

AI and You

Transcript

Guest: Alexandra Mousavizadeh, part 2

Episode 150

First Aired: Monday, May 1, 2023

Welcome to episode 150! Today we'll conclude the interview with Alexandra Mousavizadeh talking about the AI Adoption Index, curated by her new company, [Evident Insights](#), where they call it the Evident AI Index, the first independent benchmark of AI maturity in business. They started with doing it for the banking sector, but plan to expand it to other sectors, and we'll hear about the future of that in the interview.

Alexandra started at Moody's as a sovereign analyst covering Russia, Central Asia and the Middle East and then went to Morgan Stanley, later becoming CEO of ARC Ratings, a global emerging markets based ratings agency, and then Director of the Prosperity Index at the Legatum Institute. Most recently, she was a Founding Partner at Tortoise Media, where she ran Tortoise Intelligence, the Index and data business. Here, she was the architect of the ground-breaking Global AI Index which I mentioned earlier, the first benchmark to track the strength of national AI ecosystems. She calling in from London and she holds a degree in economics, mathematics and game theory from the University of Copenhagen. I really like having her on the show again because she's so generous with her time and insights about some very pivotal and very monetizable analysis.

Last week we talked about the reasons for creating and using the AI Adoption Index, and some of the conclusions you could draw from it about the banking sector. Let's pick up the second half of the interview now.

Did your research uncover anything interesting about the flow of talent between particular companies in the sector?

Yes. We take a very deep look at the talent flow between the banks, but also where the banks are hiring from. And probably too much detail to go into here, but we did find that some of the top performing banks, and that is actually both in the US, Canada, and Europe, were pulling talent in less and less from other banks and more and more from Big Tech and also from academia. And actually, there's a growing trend in the banks that are doing particularly well under AI maturity over the last year, seeing that they're pulling in actually quite a high proportion of their new talent is coming from academia, which is interesting because that goes to show that so much of what is being implemented rests on some research capabilities, not just pure research, but also applied research capabilities. Because these are very complicated, sophisticated solutions that are needed for the bank's use cases now, and they're only going to get more sophisticated and moving at a very fast pace. And in order to stay at the front end, it's a very competitive landscape. The banks are being pressured from all sides, from Big Tech that are building payment platforms, from challenger banks that have less of the legacy of the traditional banks and are building some of them sort of from scratch or AI native from the get-go so they don't have to transition; they are that. So the pressures are also coming from asset managers that are

just looking at AI application on the quant side. So the banks are getting pressurized from all of these sides so a real push to find the talent that can come up with the solutions that they need for their use cases really, really quickly. And I think a growing understanding that a lot of that talent comes from academia.

I know of cases in the past where a business's use of a particular technology was considered a proprietary secret. They did not want anyone else to know that they were actually using that because they considered it a competitive advantage and they didn't want other people to realize that they had that. Was it difficult to get any of this information on account of people thinking that that was their competitive advantage?

Well, we took the approach of doing an outside-in assessment. So in other words, what that means is that this Index and this ranking of the banks is based on publicly available data. Now that said, it has been a deep and forensic analysis of a lot of publicly available data, and actually we cover almost 150 indicators that rests on millions of data points for each bank. So while it is an outside-in assessment, there are certain things that you, and you're very correct in saying that some of the proprietary information on specific use cases and specific models that have been used in these use cases we can't see. But the footprint in the public data is there and actually can assess to quite a detailed mapping and I think a very accurate mapping of where the banks are on their maturity on this publicly available data is certainly the case. Now, we do not at this stage use proprietary data sets, but I think with the setting of this Index on the publicly available outside-in approach, it's a really good starting point to perhaps in the future look at maybe we want to allow for banks to upload some of their proprietary data in terms of the sophistication of AI models used in certain use cases simply in order to create a benchmark looking at if companies want to see how are they performing on the proprietary data. And it would be anonymized on the sophistication of their models and more specifically, the outcomes and potentially over time the return on investment. We are considering how to build towards that. But I must emphasize again that the second reason we also felt quite strongly about taking an outside-in approach using publicly available data is because there's a growing pressure from many sides - regulators and customers of the banks - for the banks to be much more visible about what they're doing with their AI activities. Customers really want to know that A, it's being done, that you've got to be on the cutting edge of this. Younger generations are actually at a point where they're trusting an AI system over one where there's a human in the loop. So that's important to then demonstrate that you are on the front end of this and are active in implementing AI where you can. On the other hand, the regulators are also pushing for visibility and reporting on how AI is used and whether it's handled accurately and responsibly. But there's also a growing pressure from customers and other stakeholders at large that you can see that banks are being vocal about how they're handling, you know, especially on the responsible AI side, and putting out statements that are expressing that they are adhering to certain principles on AI ethics and bias and explainability, and so on. So I think that the pressure is on for the banks and companies across other sectors as well, to be much more visible on their AI activities. Investors and shareholders are also putting pressure on companies in terms of being clear about what their AI strategies are because there's a clear link or definitely a very clear perceived link between AI capabilities and future financial performance. And so that link between actually

performance on this Index and future financial performance is something that investors are looking at currently. So the outside-in approach is definitely encouraging more reporting on this. And we can also see already from we launched back in January, but through conversations with the many banks that we have on a daily basis is definitely sort of a change I would say in an understanding that being more visible of these activities is definitely growing. And I would like to say it's the Index that has pushed that. I think it is many factors that are pushing for more visibility, but the Index is perhaps one component.

Well, I was just wondering as to some of the methodology, because when you said that the pillar of talent was measured at quite a granular level, I was envisaging something like company directories or org charts and I didn't know how publicly available those were, or maybe I've got the wrong idea of how you looked at that.

Well, it's interesting. On the talent pillar specifically, we draw data from many third-party data sources. We also draw data from the companies themselves. And so for talent specifically, we do look at the company's job descriptions. We look at LinkedIn, but we'd also look at what research is being put out. We look at the hiring initiatives that the banks are public about. What sort of specific AI hiring initiatives are they visible about? Also, training and retraining. Because that's something that talent, you know, if you're considering working for a bank, you really want to know. Are there skill enhancement arrangements, training, and so on so you can continue to enhance your AI capabilities within the banks? So we look at that and that is something that we would find on the company websites. But we also look at culture. We sort of draw that from sentiment data across other public sources - Glassdoor. We look at many outlets where these things are expressed. We look at LinkedIn where there's also a pool of insights into the entire talent stack of a bank in terms of roles, but also backgrounds and where the talent is coming from. And we can then also start to look at diversity. We can see gender diversity, so the gender breakdown and the AI talent pool which is also important and something we know that the banks are definitely trying to get a more balanced talent pool. But our data will show them exactly where they stand on this. And no, I'm not saying it's easy to achieve. It's something I know that everyone's trying to achieve but this gives a very clear picture of what that looks like, for example. So we do draw from many sources that are not only from the banks themselves but also lots of third-party data sources to piece together the whole talent pillar.

And hinting there at what must be some level of AI just in the analysis of that data certainly when you're doing something like sentiment analysis. How many different metrics go into the computation?

We have just shy of 150 indicators across the entire Index and we draw from its millions of data points that we pull together and we use your natural language processing. We use some machine learning. We do use those good, sophisticated tools to pull this together. And a lot of this is also in real-time, so about 60% of the indicators we track in real-time, which is interesting because it gives us a sense of direction between the two times a year where we release an update of the rankings. So we can sort of take a pulse of the direction of travel in terms of where the companies are going. Very interesting to see, you were asking me about the flow of hiring, and we can see where teams are leaving some banks and going to others, and we can catch that in

real time. And it's an interesting thing to observe where the talent is flowing to and from, and why. And sometimes we actually then can draw that back to maybe a lack of training and retraining initiatives within the bank. And banks, if they get and they hire good talent, got to really focus on keeping that good talent. And that's something we can get quite a sort of granular assessment of what's going on.

Now, you're talking about the effects of the Index, what can happen as a result of people seeing that information. And one of the things that emerged from it is that J.P. Morgan comes out as the leader in these pillars by a long way, a relatively large distance. And what do you think the reactions to that are for readers of the report?

Well, I think it's interesting to see because it really is a reflection of their focus on AI five years ago. Jamie Dimon made a decision to pivot the bank to look more like a tech company and has been very vocal about this from the outset. Talking about then if we're moving and transitioning the bank to an AI-led bank or a tech-led bank, what are the components that we need to put in place to do that? So he did a number of things, one of which was to establish a research hub led by Manuela Veloso, which was established five years ago. And that really was to emulate the tech companies' R&D hub. So Microsoft, IBM, and so on will have R&D hubs that just are in place to really generate research ideas, be innovative, that can be used both within the bank, but also outside the bank. So that was established five years ago. But also took close look at how to reorganize the bank at the leadership level but also across the divisions. And so what we're seeing in terms of the results in the Index is that coming through - a very big hiring initiative. So when you look at the proportion of AI talent to the bank size, but also the volume, granted J.P. Morgan is a very big bank, but also just in proportion to its overall number of employees, very high level of AI talent. The research that is coming out of the bank at very high volume and cited a lot outside the bank. So it's also papers that have had a lot of impact, not only in the bank's operation but also outside. So when we look at the components that go into the AI ecosystem of the bank, J.P. Morgan having started so long ago, have really doubled down on all of those things, and it's coming through in the data. But I think what's interesting to flag also is that Royal Bank of Canada came in second. Now, Royal Bank of Canada is a lot smaller than J.P. Morgan. And since we look at the data in proportion of the size of the bank, what it tells us is that you don't need to necessarily be a very big bank like J.P. Morgan is. You can also be a smaller bank. And they started out at the same time, four, five, six years ago with making that same decision. The leadership decided that they were going to become an AI-led bank over time and RBC also put in place a research lab, reorganized the bank so it was more of a hub and spoke, more nimble in terms of being able to develop and adopt AI across the bank. They doubled down on the innovation side and also did a lot in terms of reporting on the responsible AI side of things. Even though that pillar is not particularly highly weighted in our index, it does matter if you're going above and beyond the sector in terms of your reporting on responsible AI. And RBC is a great case of a smaller bank doing exactly that, making it a strategic objective, doubling down on strengthening the components in the AI ecosystem. And today it's now ranked second in the Index. Now again, we're expanding the Index. We've got banks that are coming in that can potentially rival those positions as number one, as number two. We don't have the data

collected quite yet, but we're very curious to see where the next situation will take us and give us maybe some different players in the top three positions. Let's see.

And it's certainly a surprising result to me. I bank with Royal Bank of Canada, and Canadian banks versus US banks are sort of emblematic of the general differences between Canada and the US: Canadian banks are more highly regulated, less risk-taking, seen as being more "sensible" and less out on the edge. So, not a result I would've predicted. Do you think that AI adoption could be destabilizing for the banking industry? And certainly, within AI there is this effect of the amount of data you have produces disproportionate results in the utility of AI. But is it also a potential for a winner-take-all kind of environment that capital and talent will flow towards the places where it's already dominating?

Yes, absolutely. And it's a bit of a worry because the winner-take-all has only been accelerated with the release of ChatGPT, I think. And I think what we're seeing is that if you are slightly behind at this stage in terms of just getting the rails pinned down for the AI to run on, which means organizing the data, having good content management and data management are the building blocks for AI to run smoothly. So if that is in its early stages and you're in the early stage of thinking about building and strengthening those components that go into your AI ecosystem, whether it be talent or what you're investing in, or whether you are reorganizing your business so it can better adopt AI, it's almost a little too late. I mean, I wouldn't say that it's a little too late because I think there's definitely also an argument that you can sort of skip the landline, and if you get on it now, it should be fine long term. But if you're not on it and it's not on your radar and it's not urgent, both urgent and important for the leadership to figure out how to organize your bank to become more AI-ready, then I do fear that the ones that are really ready to implement and harness the innovation and the breakthroughs that are happening in AI will be able to just move much quicker. And that means that there will be banks with their better financial performance long term. It pulls in the capital from, you know, not just capital, but it's also where the capital flows, but also where the customers go. Customers are impatient. If the customer service is lacking, which as we know, is very AI-driven today, you jump ship. I think customers are very impatient. If your fundamental infrastructure is costly and fragmented and clunky, that is something that your shareholders will have less patience with because it can be done and the AI can create really much better and more efficient systems and so on and so on. So it's sort of an exponential divide between the winners and losers, it feels like. And I'm not sure I would've said that a year ago, but I think it's become much clearer that the breakthroughs, especially on the large language models, which can be applied across many areas in banks, if you're not ready to be able to incorporate that or thinking about how to incorporate that, it's going to be very difficult to stay competitive as a bank. And it's not just large language ones. There's so many other aspects of AI and other kinds of AI that go into banks' competitive edge, but it's one of them and it's certainly top of mind.

Right. And you mentioned ChatGPT there. Has the explosion caused by that, which has caused AI to enter the mainstream conversation, late night talk shows, everything, *South Park*; has that activity over the last four months caused you to pivot your thinking in any way about perhaps

how the Index might evolve or are you seeing effects within the banking industry that are different from what you'd seen before?

What's incredibly interesting about the release was to see the banks' initial reactions to it. Some banned the use of it on the bank's email accounts and so on. There's not a single bank today that's not trying to figure out how to use it, how to use it very quickly, and across how many parts of the bank that it can be used. Now again, everyone's being very cautious. As you mentioned before, it's a highly regulated sector. You've got to do this with caution. If you use it, probably best to test it out on internal systems first and to weed out the issues with hallucinations and so on and put maybe some parameters around it just using internal data. Now, notable was that Morgan Stanley had, actually a year and a half ago, established a partnership with OpenAI and Microsoft and it was very impressive to see how ready they were and ahead of others in terms of three things really - the partnership with OpenAI and Microsoft, extremely good content management, and leadership that had doubled down on large language model capabilities. So they were really ready to use that. Now, it's still in testing mode. They're using it just on their internal data. They talk openly about how there's a constant feedback loop in terms of ironing out the kinks of it to really then sort of unleash it across the bank. But right now it's being used by financial analysts to be able to give much more informed and more quickly, advice to their clients using ChatGPT on the internal data that they have. But they also use external data to pad that analysis with, and it's really interesting to see what comes out of it. They're not the only bank. Goldman also are working on it. All the banks are working on it, but just taking slightly different approaches. So I would say, has it given some pause in terms of would we have measured things differently? I don't think so. What it has done and as we've begun back to look at, it'd be interesting to see if we look back at the data specifically in terms of talent that is large language model focused and research and patents that are large language model focused. And just have a look back and see which of the banks had sort of doubled down on that, because I would imagine that these are the banks that are going to be best positioned to harness these ChatGPT-style capabilities going forward.

Well, as we get towards the end here, I want to just comment on the image on the report and on your web page of the starlings because it spoke to me. A picture of an enormous flock of starlings, which the collective noun for is a murmuration; I don't know why. But if you've ever seen this or seen a video of starlings, they move in this amazingly organic way of thousands of birds. And the reason it spoke to me is that individually, each of these birds is doing something that has some relatively simple rules that's just reacting to the ones around it. But when you look at the whole flock, they move in this amazing large structure that constantly changes. And to me, that seemed to be reflecting what you were doing, in that it is revealing something that, at scale, becomes apparent and emergent that can't be seen in the individual level. And I don't know if that's what you were going for or not, but that seemed to be what I took away from that. Does that speak to you in some way?

I'm so glad that you picked up on that because that was exactly what we had in mind with the starling formation. And it does two things. It is to show that in all of this, what looks like a sea of data points, and it demonstrates our approach to building indices, that's sort of one aspect of it,

that in this sea of information and sea of data points, there emerges a structure that has certain contours around it. And actually, what emerges from that is a certain shape and actually a story. And the way that we talk about this internally is that you sort of have this mosaic of data and then you piece it all together and there's an image that appears from this data. And that's how we like to think about our work, that by spending all of this time in the engine room and forensically pulling all of this together that might look sort of a bit disparate when you sort of first hear, "Oh, we're pulling this data and that data and we're looking at all of these things." But when you actually end up having it all gathered, it makes complete sense. And from our conversations with the companies that we have on the Index, many of them have said, "This is interesting because all of these pieces of data points have come together and really depicted us very, very accurately in terms of our AI capabilities and our AI story." So it's nice to be able to make that loop back to the starlings that do exactly that. But it's a general AI story about how with the growing data sets and the growing compute power and the models that are emerging, that so much of this data around the world that seems so fragmented, that not only for the companies but sort of if you take it up to a more meta level that we can use AI and use the data collection to pull together stories and information and insights that we just really couldn't get to before. And with that, sort of believers of that this can have real positive impact on the world but of course, there need to be guardrails.

Right. What's next for the Index? Do you see yourself using more AI like large language models to pull out the patterns from the data for you?

Yeah, we already do that. We do a lot of correlation analysis, but it's interesting to use the large language model capabilities to see is there anything that might emerge from what they pull out that we might have missed or some correlations that are not particularly intuitive. So that's what we're using. So yes, we are using all of the innovations on the large language model platforms that are available. And what's next for the Index is that we will be expanding the number of banks on this current index to go from 23 to 60 banks, which we are going to be releasing in October. And then we're going to start looking at what sector is next and would invite anyone who has good ideas of which sector would benefit from being mapped in the way that we've done, we'd welcome any comments and suggestions for sure.

And would that interest principally come from investors in said business sector?

It would be the sector themselves. So the insurance sector has been in touch and asked if we could do them next, but we've also been approached by the manufacturing sector and the retail sector, but also the health sector, or more specifically pharma where there's a lot of AI that goes into drug discovery, but also across more broadly the health sector. So it's just a question of which sector we pick next. But by and large, over the next four to five years, we hope to cover at least seven to 10 sectors. So we'll get there eventually.

Wow. And I see somewhere in the future perhaps this has some intersection with what you mentioned at the very beginning about the possibility of us needing an equivalent to the International Atomic Energy Agency in the future for our AI governance and responsibility. So,

Alexandra, it's been fascinating talking with you. Where should people go to find out more about you, Evident Insights, and the AI Adoption Index?

Yes, please do reach out to me directly on my email, which is alexandra@evidentinsights.com. If you have any questions, please don't hesitate to contact me or go to our website, which is evidentinsights.com, and you can see all the things that we are up to and planning. We are also hosting a big forum in New York in October. We've got some interesting debates and events taking place in London in June where we talk about what's next and the impact of the large language models. We look at responsible AI. We're taking a deep dive into the talent pillar to look at what are the leading companies doing in terms of attracting talent and being able to maintain it when you're thinking about building an AI-first bank. So we're really trying to cover all of the pillars and update as we go throughout the year. So please do get in touch if you've got any questions.

Oh, thank you so much, Alexandra Mousavizadeh, for coming on AI and You.

Thank you, Peter, for having me. Such a pleasure.

That's the end of the interview. I was very interested in the possibility that the use of AI like ChatGPT could be destabilizing for the banking industry for reasons that could equally apply to any other industry. It certainly explains why there's such an incredible number of startups being created every day to bring some new ChatGPT-based service to market.

Episode 150 – it's a bit of a round number, not round enough to occasion a fanfare, but it's still an excuse to remind you that we get more listeners to the show and hence get to grow the show in ways that matter to you, when you share about the show, and especially give us five stars and a review. It's free, it takes a moment, and it makes a huge difference.

In today's news ripped from the headlines about AI, OpenAI has launched its own tool for detecting whether text was written by AI. They call it the AI Classifier, and you paste some text into it and it gives you an assessment. It's in a public beta mode – like every other piece of software these days – and it is, understandably, not very reliable, especially on less than 1,000 characters of input. Text written by AI isn't given some kind of foolproof watermark; this is text, after all, you can cut and paste it, there's just no place to put anything foolproof. So the AI Classifier and tools like it work by analyzing the word patterns in the text to see whether they look like the kind of thing AI comes up with, which usually means that it looks unnaturally typical for human writing. Humans have variation in the complexity of their writing but AI is more uniform. That doesn't sound like much to go on, and it isn't. I pasted three texts written by GPT-4 into it and it said that two of them were possibly AI-generated and one was very unlikely to be AI-generated. I pasted some of my writing into it and it also said it was very unlikely to be AI-generated. Then I put part of a document written by an engineer into it and it said it was unclear whether it was AI-generated. The same happened for part of a Wikipedia article. So this is not likely to become the tool that rescues the term paper as an instrument of assessment.

Next week my guest will be Eric Daimler, founder and CEO of Conexus, and a former policymaker in the Obama administration serving as a Presidential Innovation Fellow for AI and Robotics in the Executive Office of the President. We'll be talking about policy and safety issues with the commercialization of AI

and the role of government in all that. A thoughtful conversation with a fascinating speaker, next week, on *AI and You*.

Until then, remember: no matter how much computers learn how to do, it's how we come together as *humans* that matters.

<http://aiandyou.net>