris Calabia:

Yes. Well, I had participated in one of round tables that you organized, Jo Ann, last year through AIR. And I have to say that getting ready to participate in that round table on generative AI gave me an opportunity to do a little bit more reading. And then when I participated in the conversation, I was just fascinated by the people you had assembled leading thinkers in regulation and finance, and some people worked in technology as well. And it really inspired me to think about if I were a regulator today, what would I want to see AI try to help us do better? And so I started to write down some things on paper, and I shared a very early version of a note with one of your colleagues, Shelley Anderson at AIR, and she encouraged me to keep writing. And then I shared it on with Nick Cook and Mariama Jalloh-Heyward, and you had a chance to look at it as well, and you all encouraged me to keep writing.

And so I was delighted to have this opportunity to work with all of you and with some early thinking that Joe Adler, your former head of content had worked on as well. And it was this great opportunity to explore the topic of generative AI and think about it from a regulatory perspective and what might be the impact on regulators, and importantly, what might be the outcomes for financial consumers.

Jo Ann Barefoot:

Absolutely. So the paper at its available on our website and we will link to it in the show notes, it is called AI Transforming the Future or Triggering Fear, which I think captures the sort of binary, if not schizophrenic nature of the challenge that I think we're all feeling as we look at AI and we have some high hopes and also plenty of fear and trepidation. You started the paper with a very creative offering of a couple of scenarios and one pretty utopian and one pretty dystopian. Tell our listeners a little bit about that.

Chris Calabia: Yes. Well, as we were completing the first draft of the paper, I think several people at AIR said, "Hey, it would be really interesting if you could imagine for us what the impact would be in a very simple way. And maybe it's like telling a story or something." And I don't think your team knew it at the time, but I'm a frustrated novelist at heart. So I had an opportunity to exercise that muscle a little bit. And your team gave me great license to just imagine maybe two scenarios and what might a utopian scenario look like with generative AI helping and what might a dystopian scenario look like. And so it was a really fun opportunity to spend a couple of pages on each scenario and imagining what could go really well and what could go wrong. And so I hope that readers enjoyed that brief introduction to what the positives and perhaps some of the downsides might be of using generative AI in regulation.

Jo Ann Barefoot: I want them to read the paper and I'm glad we're teasing it, but nevertheless, give a little bit of color around-where your imagination took you.

Chris Calabia: So thank you very much. I imagined a day in the life of two different financial consumers. One is a young attorney in Chicago who trusts her generative AI app on her phone or on her computer quite deeply, and it helps her a lot, helps her manage her calendar, it helps her manage her finances and so on. And it comes up with this really interesting idea to help her save some money and at the same time pursue a long term investment goal that she has. And so there the storyline emphasizes what can happen if we trust this technology and completely, and in many ways, this technology knows a lot about this young attorney, Margarita, and it can see her calendar, it sees who she follows online, it knows about her hobbies and her interests and so on, and it blends all these things together and is able to help her manage her life a little bit differently and better than she might be able to do that on her own.

The opposite scenario reflects a big breakdown in trust essentially, that the technology doesn't have our best interest at heart ultimately, or comes a little bit out of control. And so I portrayed the day in the life of, I'll call him a soccer dad, Javier, who is a parent and he's struggling to manage his financial circumstances because something went awry with generative AI in financial services. And it kind of details some of the challenges that this huge breakdown in trust in computers and networks and so on that resulted.

Jo Ann Barefoot: I know everyone will enjoy reading them. It gets us off to a great start, and again, it expands our imagination right at the outset. I think the challenge is or the, it's challenges and opportunities we're facing. I think a great way to frame this foundational conversation is around the two, the upside and the downside. So why don't we start with the good news? If you put yourself in the shoes of a financial consumer and also potentially the financial regulator who is trying to make sure that that system is working well for everyone, what are your hopes that we might be able to reap benefits from as a result of AI and generative AI?

Chris Calabia: So if we look at generative AI in particular, this is a technology that learns from looking at data that's available to it and it can learn from a huge variety of data,

forms of data. And I think what is a potential game changer for consumers and regulators and financial institutions is that it will help us to, in particular, to process the kind of data that we haven't had good tools for so far. So if you think about the data that financial institutions generate and that financial consumers have and regulators have to work with, we have hundreds of years of experience working with what's called structured data. So these are things like quantitative information, financial information, income, expenses, those types of things. And this data can be storage in tables, it can be analyzed in spreadsheets and so on.

So we have a good number of tools we're dealing with that, but by some estimates, that's only about 20% of the data that we use in financial services. And there's 80% of the data that's not been well addressed. And this is called unstructured data, and this is everything else. So it could be pros, it could be email messages between consumers and their financial institutions, could even be complaint letters, could be social media postings, it could be notes in a loan officer's files, it could be management analysis reports and so on. So we have tremendous amounts of data and sometimes it's stored in databases and so on. But it's hard for us to access this type of unstructured data because it's not easy to categorize it, not easy to break it down into fields and into tables and so on.

And what generative AI may help us to do is to better process that 80% of data that we haven't had good ways of processing up until now because it can read tremendous amounts of information and it can assimilate that information and look for patterns in that information and perhaps help us to process it better. And that may open up new opportunities for financial consumer.

So for example, right now, consumers who want to seek access to credit have to have a record of their financial health basically. And traditional credit metrics looks at things like your ability to borrow funding and your income and so on. And if you think about people who've been excluded from the financial system, they may not have that financial records because they may not have a long-standing bank account with a particular firm and so on. There may be other data though that they have that generative AI could help with financial institution to process. And by processing this alternative data, it may open up different financial services to them. Perhaps they can get access to credit for the first time or maybe other financial products and services. So I think that's really the big game changer, is the ability to help us better process unstructured and as well as structured data.

Jo Ann Barefoot: Do you see potential upside, I'm going to guess you see potential downside, but do you see possible opportunities around dealing with financial crime and money laundering as well?

Chris Calabia: So that's a huge area of interest for me. When at the Federal Reserve, one of the topics that I focused on was anti-money laundering, and that is a subject where we are awash in enormous amounts of information. So in the report, for example, I cited, this is one of the three areas where I thought generative AI could make a difference. And anti-money laundering matters because it's a

crime that affects so many different aspects of our society and some of the impact of these crimes can fall on the most vulnerable in society, things like human trafficking or illegal trafficking and narcotics and so on. And we have tons of data yet we don't have a good way of... We haven't demonstrated we have a good way to stop this. So the best estimates are that of all the money laundering that takes place in the world equivalent to something like 2.7% of global GDP is thought to be laundered funds, we catch only about 1% of that.

And even though we have millions of suspicious activity reports and we can monitor transactions almost actually in real time nowadays, we haven't had a good way to tackle this problem. And so I do think the generative of AI is going to be an important tool in helping us to explore and study this data potentially in real time. And in the report I detailed some research that the Bank for International Settlements has been doing to make use of a variety of forms of artificial intelligence to help us better understand this crime and help us to reduce the impact on society of money laundering and financing a terrorism.

Jo Ann Barefoot: So Chris, we know there's a lot of upside opportunity and we know there's a tremendous amount of downside risk. As you worked on the paper and have been thinking about these issues, what are the main risks that you're concerned about in the realm of financial services?

Chris Calabia: Yes, and there are some risks. And I know that we're emphasizing today some of the positive outcomes, but we should also think about those downside risks because this is a new technology. And because it's a new technology, we don't really still have a good understanding of precisely how it works. And so for example, we know that GenAI is able to survey an enormous amount of information, both online or in training data sets and so on. But we're not always sure how it's spotting what the answer should be to questions based on those data sets. So for example, there may be times that we don't understand a decision or a recommendation that a GenAI tool is making. So as a simple example, maybe a GenAI tool says that we should not offer a line of credit to a particular consumer. Well, in many countries, including the US, when you turn down someone for credit, you have to give them an explanation as to why you turn them down so that person can address it and work on it and perhaps qualify the next time.

But if we don't know why the tool is telling us we should turn them down, it puts the consumer in a really difficult spot in terms of trying to address it. And from a regulatory standpoint, firms might say, "Well, it wasn't our fault. The tool told us we should turn it down." And that's not appropriate. You can't outsource your responsibilities and so on. I think the other thing is that because GenAI is able to create very convincing answers in prose, and also it can create really amazing images and so on, we may trust the outputs of GenAI when we shouldn't. And there are many examples of times when people have asked questions of a GenAI tool. For example, in the paper we covered a story of a lawyer who asked it for an analysis of a particular legal issue. And that lawyer unfortunately submitted

an analysis that a popular GenAI tool developed of case law, and the case law turned out to be completely fabricated. It was what was called a hallucination.

And so we may trust the tool when we actually shouldn't, and we may not be able to tell when we can trust it and when we shouldn't be able to trust it and so on. So I think those are some of the downside risks. It's just not really understanding how these decisions are made, and not being able to explain these decisions. And there are many firms trying to do research to improve explainability and so on, but that is an immediate downturn side risk. But I think a big and overarching concern that I have is that GenAI relies on enormous analyses of data. Past data, existing data, and so on. And just by defining what data it's looking at, we may be exposing the tool to data that reflects biases from the past or decisions that other firms have made in the past.

So for example, if you were trying to train a GenAI tool how to assess the credit worthiness of a particular consumer, if your data training set includes decisions that were consistently made against people from particular neighborhoods or particular types of backgrounds and profiles, that bias might get replicated in the future. And you may not realize that. And so understanding really well what types of data you're exposing the tool to and how you can put guardrails around that because it's probably impossible to develop a completely bias-free data set, but at least understanding what biases might be there will be very important. So understanding how the decisions are made, knowing when we can trust it better and when we can't trust it, I think those will be important things for us to learn more about in the coming years.

Jo Ann Barefoot: And in your research, do you feel like you could see a pathway to resolving, let's start specifically with the bias issue. What should people do about that concern?

Chris Calabia: Yeah, the bias one is a really hard one, because again, if for generative AI to work we have to train it on data sets, we have to provide it with useful data. And so some models just go out and look at the internet and can pull in lots of information, some of it will be factual information, some of it will be information from sources that maybe we shouldn't be trusting and so on. But with training data and when we're developing these foreclosed use within, say a financial institution, we have to provide that data. But how do you find a bias-free data set? Just deciding what data you want to record might reflect a bias that we don't realize. And so I think it's going to be very important that we think very clearly about ensuring that we're aware of what biases may be built into data sets, because we probably can't create a completely bias-free data set, but we can at least be aware of them and maybe build safeguards and guardrails around those so that we are not necessarily going to replicate the same problems from the past.

Jo Ann Barefoot: Looking at the fraud and money laundering side as well, I think maybe this is one of the areas where people have been most concerned about the advent of voice cloning and deep fakes and so on, tricking people into making decisions that they shouldn't make. How do you think about how to counter those kinds of

problems? And are we going to need a, I guess I have a working theory that we're going to need AI to protect us from AI that's attacking us, but what do you think's ahead there?

Chris Calabia:

Yes. Well, without giving too many spoil alerts, that was the second story. The dystopian story is the potential for AI to send very convincing information and messages to us, and perhaps we'll listen to doing things that we might not otherwise do. So I'll leave it to the readers to read that very short story to get some insight into that. But it's a big challenge. And if you think about the recent controversy with OpenAI and the potential allegation that they may have replicated a famous Hollywood actress's voice or use in one of their tools, that's a concern. And if someone with that much ability to have resources and access to try to stop this type of thing if she's facing challenges dealing with these threats and these issues, it's a concern for I think all of us as well. I don't have a good answer for exactly how we can not monitor that and manage that.

I did come across a good deal of research and talk to some people who are looking at ways to audit for gen AI developed content and to try to determine what might be behind some of the answers that it's giving us. Some firms are doing research on how to improve the explainability of their answers so that at least you can know what sources it's relying on when it's coming up with some of his answers. And you can follow up and make sure that if it's, for example, citing case law, that that case law really exists. And there was a famous story from last year of a lawyer who unfortunately relied on a gen AI tool and it developed case law that wasn't true. It was all completely made up and so on.

And so I think for firms to put those guardrails around, help us understand the sources that the AI tool is relying on, and again, back to the data sets, thinking much more about the types of data that we're going to introduce to a tool and making sure that we understand what biases and what information and what decisions these may reflect will be very important.

Jo Ann Barefoot:

We touched already on the opportunities for regulators to be using more unstructured data, getting a richer set of information as they're looking at risk and compliance issues. Can you talk more about that? Do you think that there is likely to be transformation of the financial supervisory process and what might that look like? What should we strive for? Are there some principles that you think we should have top of mind as we move along this journey?

Chris Calabia:

Yeah. So if I think about consumer protection in particular, I think that there are a lot of possibilities for us to change the way we approach this as regulators. And so in particular, consumer protection involves trying to make sure that everyone is treated fairly and that institutions aren't biased against particular groups of people or aren't consistently making the same mistake with regard to consumers and treating consumers differently or in a manner that's not consistent with laws and regulations. And right now, we have to rely on complaints as regulators to come in. And when I was a regulator, I oversaw a consumer complaints section, one of my responsibilities. And so someone had to read every email that came in

and to make sure that we understood what the complaint was and also whether it was something that we could address, because sometimes people would send in complaints about things that we had no jurisdiction over or no responsibility for. And so we couldn't necessarily address those things.

But that take a lot of time, and it also takes time for the consumer to write down and send a letter or an email or maybe to try to call somebody. And so there's some really exciting work being done, including in the Philippines, to make use of chatbot technologies to make it a little bit easier for consumers to interface with the relevant regulator and to engage in an online chat with a chatbot about their concerns, and the chatbot can help ask questions and narrow down what the issues are. And then if it's something that the central bank or regulatory authority can address, it can forward that complaint to the right department within the central bank and maybe the right person who's an expert on that. And so I think there's some really exciting opportunities there.

In addition, so much is on social media these days, and there's so much information where people may post on a popular messaging board or a website, a concern or complaint that they have. And right now we don't have a good way of going through that. I mean, you could pay someone, I imagine, to look at social media every day and look for keywords, and we can even now do that in an automated fashion. And when I was a regulator, I did try to do that. I set up an alert on a popular search tool to look for the names of the firms that I was supervising and try to find newspaper articles about them that might be of interest, but often return so much information that I couldn't read every article. And it was hard for me to keep up with that.

With gen AI, it may be possible for us to look for keywords or to look for sentiments that consumers are sharing online. There's social media posts or other places, and look for patterns and see, "Oh my goodness, there are a lot of questions about one particular firm and how they handle some particular consumer issue. Maybe we should spending more time looking at that particular firm." So I think that there may be clues in social media and other outlets that reflect unstructured data, texts and pros that we haven't been able to process very well before, and Gen AI may help us to do that better.

Jo Ann Barefoot:

Yeah. We have another white paper coming out shortly authored by my colleague Nick Cook, who is our chief innovation officer and had been the director of innovation at the Financial Conduct Authority in the UK. And Nick has written a paper about how the regulators can navigate their own odyssey, as he calls it, through the tech transformation that's underway. How do they need to think about changing their own technology? How do they need to think about getting and developing the talent that they need with the requisite tech skills? And how do they think about culture change? So Nick draws on his experience as a financial supervisor. And so we're excited about this. I know we've got a number of regulators very interested in reading it. It'll be out soon, but let me ask you about this. What is it going to take for the regulatory bodies themselves to build the capacity or be sure they have the capacity to oversee this



transformation of the financial sector, which is happening due to technology in general? And I think AI in particular. Do you have advice for your former brethren in the supervisory world?

Chris Calabia:

Yeah. So for my financial regulatory brothers and sisters, very happy to share some thoughts about that. I think like with any new technology, a big thing is education, learning to understand it, learning what its abilities are, what some of the downside risks and so on are, and just educating oneself. And I must say that regulators have hard jobs because they're dealing with very large institutions and sometimes a large number of institutions. And these institutions are stewards of society's wealth, and so they want to make sure that they're very careful.

And so sometimes regulators can be accused of not moving as quickly as the industry does in thinking about adopting new forms of technology and so on. But this is an area where it's catching on so quickly, it seems, that it's going to be important for regulators at least to educate themselves about the technology and how it may transform institutions, financial institutions that they regulate, but also to think more about how they can use it themselves as regulators to help them tackle the information that they already have at hand. And that's what I think one of the things where gen AI may help regulators to make better use of the data that they already have access to or can gather pretty easily.

I think the second thing is that we are going to need to think a little bit more about the skill sets that we'll need to develop internally. Just as a simple example, it's one thing to gather data and put it into spreadsheets and tables and so on, and then to start to analyze it. And for some regulatory authorities, historically, the biggest challenge was putting the data in a useful format. And so if you think back several decades ago in the US but even today in some countries, regulators would receive a lot of information on paper. And so they had to move that information into a database and that took time. And even if they receive things like in PDF forms or emails and so on, it takes time to move that from an email or from a website into a database and manipulating it. And they have large numbers of staff who in some cases have to do that.

Gen AI may simplify that. It may become easier for regulators to gather that data and put it into a useful format, which means that their staff, instead of having to do the tedious work of moving data and making sure that they got it right and that they didn't make any mistakes and so on, and putting it into tables, they'll need now to analyze that data. And that is a different skillset, and it may require hiring more people who have that skillset of being able to work with and analyze the data as opposed to simply people who are making sure that the data is correct and captured in tables and data sets appropriately. And those are really important jobs, but gen AI may make this easier. And so we may need to think about the skillset that we have in regulatory authorities as well, people who can spend more time analyzing the data and analyzing the output because the gen AI is going to be very helpful in getting us started, but we're probably still going

to need human intervention and analysis, especially with the most important highly sensitive issues.

Jo Ann Barefoot: It seems like there's a certain amount of emerging consensus, and this might be overstating it, that there's going to be a need for the so-called human in the loop on practically everything in the space at least for the foreseeable future. It's hard to imagine just unleashing AIs without having a human being there to check what they're doing.

The scale side issue is significant. And I know there are people who worry that the robots are going to take the jobs away, and I'm sure robots will take plenty of jobs away. But I do think that with financial supervision and oversight, that mostly what we're going to see here is an empowering of people to really use their expertise. You and I both know many, many financial supervisors and regulators, and they understand these complex systems deeply, but I know a few years ago the FDIC estimated that their examiners spent 400,000 hours a year just doing data entry into spreadsheets. If you can have the robot do that, and then the expert can spend the time analyzing that information, how much more could we find in terms of finding risk earlier or finding non-compliance early and preventing the harm that develops if these problems take hold and are left neglected, which happens throughout the industry, it seems to me.

Chris Calabia: I think the hope is that automation and tools like artificial intelligence will take away the drudgery of data management and give us time and, more time rather, for more creative and analytical thinking, which I think will make the role of a supervisor or financial regulator even more interesting. If you spend a lot of your day gathering data and making sure that it's correct and so on, that's tedious and that gets to be not so interesting. But if you can spend more of your time analyzing the data and trying to understand the story that the data might be telling us, I think that makes for a much more interesting job as a regulator.

Jo Ann Barefoot: Yeah, I couldn't agree more. As you did your research on this, Chris, was there a moment that you recall that surprised you, that really caught you by surprise and made you say, "Oh my goodness, this is a game changer?"

Chris Calabia: I think one game changer certainly for consumers is my sense is that if this goes well, consumers may feel of stronger connection to their financial institution. So right now, if you have a concern or complaint and you call, you might be dealing with a chatbot that can only handle very simple types of problems and issues, or you'll be put in a really long queue and you have to wait for a human to come and answer your question or talk to you. Maybe they have to look something up and they have to get back to you later. And then when you call back, you have to talk to somebody else who who's not familiar with their case and so on.

Gen AI, if it's done well and if it works well, potentially you could have almost direct access to the automated equivalent of a private banker, someone who knows your file and is familiar with the communications you've had with the firm and can take the time to explain things to you and answer questions you

might have about the products that you're purchasing from the firm or about how loan terms work and so on. And you might be able to put at least some basic questions to it, and it might be able to give you at least some early answers.

Now for more sophisticated things, it might still require human to double check the information and maybe someone to walk through with things. But for common types of questions and common types of concerns, it may be a lot easier for you to speak with something that sounds like a person and can answer questions in ordinary language. And if you don't speak the dominant language in your country, perhaps it can speak a language that you do speak well and that you understand well and translate the information for you and so on. So I think that could be a real game changer that consumers may feel a stronger connection to their institutions and feel better about the types of products and services that they're buying.

Jo Ann Barefoot:

And another dimension of what you're saying is the potential to democratize finance more fully and make it profitable and affordable for providers to serve people who have smaller balance accounts or less wealth or less income, and are maybe not always getting that customized service, as you say, in the cost structure that we have today. I mean, you've been part of, including in your work at the Gates Foundation, part of this huge sweep of transformation that has been bringing more financial access to more people through the cellphone and through digital financial services, people who were never served by traditional banks around the world, because again, they didn't have enough money to be worth building branches to serve them and all of that. If these tools could really drive down the costs of giving people good service, it could really be transformational.

Chris Calabia:

So, Jo Ann, you're absolutely right. I had a chance when I was working on financial inclusion issues, including while I was at the Gates Foundation to meet with some people whose lives were changed when they got access to digital financial services for the first time. These were people who were typically from low income communities, may not have had high levels of education or even literacy, and they weren't able to open up financial accounts in the past because the fees were too high or no bank wanted to build a branch in their neighborhoods or in a remote area because it wasn't profitable for them to do so. But because they had access to mobile phones, mobile network operators made it possible for them to engage in very basic financial services. So these weren't truly banks, but they were able to make and receive payments.

And I met especially women whose lives were changed as a result of being able to receive payments digitally or make payments digitally. And so instead of having to wait in line to be paid if they worked at a factory and lose potentially up to a half a day's wages or a day's wages because the lines are so long, they would get paid instantly, and that meant that they had more time for other things. And some of the women that I met started small businesses and did other things. And so seeing how that technology changed their lives really got

me interested in the topic about how we can empower people using technology. And so gen AI might be one of those tools that can help us continue that story. And I say it might because as we talked about earlier, there're also some downside risks to the technology. Of course it's not a panacea.

Jo Ann Barefoot: I'm going to link in the show notes to a recent episode that we did on AI ethics, and it included looking at the potential of so-called agentive AI, or having an AI agent that can be optimized to have the best interest of the consumer at heart, for example, and potentially guide people and help them, again affordably advise them or help them make good choices or alert them if a problem is developing. And again, democratizing those kinds of capabilities beyond anything that we've ever been able to imagine before.

I want to say that your paper, as I mentioned, is the kickoff of a series of work that we're doing at AIR. We are going to be putting on a tech sprint in West Africa focused on payments fraud later this year, and it's got an AI theme running through it. We envision doing a series of text sprints over the next few years on these issues and additional thought leadership, and including that we have out currently a call for papers and are inviting anyone and everyone to submit papers that are looking at both the benefits or risks and solutions around AI in the financial sector. So I'll link to the instructions and criteria on that in the show notes and hope people will consider participating and getting involved with us.

As we move toward wrapping up, Chris, is there anything that you would like to add?

Chris Calabia: Yeah. So I had a lot of fun writing the paper, and thank you so much for the opportunity to work with your team. And I had a chance to meet some people in the field and the industry and so on, so that was a lot of fun. But I hope when people read it, they realize that I identified more questions than answers, frankly. And so that's why I'm really excited about your call for papers because I think that we'll have an opportunity to hear from people all over the world in different aspects of financial services and regulations who may have identified both opportunities and risks that we didn't talk about in our paper. And so I'm really excited about that and hope to get an opportunity to talk to some of those people who worked on in the future.

Jo Ann Barefoot: That would be wonderful. I'm looking forward to that too. Financial services, finance, money, it's always been about technology one way or the other. You can look back to the invention of metal coinage and all the way through, but there's never been a phenomenon that has been so profound and so fast moving as generative AI. The experts are saying it's the most rapidly adopted technology in the history of the world across the board, and certainly in finance. It's got the potential to not only generate solutions but then generate on top of that, it could consistently use what it's learning to help us learn more and more. And in my years, many years of working in this space, I've never seen a change, a tech trend that hit the top short list of every leader in the whole ecosystem in such a

short time as this. If you think about cryptocurrency, for example, it had years of developing sort of in a stealth mode completely off the radar, gus stealth or whatever you want to call it, but not on the radar of the financial regulators.

When generative AI came out at the end of 2022, all of a sudden, every corporate executive, every bank examiner, every member of Congress, you name it, everyone was experimenting, was getting the briefing, was thinking about what is this going to do. Having that come onto everyone's plate at the same time just creates this potential for explosive rapid change, both good and bad, and we want to do our part to try to help figure out how to manage that, how to channel this to the good and away from the bad. And your paper is absolutely, it's a fantastic ground building ability to tee up the issues. And then we're going to go from there.

Chris Calabia:

Thank you so much. Well, if I could just pick up on that comment, I think it's true that a lot of people are very excited about this technology. Some are also very anxious about the technology, and our paper talks a little bit about that dichotomy. And I was just at a conference recently with some senior bank executives and they were asked in a poll, how many of them are making use of gen AI today or artificial intelligence today? And most of them said that their firms were not using it. And my response to that poll live was, "I suspect if you look carefully, you are using it. You just don't realize it."

I think that even some well-meaning staff may be experimenting with these different tools which are free and online and easy to access, and that is both exciting and a concern. It's exciting because it means your staff is trying to figure out how they could use this tool better. It's a concern because we still don't really understand this technology that well. And especially if people are putting in sensitive information into these public tools, we don't know where that information goes. We don't necessarily understand how these different bots and chat tools and so on are processing that information. We can't always explain the decisions that they're making.

So I think it's really imperative for financial services executives as well as financial regulators to take a careful look at what their staff may already be doing and make sure that people understand that you have to adopt a safety first type approach with this new type of technology. It's exciting, but you want to make sure that there are guardrails around it so that you aren't inadvertently subjecting yourself to all the downside risks that we mentioned in the paper. You and I tend to be optimists, I think, and we want to be optimistic about this, but there's a lot that could go wrong, and it's really important for people to pay careful attention to this in their own organizations today.

Jo Ann Barefoot:

That is a great note to end on. I really, really appreciate it. We will put information in the show notes about where to get the paper and how to reach you, and that is at [regulationinnovation.org](https://regulationinnovation.org), so we hope everyone will join us for those next steps as well. Chris Calabia, I cannot thank you enough for being with me today. It's been fantastic talking with you.

Chris Calabia:

Thank you so much for having me. It's always inspirational to speak with you, Jo Ann.