



Artificial Intelligence and Prudential Supervision

Speakers:

Manoj Singh

Deputy General Manager, Central Quality Assurance, Reserve Bank of India

Kangyu Wang

*PhD Candidate, Department of Philosophy, Logic & Scientific Method,
London School of Economics*

Host:

Wynnell De Landro-Robinson

Program Director, Toronto Centre

Date:

Jan. 22, 2025

Transcript:

Wynnell De Landro-Robinson:

Hello everyone, and welcome to another episode of Toronto Centre's podcast series. I am Wynnell De Landro-Robinson, a Program Director here at Toronto Centre. Artificial intelligence is rapidly permeating and invading all aspects of personal and public lives. Initially, some believe that eventually, AI would replace many humans in the workforce. This belief has prompted responses from governments and industry. The question of AI's significant impact on society has already been answered. The real question is whether this impact will be positive or negative, and to what extent. Scholars have stated that AI is a functional tool for emancipating people from routine work tasks, thus expanding the possibilities of their self-realization and the utilization of individual interests and aspirations through more meaningful spending of time. On the other hand, some scholars caution that the risks associated with increased or excessive machine autonomy and limited human control are unquestionable and are compounded by the insufficient ability to monitor the performance of these systems and prevent errors or damage.

The Toronto Centre, for your reference, has produced a tactical note on the Supervisory Implications of Artificial Intelligence and Machine Learning, which is available on the Toronto Centre website. This particular podcast will focus on using AI and prudential supervision and the possible benefits and detriments.



We will discuss AI from three perspectives:

1. A tool for supervisors
2. Supervising its use by the industry
3. Regulatory changes required

Today, I have joining me two panelists who have decided to look into AI and its use with respect to prudential supervision and generally. First, I have Manoj Singh. Manoj Singh has experience in supervisor policy work at both the national and international levels, and is currently working as a deputy general manager of the Department of Supervision at the Reserve Bank of India. At the outset, it is mentioned that the views expressed here are personal, and should not be interpreted as the views of the organization where Manoj is currently working.

Second, I have Kangyu Wang. He's a PhD candidate in the department of Philosophy, Logic, and Scientific Method at the London School of Economics and Political Science. He works on decision theory and practical reason, moral and political philosophy, philosophy of economics and finance, and philosophy of artificial intelligence. He's also a research affiliate at the Machine Intelligence and Normative Theory Lab, Australian National University. We have two experts with different points of views.

As the financial sector is getting driven more by technological tools like AI, we might be getting in an era where the same technology is also getting deployed as a supervisory gatekeeper.

Manoj Singh
Reserve Bank
of India

So, let's start our podcast, our questions to both panelists. The first section is about a tool for supervisors. Manoj, what do you perceive as the role of AI in prudential supervision?

Manoj Singh:

First of all, thank you very much, Wynnell, for inviting me to this podcast. I think your very first question is very important for starting the current topic because, as we know, prudential supervisors have been using new technology under a broad umbrella of what we call SupTech. In general terms, soup tech is understood as the use of innovative technology by supervisory agencies to support their supervisory work. There are several enabling technologies that include artificial intelligence, or AI, that make the SupTech initiatives possible. Many supervisors have been using systems or applications that automate various processes like data monitoring or handling of customer complaints, from the time when the term SupTech or AI itself was not that popular.

So, currently one of the focus areas for supervisors is to understand how well they can discharge their core mandates by using cutting edge IT tools like AI. As the financial sector is getting driven more by technological tools like AI, we might be getting in an era where the same technology is also getting deployed as a supervisory gatekeeper. That's quite interesting because I think that the supervisors are equipping themselves with new tech solutions powered



by AI that can be relied upon to monitor all kinds of risks arising from both the traditional sources as well as new tech innovative sources.

Wynnell De Landro-Robinson:

Thank you, Manoj. I agree with you, and as you know, in supervisory work, saving of time and being able to focus on other things as in a risk-based supervision framework is always a welcoming feature. Now, we know that AI is currently being used at the Central Bank of Netherlands and CHATDNB. What other uses do you foresee for AI menage?

Manoj Singh:

Yes, Wynnell. Actually, one of the studies published by the Bank for International Settlements (BIS) points out that many of the central banks and supervisory agencies have been early adopters of AI in various areas like market surveillance, payment system oversight, and supervision, among others. But the current experience shows that many supervisory agencies are already reaping the benefits of technological acceleration. For

instance, as you said, CHATDNB, in its current form being used by the Central Bank of Netherlands, uses Generative AI to answer questions which reduces the time it takes to find relevant information and documents. That can be very important for supervisors given the time constraints that they face. As reported, we understand that the tool is also capable of offering interpretations of regulations when their original meaning is not clear.

Many of the central banks and supervisory agencies have been early adopters of AI in various areas like market surveillance, payment system oversight, and supervision, among others.

Manoj Singh
Reserve Bank
of India

There can be several other use-cases that can be very valuable for the supervisor. For instance, AI can be very useful in making the existing fraud and misconduct detection tools more efficient. Furthermore, AI can help in designing new virtual assistance, capable of providing dashboards of risk indicators or early warning systems as we understand it. Another area worth exploring would be how AI can be used in making current regulatory architecture machine-readable. That will be quite interesting, because forecasting tools used for credit and liquidity risk can be made more efficient through AI-powered high-frequency data analytics. Overall, I think it can be said that AI can be used to either radically improve the existing tools or develop new tools altogether.

Wynnell De Landro-Robinson:

Thank you, Manoj, I agree with you that AI can help with some of the routine tasks for supervisors, but I also like the fact that it would not take away from the judgment aspect of the supervisory work. So, these tools, these innovative uses of AI, can prove to be positive. However, what are the detriments for supervisors? Do you want to chime in here, Kangyu?



Kangyu Wang:

Yes. First of all, thank you for having me here, everyone. I'm super delighted to be joining you and I agree with what Manoj had just said. Perhaps what I want to add to that and answer this question is something like this. So, as far as I can see, the use of AI tools by supervisors, like the use of virtually any new technology by supervisors, is likely going to face a kind of "Catch me if you can" problem, but this time it could be specifically tricky, because the industry always adopts new technology faster than regulators, as they often have more economic incentives, and, in many cases, more resources, especially human capital. The supervisors, however, must also be more cautious and conservative, both because they are the final safety net of the entire system, and so they really need to make sure that their technologies are going to work in an expectable manner even in unexpected circumstances, and also because there will be more discussions, negotiations, and political debates and those are very understandable, but this can slow down the process.

The industry always adopts new technology faster than regulators, as they often have more economic incentives, and, in many cases, more resources, especially human capital.

Kangyu Wang
London School
of Economics

Moreover, in the case of AI, one concern I have is that the industry and the supervisors may need different types of data to train their AI models, which can further broaden the gap between the industry and the supervisors. The supervisors may probably care more about systemic risks rather than let's say everyday trading data, but while you have enough everyday trading data to train good AI trading algorithms, it may turn out to be much harder for you to collect enough data about systemic risks because they do not happen frequently. So, this data-thirsty nature of AI technology may make it easier for some parties to train good AI algorithms than for other parties. In our context, my worry is that this dynamic may not necessarily work in favor of the supervisors.

Wynnell De Landro-Robinson:

Thank you. Kangyu. Manoj, do you want to add anything to this?

Manoj Singh:

I would just like to add that the BIS study notes that AI presents both opportunities and challenges. The use cases of many supervisory AI tools are still in the early stages, and the results can be fully assessed only after some time. Moreover, any use of general-purpose technology by the supervisors needs to be preceded by a very comprehensive understanding of technological limitations so as not to set any wrong expectations for the stakeholders. Though active participation by supervisors in the use of new AI tools can definitely serve as a catalytic pillar, the future growth part of the technology is not entirely in the hands of supervisors. For instance, getting access to high performance computing power in the face of a limited supply of



high-performance chips is already recognized as a significant bottleneck. Further, supervisors may also need to put equal emphasis on building human capacity for a vast changing technological environment. It would be hard to integrate risk and AI without commensurate support and buy-in throughout the entire value shape.

Wynnell De Landro-Robinson:

Thank you. Let's go on now to AI and the environment. So, AI mechanizes the environment in which it is used. It enhances supervision. How can this be achieved without removing the supervisor's judgment and responsibility? Manoj?

In machine-driven information processing, errors can arise from many sources, so putting a human in the loop forms the core part of any discussion on fixing accountability arising out of the responsible use of AI tools.

Manoj Singh
Reserve Bank
of India

Manoj Singh:

I think no matter the current state of developments in AI, the supervisor's judgment and responsibility are here to stay. In machine-driven information processing, errors can arise from many sources, so putting a human in the loop forms the core part of any discussion on fixing accountability arising out of the responsible use of AI tools. And, as you know, even the EU AI Act recognizes that AI systems must be designed in such a way that they are effectively overseen by natural persons as a safeguard against malfunctions to allow human control or interference as warranted. The outcomes generated by AI need to be clearly understood and challenged by an expert human when needed. Further, we should not forget that as a part of the basic principle for governance, good AI needs to be explainable. We must also factor in that regulatory guidelines are drafted in a certain way, leaving their interpretation subject to evolving circumstances. Many times, we need to read between the lines. Now, AI tools are becoming smart enough to give a convincing reply. They are also prone to hallucination.

In many situations, I think the best bet for the supervisors would be to rely on expert human judgment for taking a final decision.

Wynnell De Landro-Robinson:

Thank you. Kangyu, do you have anything to add to that?

Kangyu Wang:

Well, just like what Manoj has mentioned, whenever AI tools are used in decision making, there will obviously be an interpretability problem or shall we say, explainability problem, that is the famous "black box" problem, which I'm sure everyone understands, and also a responsibility problem, which Manoj also mentioned. Both issues are well known and long studied in the industry as well as in academia. But insofar we have still not found much, too much consensus



even on some basic points from a philosophical point of view. There are many AI ethicists working on the responsibility problem, and why or whether at all, we should always put a human being in the loop. It is really morally necessary, more important, or is it just for our psychological certification? Most people think that there is some kind of moral importance of the human in the loop, but it's going to be challenging to articulate what exactly the reason is for that. Well, I'm not working too much on this responsibility issue, so I wouldn't say too much about it. What I want to mention a bit more is that it's not even entirely clear and that least remain debatable for what reason and to what extent interpretability matters, and there are of course many practical and contingent reasons, of course:

people, consumers, regulators, voters, et cetera, just want to sort of get some kind of explanation. But from a philosophical point of view, it's debatable and the way we understand this problem could influence what we do about it.

Now on the positive side, the interpretability problem of AI may not necessarily be worse than other problems with non-AI computational models, which can be also extremely complicated and

extremely hard to explain, for instance, and this can be true for say, weather modeling or natural disaster modeling. So, this makes us wonder, what's the difference, the real difference between AI and non-AI computational models, which are already hard read, interpret, and do not forget, human beings are also very hard to interpret. We are also black boxes in a certain way! We can generate explanations when asked to do so, but to what extent should people trust our human made explanations? That's another question. Now, coming back to the role of supervisors and the central bankers, I guess leaving all those issues we have been talking about aside, there may be something additionally interesting. It's already hard enough for the market to guess what the central banks may do in their decision making but suppose that the market now knows that the central bankers are using AI models to facilitate their decisions. What will the market do? I have no answer. There are just lots of question marks.

It's already hard enough for the market to guess what the central banks may do in their decision making but suppose that the market now knows that the central bankers are using AI models to facilitate their decisions. What will the market do?

Kangyu Wang
London School
of Economics

Wynnell De Landro-Robinson:

Thank you. I agree, we're still working out all the kinks in the AI world. Let's go now to supervising the industry with AI and its uses. How to regulate new insurance products or any financial product that is a result of AI. This is going to be quite important going forward. Manoj?

Manoj Singh:

Most of the time we use cases of any new technology by supervisors plays an important role in setting the tone for the entire market on various aspects like accountability, ethics, capacity



building, and transparency. So, in general, any financial product or service that comes out of AI would need to go through the same set of basic principles which are broad enough to be applied to different sectors and industries. For instance, as you know, DNB issued in July 2019 the basic principles for responsible use of AI, namely soundness, accountability, fairness, ethics, skills, and transparency or what we call "SAFEST." These principles are broad enough to be applied to any financial product or service that is relying on AI. Similarly, the Monetary Authority of Singapore also has introduced the principles to promote Fairness, Ethics, Accountability, and Transparency, or what we understand as FEAT in short, which can be relied on for designing any product or service that relies on the use of AI. However, a few implementation challenges may need to be tackled.

Any financial product or service that comes out of AI would need to go through the same set of basic principles which are broad enough to be applied to different sectors and industries.

Manoj Singh
Reserve Bank
of India

For instance, the regulatory and supervisory authorities would need a framework for actually applying the basic principles to live examples. Based on experience, what I see is that we may need to further harmonize the standards at the global level.

Wynnell De Landro-Robinson:

Thank you, Manoj. So, Kangyu, let's hear your point of view on this topic.

Kangyu Wang:

Thanks, and I just want to add a bit more on the AI fairness issue. So, AI fairness is a big issue and part of the difficulty is to figure out the correct or appropriate set of metrics of AI fairness, and probably even just fairness in general. For a quick example, when there is a base rate difference between two groups of people in terms of their probabilities of doing something, and we need to find a somewhat accurate way to label those who are having higher probabilities of doing it than others, it's usually mathematically impossible to achieve two things at the same time. You can't make members of the two groups have equal probabilities of being mislabeled - which is a fairly plausible way to understand the concept of fairness - and on the other hand, at the same time, make them have the same, equal probabilities of being mislabeled conditional on being labeled for "high probabilities," which is also another good way, fairly plausible way to understand fairness. So, of course, we have this problem with or without AI technology, but this means that it's probably also impossible to solve this kind of ethical problem purely by improving the technology. Philosophers love puzzles, but for those living in the real world, we need solutions.



Wynnell De Landro-Robinson:

Yes, we do. Thank you so much. So, now we go on to cyber risk. Cyber risk is a risk that supervisors are already facing in that they have to understand this risk and assess their regulatory companies accordingly. So, how will AI add to the threat associated with cyber risk? Manoj?

Manoj Singh:

That's a very important question, and I think currently, cybersecurity already forms one of the core priorities for the regulatory and supervisory authorities, as any significant cyber incident can also have financial stability implications. So, what we are seeing currently is that regulatory authorities are also conscious of the unique challenges posed by AI-related cybersecurity incidents. For instance, you might know that recently, the US Department of Treasury published a report detailing the current state of AI-related cybersecurity and fraud risks in financial services, along with the AI-led use cases supporting cybersecurity and anti-fraud operations. So, the US Treasury report notes that concerns identified by financial institutions are mostly related to lowering the barrier to entry for attackers, which will, in a way, increase the sophistication and automation of attacks and will decrease time to exploit. So, existing threat actors can develop and pilot more sophisticated malware by using AI tools. For instance, the use of AI tools for creating deep fakes can be misused for getting access to financial systems. Overall, I think because the nature and the scale of the challenge posed by AI-related cybersecurity incidents appears significant, and the solution lies in integrating AI-specific risk-management tools within the existing risk-management framework. The focus needs to be on identifying AI solutions that are useful in identifying and supporting cybersecurity efforts at the global level.

Wynnell De Landro-Robinson:

Thank you so much, Manoj. Yes, cyber risk I know is an important part of the priorities, as you say, for supervisors, and I've seen many frameworks adding this either as an inherent risk or part of the operational risk for their frameworks. Now with AI, and we know that supervisors require data to do their work, do you foresee new requirements from the industry, such as additional data and maybe more frequent filings, to help understand the use of AI. Manoj?

Manoj Singh:

In general, with the availability of cutting edge SupTech tools, data filing will be less dependent on any kind of manual intervention, and currently, we are seeing that many jurisdictions are already moving towards data automation, which will greatly reduce the need for manual data filing. The current data strategy of many financial authorities focuses on making better use of data by using advanced analytical tools and centralized data repositories. With the rise of AI tools, I think we may see an increasing trend toward direct access to data sources.



Wynnell De Landro-Robinson:

Thank you. Kangyu, I see you're particularly interested in this area.

Kangyu Wang:

Oh yeah. I am actually particularly interested in one relatively specific aspect of prudential regulation and the data-related challenge of it, which is this: the calculation of capital requirement. So far, for whatever reason, the potential regulation frameworks are on the table and basically all risk management frameworks of financial institutions we have at the moment have been built on frequentist statistics. We talk about, for example, 99.5% confidence level in EU's Solvency II and 97.5% percentage confidence level in Basel III, and the reference to confidence level implies that we are adopting the frequentist approach instead of the Bayesian approach. Otherwise, we'll be talking about things like credibility level, et cetera. Now, we can encounter the uncertainty quantification problem, or the UQ problem, of machine learning. The problem is that when AI-based risk assessment models are used by insurers, reinsurers, and banks to make predictions on the level of risks, their predictions do not come with confidence intervals or anything similar, any kind of conventional measure of uncertainty.

We know that uncertainty exists and that those AI models can err on both sides, but we do not know how much, how uncertain we should be about their predictions and prediction errors. And thus, it can be hard to determine how much capital will need to be reserved in response to the uncertainty of predictions and prediction errors of AI models. Now, there are of course many methods developed to at least partly solve the AI uncertainty quantification problem, but most solutions are adapting the Bayesian framework, which is not the basic framework, the current potential regulation framework is adopting, and the non-Bayesian solutions are either, at this stage, too computationally expensive or having all kinds of questionable reliability problems or practical ability problems. Moreover, statisticians, as well as machine learning researchers, have so many disagreements with charter on this UQ problem and those different UQ measures and also between the frequentist camp and the Bayesian camp. So, it's highly unlikely what the regulators should do: should they somehow convert to Bayesianism or somehow select one or several approved UQ measures that are suitable for financial purposes award. Perhaps nothing at this stage, as the use of AI-based risk assessment methods are still not very wide at the moment, but because of the expected accuracy and efficiency and the granularity of AI-based risk assessment, I suspect that this problem will need to be handled at some point in the future.

Wynnell De Landro-Robinson:

I agree with you, and capital is the key point of prudential regulation for both banks and insurers. For insurers facing climate change risk, the granularity that AI and the methodologies AI can offer is going to help them in understanding this risk better and providing direct capital requirements for it. Thank you, Kangyu. Now we go to regulatory changes. Supervisors get the authority from legislation. What will be the changes to legislation that would help supervisors to be effective in supervising AI as a tool or as a product development tool? Manoj?



Manoj Singh:

Yes, in this respect, the situation would vary from country to country, and we have seen that in many countries, the supervisory powers form part of the overall legislative framework of the central bank of the country. And in many countries, regulatory and supervisory powers are derived from the specific legislation for the financial sector. So, most of the legislative frameworks are tech-neutral and therefore they may be effectively applied to different scenarios arising from the rise of new technology. However, to deal with specific situations like fixing accountability and maintaining a certain level of transparency, non-discrimination, and many other principles, supervisors may need the backing of legislative provisions tailored to the specific requirements of AI. And in this respect, the EU AI act is a good example. As you see, it specifically designates AI credit scoring systems as high-risk for their potential for unfair discrimination. Such AI systems must meet the highest standards set in the EU AI Act, so we may see a similar approach develop in other jurisdictions also.

One thing that can be helpful is alongside legislations, rules, and guidelines is more discretionary power authorized by the lawmakers with which the supervisors can take necessary, appropriate, and flexible actions when needed.

Kangyu Wang
London School
of Economics

Wynnell De Landro-Robinson:

Thank you, Manoj. Kangyu, what would you like to add to that discussion?

Kangyu Wang:

I'd just like to add that when dealing with uncertainty, it's usually a good idea to keep things flexible. Given how much, or actually, how little, we know currently about the potential, the future potential of AI models, perhaps one thing that can be helpful is alongside legislations, rules, and guidelines is more discretionary power authorized by the lawmakers with which the supervisors can take necessary, appropriate, and flexible actions when needed, and you never know when they will need this kind of power without asking for more new approvals when the problem appears. If something goes wrong, and it involves both AI and finance, the situation can develop very quickly and time matters. Now, this of course leads us back to the problem of public justifiability, democratic accountability, and political legitimacy, but those are general problems, which you can see everywhere from pandemic response to military actions. I guess what can be helpful at least is to make politicians and the public understand the gravity and the difficulty of this particular AI for finance problem better.



Wynnell De Landro-Robinson:

I agree with you that legislation cannot be too prescriptive because it would then stifle responses from supervisors. Now, some institutions operate internationally, so what do you recommend in the way of legislative changes for those involved in the regions? Manoj?

Manoj Singh:

The bottom line for new advances in AI is that we are all in this together. So, we can say AI transcends borders, and therefore it requires a global response.

Global cooperation is required for achieving the best possible outcome both at the national level and international levels. Thankfully, we are seeing that international agencies like the OECD are laying the groundwork for widely

accepted AI principles that can form the basis for national legislation. Recently, the OECD published a report on the state of implementation of AI principles in different jurisdictions. As of now, we are seeing different jurisdictions have shown different preferences in choosing the right kind of regulatory architecture for their jurisdiction. It seems that some jurisdictions are taking a cross-sectoral kind of "horizontal" approach to AI regulation, while others are considering a more sectoral, "vertical" approach suited for their jurisdiction. In many jurisdictions, the regulatory authorities are just in wait and watch mode. We may also see so many differences in outlook that there can be challenges in relation to global interoperability, particularly for institutions that operate internationally. But overall, I will say that at this juncture, promoting cooperation at the international level becomes most important for all the stakeholders.

When we take the global point of view, there will be technological diffusion of course, but given the speed of the development of AI, the worry that if a country misses out for several years, it may never be able to catch up.

Kangyu Wang

London School
of Economics

Wynnell De Landro-Robinson:

Thank you, Manoj. Yes, I think at this juncture we have to have cooperation internationally. Kangyu, can you tell us what your points are, what your point of view is?

Kangyu Wang:

Yes, I think there is certainly a worrying asymmetry that we should be talking about when it comes to AI, because AI capacities as well as financial powers are so concentrated in a small number of countries. When we take the global point of view, there will be technological diffusion of course, but given the speed of the development of AI, the worry that if a country misses out for several years, it may never be able to catch up. Thus, a small number of sovereign states and institutions, subject to their jurisdictions, will have the power over everyone else because of this concentration of technology, and the likelihood is that they're going to leverage this power.



This situation could be kind of similar to what we have with nuclear weapons. Meanwhile, as Manoj has mentioned, we are in this together, at least in the sense that AI risks are for everyone, although the benefits may not be shared by everyone, and the distributive impacts of AI will influence everyone. Actually, just a couple of weeks ago, an economics friend of mine had just told me that some ongoing studies have already been trying to understand whether the development of AI in certain developed countries has already caused unemployment in certain developing countries, relying on export-oriented service sectors, and it's hard to tell what should or could be done about this. Of course, the global point of view is important, but when it comes to the practical difficulty, it's hard to say what can be actually carried out.

Wynnell De Landro-Robinson:

Thank you. So, any last words that you may have to caution supervisors when it comes to AI supervision, the use of AI, and the supervisory work involved. Manoj?

Manoj Singh:

Yes, this is very interesting. After more than six decades of research work on AI, we are now just seeing the beginning of the AI age. So, the rise of AI provides further push to the potential of innovative technologies as a transformative force in the financial industry. Even supervisors are not far behind in reaping the benefits. However, the power of AI to drastically transform prudential supervision remains an important question. One of the papers recently published by the Financial Stability Institute mentions that when it comes to the actual deployment of SupTech tools, some evidence points to challenges that prevent deployed tools from being effectively embedded in the supervisory processes. So overall, when it comes to the deployment of AI tools, supervisors need to have a clear understanding of their limitations in the long run. I think a truly transformative technology becomes omnipresent on its own, and no matter how we make our choices, the outcome looks inevitable.

Wynnell De Landro-Robinson:

Thank you. Kangyu?

Kangyu Wang:

I just want to add that although I have been raising all kinds of problems, that's just what I do as a philosopher, an infamously annoying job, raising more questions than we can solve. Really, I'm sort of an AI optimist and the fact that we do not know how to solve a problem now or even what the problem really is, in my view, is not a reason for us to be worried about this problem. We should not be worried too much about something that we do not know, but we should rather be worried about the fact that we still do not know it. Therefore, more research must be done. I think that's true for most parts of life, and I think it's probably also true for what we have been talking about today.



Wynnell De Landro-Robinson:

I would like to thank you both for your points of view, Manoj and Kangyu. It's good to see the practical side and the philosophical side, but we are all in this together, like you say, and we all have to deal with it. The benefits and the detriments of AI will unfold as it becomes utilized more in prudential supervision. But as you have pointed out, there is enough engagement in AI and supervisors to start being aware of the inherent risks associated with its use. Thank you again.

Manoj Singh:

Thank you very much.

Kangyu Wang:

Thank you very much.

Wynnell De Landro-Robinson:

Okay, so I am here today with Manoj Singh and Kangyu Wang, and you have been listening to the Toronto Centre podcast series on Artificial Intelligence and Prudential Supervision. Thank you for joining us today and stay tuned for the next episode.