



The Power of Data, Technology, and Collaboration to Combat Money Laundering Across Institutions and Borders

Speaker:

Beju Shah

Head of Nordics for the Bank for International Settlements' Innovation Hub

Host:

Demet Çanakçı

Senior Program Director, Toronto Centre

Recording Date:

June 21, 2023

Transcript:

Demet Çanakçı:

Hello everyone. I am Demet Çanakçı, a senior program director at Toronto Centre. My guest today is Beju Shah; Beju is Head of Nordics for the Bank for International Settlements' Innovation Hub. He leads a team of experts in cybersecurity payments, artificial intelligence, and central bank digital currencies to deliver trail-blazing projects and insights on technology to the global central bank community and policymakers. Before the BIS, Beju was at the Bank of England for 10 years in several roles, leading large data, digital and business transformation programs, innovation, as well as international initiatives. Beju, thank you for taking the time to talk with us today.

Beju Shah:

It's a pleasure. Thank you for having me and for inviting me to speak about this today. Thank you.

Demet Çanakçı:

First of all, congratulations on successfully completing Project Aurora. Before we dive into the project, I'm curious as to how innovation hub projects are named. They all have really cool names such as Icebreaker, Polaris, Aurora, just to name a few.

Beju Shah:

That's a very good question to start with. How they're named, I think it is down to the creativity of the teams and the center heads. In our case, our first projects, we asked our host central banks to name them. We ran a little competition and then from there, the team selected their favorites. So, it's quite a democratic process, but actually, what came out was some really, really good names and I think they're very fitting for the projects we've just delivered.

Demet Çanakçı:

Very good; it's a good combination of democracy and creativity. Okay, let's dive into the project then. Money laundering is a global problem that undermines the integrity and security of the global financial system. Can you tell us the main objectives of Project Aurora?

**Beju Shah:**

Well, the main objectives of Project Aurora were to really look at this very complex problem of money laundering and treat it as a data problem. So, the first objective is to improve the performance and the effectiveness of AML analysis in a way that still protects privacy and sensitive information. The Financial Action Task Force (FATF), in some of the work they've done around digital transformation for AML efforts that they published over the last couple of years talks about several different technologies that could be used to enhance AML efforts and they include the use of privacy enhancing technologies, advanced analytics including machine learning, federated learning, network analysis and graph data structures, and also the power of collaborative analytics to do that.

So, we wanted to, and after speaking to many different experts and companies and authorities, as we started to get a real grip, a better understanding of the problem, it was clear that these technologies needed to be tested, and so we created a proof of concept to give a platform to see, actually, if could you do this? But also, we wanted to explore the potential of connected payments data and leverage the power of that because today, that's not really used or used nearly as much as it could be, and also, to look at the problem right now, which is where you have a siloed approach to transaction monitoring analysis and the nature of money laundering is actually cross institution and cross-border.

So, if we're not looking at the problem for what it is, you're not able to find solutions to tackle it. So, we wanted to test and compare the performance of these technologies to see, actually, can we, or are we able to detect more than we can do today by testing these technologies out in different modes and actually compare the results to say, can we achieve a better outcome than the way that things are working today which, frankly, is ineffective.

Demet Çanakçı:

Thank you Beju. You touched on a little bit of the process when you are talking about the objectives, but can you elaborate a little bit more on how a project is run?

Beju Shah:

The way we constructed the project was that it wasn't going to be feasible for us at this stage to get real world companies involved for a number of different reasons. The legal process and other factors would be incredibly complex. What we wanted to do is, with a very small proof of concept, test out the technologies I've just mentioned and to really showcase what was possible. So, what we did was we first generated a synthetic data set; this was based off of real-world transactions and represented transaction flows across six countries, 29 financial institutions, and over 150,000 accounts. We embedded complex money laundering patterns within that data, and we built that on the advice of the partners that we were working with and some external experts as well. That meant that we had something realistic that used a minimum set of information that was required and also represented different views of the data. So, you could have an institution, a siloed view of the data, and you could also have a national perspective and a cross border perspective too. That was the first part.

Once we had that data set, we were able to then apply different machine learning models over that data set to analyze it. We took, I think, four or five different algorithms and we analyzed the effectiveness of each of those algorithms against a siloed view of the data, a national view of the data, and then a cross-border view of the data and compared each for their effectiveness, not only the machine learning models in each of those scenarios, but also we could compare how many more money launderers you could detect in each of these more broader views of data and connected data sets.



Then, the third part took the more optimal machine learning models from the second part and then applied them again. But this time, we'd applied privacy enhancing technologies over the dataset. So, these were then encrypted, where sensitive and private information were protected, and then applied the optimal machine learning models over that to see, actually, how they compared as well, both in each of those scenarios. So again, in a centralized view where we simulated, the data could be collected at a national level and in some sort of national transaction monitoring utility. We then did the same thing, and this is more of a thought experiment to say, okay, well what if there was a cross-border international utility? Now, being frank, something like that would be quite difficult to achieve, but actually, from a comparison and a thought experiment point of view, we did that. The next one part of comparison, we then looked at a decentralized approach there we use federated learning. So, in that instance, no data would leave a country or a firm and actually countries and/or firms would work together to collaboratively train a machine learning model, and then train that on the local dataset and then share the insights around. So, that's a very different approach.

And then the last one we used, was a hybrid approach, where data would be, let's say, collected on a national basis in a centralized way, but then internationally you would collaborate in a decentralized way where you then train a machine learning model between countries to do that. And then by comparing those different scenarios, you then were able to say which of those collaborative analysis and learning models, as we call them, worked. So, these collaborative analysis learning models are when data has been shared or you're able to use it in a decentralized way, and also that's where the data's been protected using privacy enhancing technologies. Through that, we were able to get a lot of different comparisons of what is the most effective approach and really lend weight to the fact that, actually, there is a better way of combating money laundering, and actually you can use things like privacy enhancing technologies. If there was more appetite to move towards, even at the very least, some sort of national transaction monitoring utility that employs these technologies, then actually that's still much better than what is going on today.

Demet Çanakçı:

Thank you very much, Beju. This sounds quite comprehensive and also obviously a lot of work. Thank you for sharing. Maybe a follow up question on that, what are your findings and main takeaways from Project Aurora?

Beju Shah:

So, the main takeaways are really looking at national and cross-border systems. So, the broadest data, the more holistic view of the data you have, the more money laundering events that you're going to find, basically. National and cross-border analysis is more effective in detection, but also in the reduction of false positives than the approaches that are taken today. Machine learning and network analysis is more effective in detection and reduction of false positives than rule-based approaches that are used today. Privacy enhancing technologies, in particular, when you combine them with machine learning and network analysis, become a game changer in that because you can still protect that sense of information and enable collaborative analysis and learning to take place. So, this starts to shift the dial. Now we've seen also, particularly in the last few weeks and certainly this week, increased emphasis on the use of privacy enhancing technologies where the UK's ICO have published a very comprehensive report underscoring pretty much the same things that can be used to potentially meet certain data protection requirements, protecting sense of information, and allow this type of learning to take place.



Certainly, we know that the UK and the US are also collaborating on a privacy enhancing technology competition where they're looking to advance innovation in that space, and it's a fast-growing field. So, actually, even some of the technologies that are available today, if you fast forward three or four years, they're going to be really powerful and more mainstream by then because I think there is such a demand and huge potential in that. The last sort of I would say key finding, is that cross-border collaboration with PTs (privacy technologies) in this experiment can find three times more money launderers than conventional approaches. I mean, if you go into the report, there's many other comparisons, but essentially the bottom line is that the broader view, the more holistic view, the more you can connect different data sets, the more you're going to find, and actually with any issues around data protection or data sharing or even international kind of collaboration, PTs, privacy enhancing technologies, can come in and actually can really catalyze some of that change.

Demet Çanakçı:

Thank you Beju. Thank you very much. This is all very interesting, I mean, it is great to see that we can actually leverage technology to the benefit of enhancing cross border payments. So, CBDCs are also one of the initiatives to enhance cross border payments. On the one hand, the potential introduction of CBDC systems offers novel possibilities to strengthen domestic and cross-border payments. On the other hand, CBDCs could provide a new vector for money launderers. Some EMDEs, as we all know, have already launched CBDCs, and some others are working on it. So, how can they benefit from the findings of Project Aurora?

Beju Shah:

Yeah, we touched on that in the report for exactly that reason because, yes, you're quite right; they will invite new vectors for money laundering and things that we haven't seen yet. Or you could potentially see examples from either the crypto space or even instant payment systems that you could probably extrapolate to a CBDC system. What people can learn from Aurora is the fact that with CBDC systems, again, you have more of a holistic view of that data and therefore you can use that for this type of analysis too. Even if you did it at the bare minimum, you didn't have some national utility, but you're just doing it at the CBDC or instant payment system level. You can still use that to find quite a lot. The other aspect is, one thing that we do mention in the report is to say, "Well, we have an opportunity now to get on the front foot."

Now you're not going to be able to predict different typologies because they evolve all the time. But what you can do is go, well, what could you see? What are you likely to see? What can we learn from what we've seen in crypto? What can we learn from and come on the front foot to start thinking about these things now? So, you're not on the back foot. I mean we're, we're always on the back foot because these are very creative. I remember, there's one newsletter I follow and some of the new typologies that emerge sometimes are quite bizarre, and you're like, okay, that's interesting. So, I think it does give you that view of you can use that payments data, you can build in this analysis from the start and, actually, you can get on the front foot if you start thinking about some of the things that you're likely to see.



And I think that's a lot to learn just from there. I think also it is how this also feeds into the broader ecosystem, because it's not just a CBDC system. It could be the existing instant payments networks, it could be other pools of payments data, and, by bringing these together, you'll be able to see the flows that go in and out of each of these systems. So, you may move from, for example, a CBDC to a traditional bank account, into crypto, into something else, and move it around. Unless you're connecting all these dots together, you're not going to be able to see those flow. So, this is just going to be another vector, another system that illicit money may (will likely) flow through as it does anyway, and, actually, we can be on the front foot around some of these things.

Demet Çanakçı:

Great, thank you so much Beju. We heard from you how Project Aurora contributes to both national and international discussions on leveraging connected payments data to combat money laundering across institutions and borders. Now that the first phase is successfully completed, what is next?

Beju Shah:

Well, that's a great question. I think we've achieved what we set out in this first phase, which is to give a platform to the technologies and show what is possible, and I think even some of the results went beyond our expectations too. So, what's next? It's more of things like this where we're speaking to yourselves, we're giving more platform to the work we're discussing, sharing these ideas. Hopefully that leads to either people taking the ideas on and applying them in their other work, or thinking how it can be used for influencing the debate. Or it can also lead to more of an appetite to do some more real-world pilots, which is a call to action we also have in the report as well.

On both sides, I mean it's only been, what, about three weeks or so since we published the report, and even in that time there's been huge interest from around the world. So, we're speaking to many different groups about this, even if it's just a briefing on the work, but also some groups have suggested, "Oh, actually we might think about maybe doing some sort of real-world pilot." So, right now it could go in a number of directions. Also, in the report there's probably more room for the human side of it because it's fine to do the technology, but there has to be the people driving it. I mean, earlier in April, we had a Nordic Baltic symposium to collaboration to tackle economic crime. That addresses the people as part of the public-private partnership that the technology enables, but unless there is that people coming together, it's not going to work. So, we may have more round tables to discuss some other aspects around privacy or data or other aspects, legal aspects, and that will just inform more of a discussion. But also, within that there's a few other values and technical things around applying knowledge graph technology to represent typology information, and that could be quite useful too. So, it could be a few small things and it could potentially be some bigger things, but I think we need to talk more and see how we take this to potentially the next phase. Certainly, the findings from it may influence the other work of the innovation hub or other things that are going on within the BIS too or the central banks as well.

Demet Çanakçı:

Great, thank you Beju. I think this is a good place to conclude the conversation. It has been a fascinating discussion. I learned a lot about Project Aurora. Do you have any final comments?



Beju Shah:

I think the main thing is that we continue this dialogue, and we work out how we collaborate. There are some big questions to be addressed around data protection and other things and legal matters, what public-private partnership means, actually. They need to continue and happen, but also, I think, as I said, more real-world experimentation does need to happen. If people are serious about really combating money laundering, then we have to leverage the data technologies and the opportunities we have today that we didn't have 30 years ago, for example, during the BCCI scandal. If we had the data and technology we have today, that might become more apparent. But I think we need to use what we have today, collaborate, and come together to re-tackle this, and if we're serious about this, then these technologies can really help.

Demet Çanakçı:

I couldn't agree more. Thank you Beju. I encourage participants to read Project Aurora and other initiatives by the BIS Innovation hub. They're all available at the BIS website. I'm here today with Beju Shah and you have been listening to Toronto Center's podcast series. Thank you for joining us today and stay tuned for the next episode.