

**iLLUminate Blog Transcript: Don Bowen on AI's Impact on FinTech**

Recorded March 22, 2024. Listen to it [here](#).

- ANNOUNCER: 00:01 [music] This podcast is brought to you by iLLUminate, the Lehigh Business blog. To learn more, please visit us at [business.lehigh.edu/news](https://business.lehigh.edu/news).
- JACK CROFT: 00:13 Welcome. I'm Jack Croft, host of the iLLUminate podcast for Lehigh University's College of Business. Today is March 22nd, 2024, and we're talking with Don Bowen about the impact that artificial intelligence, or AI, has already had on the fintech industry and some of the ways that AI offers both future potential benefits and threats to fintech. Dr. Bowen is an assistant professor in the Perella Department of Finance who teaches Lehigh Business's Capstone course in FinTech, among other courses. He also is co-director of Lehigh's FinTech Minor. His research focuses on corporate finance, with particular focus on the interplay between corporate investment, innovation, patent markets, venture capital, and IPOs. Thanks for joining us on iLLUminate today, Don.
- DON BOWEN: 01:07 Thanks for having me. Happy to be here.
- CROFT: 01:10 So as, yeah, I mentioned in the introduction, the impact that AI has already had on the fintech industry, let's start there. What kind of impact has it already had that you've seen?
- BOWEN: 01:27 Tremendous. It's been pretty profound already, and it's only going to accelerate from here. AI, I think of it as having two main components. One is a set of tools that make it easier for firms to do predictive analytics. So they're trying to understand some outcome variable, like will this customer pay back their loan? So that's one facet, predictive analytics. The other facet is generative AI. So these large language models, such as ChatGPT, and integrating these language models into how the firms do their many different tasks across the firm. On the latter - on using generative AI tools - it's going to change not just the way employees at firms do their job, but it's going to change the way customers interact with firms in tremendous ways.
- BOWEN: 02:23 The first one thing that's already underway is the broad adoption of chatbots for customer service. So Klarna is a loan provider, and they recently disclosed that for about a month, they've been using an AI assistant for customer service chats. And in one month, it handled two-thirds of their customer service chats. Basically, it was about 700 employees' worth of work. And humans, the customers, gave it nearly equivalent to satisfaction scores. It did a great job solving customer inquiries. A 25% drop in repeat inquiries. And it was solving customer problems much faster. So I think the number they gave is it was about 80% faster at fixing customer problems. Because it's a computer, it's on 24-7 around the world, in a lot of languages. And Klarna thought, OK, well, this one use of AI - one use - is going to help them increase their profit by \$40 million because obviously they're going to reduce their customer service organization. So the wage bill will fall. But also, you can bake in some value from the quicker service and the happier service that they get. And that's just one way that these GenAIs are going to be useful.

- CROFT: 03:59 Now, in terms of AI - generative AI - taking over some of the customer service responsibilities, there's been a lot of discussion about AI as certainly a first step into moving into a lot of companies, taking over what some of the redundant functions that people have been doing for decades. And--
- BOWEN: 04:27 Repetitive.
- CROFT: 04:28 Repetitive. Yes. Yes. And the question has been, of course, that that's probably one of the first workforces that's going to get hit. So if you could talk a little about that balance between a more efficient, faster function with AI and the impact it has on the human workforce.
- BOWEN: 04:59 The first thing that I would think about in considering how AI is going to impact workers and people more broadly through their other relations with the firm production, it's going to make firms more productive. They're going to produce more goods at a lower cost, and consumers are going to benefit from that. This is true if I'm talking about AI now, but it's also true if I'm talking about automated assembly lines or talking about using cars instead of horses, right? So new technologies invented this changes the way things are made, and the way things are done. And the first and largest impact is this makes everybody better off. We are able to make more things with less. So the pie is a lot bigger now. There are some questions about-- there's always some questions about who is affected by that, and how long are they affected, and how deep is the effect.
- BOWEN: 06:06 So this changes from generation to generation as we have new technology that comes online maybe has slightly different fundamental attributes than prior ones. But it also depends on the macro economy. Right now, if you get displaced by generative AI, you're going to be on the job market at one of the best times in history to find a job. Our unemployment rate is historically low. And job seekers are having an easier time than ever at finding new employment. So that at least makes it likely that the impacts are going to be less severe than they might otherwise be. But it's a very interesting topic to think about. I mean, every time there's a new technology that can do the job of some humans, this is a question that's going to come up going back as far as the Luddites and the cotton textile work that they were doing.
- CROFT: 07:10 All right. Now let's talk then about some of the ways that you see so far that AI holds the promise to improve the fintech industry. What kind of things that it may not be capable of doing today, but the trends are that in a short time it will be able to do?
- BOWEN: 07:35 Yeah. Yeah. So just a few promising avenues that AI could be used to influence the industry. One, enhancing customer experiences, which we just talked about one example of. Two is refining risk management processes at firms. And the third one is hopefully improving financial inclusion. One of the first things I said about AI is that one of the things that AI means to me is predictive analytics to try to understand some outcome, like is a borrower going to pay back their loan. Financial firms have used AI techniques for quite a while to improve their risk management processes. So they have a lot of data. It helps them understand which customers want which products and which services. And it helps them improve their pricing models.
- BOWEN: 08:35 And the nice thing about that is that, in a sense, two things happen. One is that more people are offered credit. So some people who previously weren't offered credit are

going to be offered credit. It might be at a higher rate, but their access to credit will expand. And that can make the lending space more competitive, which is better for customers. And the access to finance, we know from a long history of research, the access to finance is very, very important for the dynamism of an economy and helping new firms get founded and grow, helping people manage their personal liquidity so that they can concentrate on other things that matter to them, right? So if you're worried about making your next bill and you're having trouble raising funds, well, this burden is going to make it harder for you to concentrate on your family and your children. Finance has always been a backbone of the economy, and things that improve it are generally going to have tremendously positive first-order effects.

BOWEN: 09:48

In the finance space, we get risk management techniques that are helping firms assess loan applicants. And also, fraud detection was very, very-- in the first wave of AI implementation in the finance industry, fraud protection expanded a lot. My first experience with this was back in 2008 when Wells Fargo told me somebody was trying to use my credit card in Lafayette. I had never been to Lafayette. And they saved me that day \$300. And I was an undergraduate student at the time, and that felt like saving \$1 million. [laughter] Nowadays, these fraud detection mechanisms are much more sophisticated. When you leave the country, you don't need to tell your bank that you're leaving the country. You used to have to tell them, and now you don't. And when you go abroad, your credit card's going to work. But if somebody in a different country were to use my credit card right now, it's more than likely that my bank would flag it for me with relatively fast speed and protect me. So that's the risk management.

BOWEN: 10:55

I guess the last thing that I mentioned at the top was the idea that AI can improve financial inclusion. So one thing I said earlier was that by allowing firms to do a better job of understanding the risks for borrowers, that makes it easier for them to open access to more borrowers because they can price their loans, that credit risk, appropriately. The other thing that's really exciting is the expansion of robo-advisors and hybrid advisors. A lot more Americans than in the past have access to retirement funds and savings. But making choices for your retirement accounts is a very complicated thing for most people. And the people who study this find that a lot of people leave a lot of money on the table by making some relatively simple, easy-to-correct errors, to be blunt, with how they set up their retirement accounts.

BOWEN: 12:02

So there's been a large growth in the use of robo-advisors, or nowadays, hybrid advising systems. So it's an AI system that's suggesting that, "Hey, Jack Croft, he should probably consider reallocating his money from one investment type to another." But what happens at several fintech firms or firms that are adopting this technology is that actual advice is sent to a human advisor to tell the customer. So there's a human in the loop. So the generative AI and AI system that they have kind of analyzing your retirement account and the way you're set up will generate some insights, some suggestions. Those will be sent to a human who can review them quicker because the human doesn't have to come up with all the ideas from scratch, but simply sits there as a reviewer and then tells that to customers, which customers love. People like advice from LLMs, but they like it less when they know it's from a robot. They like advice from robots more than advice from humans. But once you reveal that the advice is from a robot, people get a little bit hesitant. So what finance

firms are doing these days is changing the way it works to often have a human in the loop.

BOWEN: 13:22

So for example, at Vanguard, if you're in one of these systems, the advisors who are doing this kind of hybrid advising, they don't service the same number of customers they used to service. They service a lot more because the AI system makes it easier to do their job. So if they want to work with one customer, it's a lot quicker. And so they work with more customers, which again expands access to that advisory service that did not used to be accessible before. There used to be a very high net worth minimum to get into those programs. And now it's lower because they can do their job more cheaply. So that expands access. And not only did that expand access for customers, that actually expanded the amount of employees that the firm needed. Because now, instead of only servicing some elite customers above some high net worth, they're now able to service customers at a much lower net worth. And there's many more of those customers, which then requires more advisors. So this kind of gets back to what we were talking about earlier, the way AI people are worried that it might displace workers. But I'm not personally so worried about that. It's going to change what people do. The skills you use might be different. The job tasks you have might be different. More people than ever will work with GenAI in their day-to-day work. In fact, the adoption of GenAI in the economy has been faster than just about any recent innovation. It's going to change how people do their work, but it's not going to put people out of work. It's just going to change what they do. That'd be my prediction.

CROFT: 15:08

Now, on the flip side, any new technology obviously also comes with the potential to disrupt industries in harmful ways. So what potential do you see, particularly in the fintech industry, for that to be the case?

BOWEN: 15:31

Thinking about the risks and challenges has got a lot of attention, not just from media thinking about what this means for employees across the economy, but also from regulators. So one thing that a lot of these AI systems do, well, all these AI systems, they learn from historical data. So one and very natural thing to think about once you start from that premise is that if the data that they're learning from has some biases inside of it, that might be present in the AI systems. And most of - well, all of the - companies building these foundational AI models, the leading models that you've heard of, have done a lot of work to reduce bias in the back end.

BOWEN: 16:24

But it's work that requires a lot of care. And frankly, these systems are so complicated, nobody really understands exactly how they work. There's a lot of very interesting examples of unintended consequences from changing small instructions that you give these robots. So if there is discrimination in the data you're starting with, it could lead to discriminatory lending practices, or insurance premiums, or investment advice. It's possible. So what this means is that firms that are using these models in production and regulators that are looking at these firms are going to spend a lot of time on the back end to ensure fairness and to do their best to increase transparency in the process.

BOWEN: 17:18

One of the simplest things that finance firms do when they're using more sophisticated models is be very, very stringent about the data that you give it. So if you're asking a model that's designed to help a bank make a lending decision, if you give it information about where the lender is in America, that's going to be heavily correlated with race. And so a lot of lenders try when they give the data about the

prospective applicant, when they give that information to the AI systems, they'll purge any information that they think is, first of all, explicitly forbidden by law, but they'll also exclude information that's highly correlated with those characteristics, all right? So they're trying to both prevent explicit discrimination on things that are protected, like race and age, but they're also trying to prevent inadvertent discrimination due to correlations.

CROFT: 18:31

Yeah. And I think that brings us to moving kind of from the general industry to the specific, the FinTech Capstone course that you teach at Lehigh Business. And last year, I think there was a very interesting example that illustrates what you were just talking about, where you incorporated generative AI through ChatGPT in a project that was looking at fair lending practices with the students. And so if you could just kind of walk us through what your students were looking for and what they found. Because, as I recall from talking with you about this before, there were several layers, and the presence of some bias in the data didn't show up until you got several layers down in this project.

BOWEN: 19:28

The project was relatively simple, but maybe at least I thought pretty ingenious in the way it was set up. So what they did is they took real loan data, where people applied for a loan and they were approved or denied, where the students were able to also know what the interest rate was that the lender gave for that loan. And they basically gave a bunch of information about that loan applicant to ChatGPT. And they said, "Hey, should this loan be approved? And if so, at what interest rate?" Then they collected the information. and they did this for a bunch of loans.

BOWEN: 20:10

And then they went back, and they did the same thing again, but they changed one thing. They told the robot that the borrower was white. And when they did that, the denial rates did not change and nor did the interest rates that were suggested. So the first time they did it, they didn't say anything about the race of the borrower. The second time they did it, they said that the borrower was white, and there was no change. The third thing that they did is that they sent all the loan information back to the robot, asked it for loan decisions and interest rates, but they said that the borrower was Black. And when they did that, the denial rates didn't change very much. The sample was small, so they weren't able to tell if it was statistically different. But what did change is that when they told the robot that the borrower was Black, the suggested interest rate went up. And it went up for Hispanic borrowers as well, well, when they told the robot that the borrower was Hispanic. So this suggested that there was some baked-in bias in the model. And this was ChatGPT-3.5. So this isn't the most current leading model that they offer, but this is the original one that kind of broke through and was popular. So that was the first set of tests. So it looked like if you revealed to the robot the race of the borrower, it changed its suggestions.

BOWEN: 21:47

The next thing that they did is that they just repeated this entire experiment all over again, but they changed one instruction. Instead of saying, "Should this loan be approved and at what rate?" They said, "Should this loan be approved and at what rate? You should not allow any bias to affect your decisions." They didn't say anything about the fair lending laws, or they didn't say race is a protected class. They didn't say age is a protected class. All they said is, "You should make this decision without any bias." Which it's kind of incredible because it's a very vague statement. It's not specific. But when they did that and they repeated the experiment - white borrower -

when you send that instruction, you send the loans, it says the denial rates and the interest rates. You do it again, but you tell it that all the borrowers are white and there was no difference, just as before. But what's interesting what happens when they told the robot that the borrower was Black the interest rates that were suggested were the same as the interest rates for white borrowers.

BOWEN: 22:54

So the upshot here is that ChatGPT-3.5, at least in their study, showed some bias in how it assessed the risk of the borrowers, but it was also relatively easy to mitigate. So this is something that might be relevant for regulators, even thinking about how lenders and financial firms should go about using these tools in their processes. So they were simply able to tell the robot, "Hey, don't have any bias." A more sophisticated instruction might be to explain a little bit about how fair lending laws are in the U.S. and what they are allowed and not allowed to do. But that might not even be necessary.

CROFT: 23:42

And then, in terms of other takeaways from the results regarding the use of AI in fintech is there anything else that kind of bubbled up from that project that you think might have some broader implications for the fintech industry?

BOWEN: 24:03

Yeah. Well, first of all, you have to assess a firm's need to be diligent to collect enough information so that when they use their AI systems, after the fact, they can assess whether it looks like they have some bias that's emerging, whether it looks like it's starting to make different decisions across some protected classes. Again, race, age, for example, gender. And if you find any instance where it looks like that's occurring, there are some remedial steps that you can take. And the nice part about the remedial step that the students took, just to tell the robot to make the decisions without bias, is that instruction is at least very transparent.

BOWEN: 24:47

If you're making loan decisions with ChatGPT, which I doubt a bank would do at this point, but if you were-- maybe a fintech that's young and brash might do it. If you were using generative AI to suggest what an interest rate should be, you can't turn around and tell a regulator how it's making that decision. And regulators want to understand how your lending process works. But if you're going to use a generative AI, you're not going to be able to really tell the regulator what's happening behind the scenes. All you can do is toggle under, "Hey, we give the robot these variables and these instructions." The variables that you give it are transparent, and the instructions you give it are transparent. So the students viewed the exercise as mostly about understanding whether or not there were some ingrained biases in the data and the way the model was written. And at least for ChatGPT-3.5, it looked like they found it. But it also seems very easy to correct. And that's relevant both for firms that use it, and regulators when they evaluate how firms are using these models.

CROFT: 26:07

There's a couple of recent studies you've done - research - that I'd love to talk about, but we're out of time. So I think we'll--

BOWEN: 26:20

That's all right. We have some meat on the bone.

CROFT: 26:22

Yep. And I think this has been a good kind of overview of some of the issues involved. And we'll be happy to have you back to talk about those studies as well. And I guess for closing, I'll leave it with, is there anything else that we haven't talked about regarding fintech and AI that you think our listeners should particularly know?



BOWEN: 26:50

These AI technologies, they're going to change dramatically how financial firms do their jobs, and how they interact with customers, and how customers interact with them. It looks very much like we are on a path towards a lot of financial tools being more accessible to people. The most exciting thing is increasing the access to financial advice and advising. This is the kind of thing that has tremendous impact as people get older and they get closer to retirement. And so hopefully, this will have a positive impact on many generations to come. It's going to change the way people do work. Their tasks will change, their jobs will change.

BOWEN: 27:40

But it looks like a lot of these generative AI tools, especially in the finance space, they work in a way that expands the needs and the productivity of employees. So it seems like in a lot of settings, while this might reduce employees in some use cases, in other use cases, it increases the number of employees that you need. But regardless of the direct and short-term impact on how labor is allocated economy, it does look like these technologies are going to make for a world where we get a lot more with a lot less. And that's pretty exciting.

CROFT: 28:23

Well, Don, thank you so much for being with us again today.

BOWEN: 28:26

Thank you so much, and I'm very glad to be here.

CROFT: 28:29

Don Bowen's research has been published in Management Science and presented broadly, including the Western Finance Association, NBER, UBC Winter Finance, Finance Research Association, and the Midwestern Finance Association. This podcast is brought to you by iLLUminate, the Lehigh Business blog. To hear more podcasts featuring Lehigh Business thought leaders, please visit us at [business.lehigh.edu/news](https://business.lehigh.edu/news). You'll also find links there to follow us on your favorite social media platforms. This is Jack Croft, [music] host of the iLLUminate podcast. Thanks for listening.