

The Impact of AI on Macroeconomy, Growth, and Financial Stability¹

Jason Furman

Tobias Adrian

Abstract: Mr. Tobias Adrian believes that AI has brought positive changes to the financial sector—algorithm-based financial transactions have increased market efficiency and reduced financing costs across sectors; using AI to improve credit allocation enhances financial inclusiveness. Theoretically, with the emergence of general AI in the future, the financial sector may also achieve autonomous decision-making by machines.

However, from the perspective of financial stability, AI increases the difficulty of risk management at the corporate level, and can also expand the scale of attacks (against financial institutions), increasing the efficiency of attacks, with cyber risks becoming a potential macro critical risk. Meanwhile, although generative AI may be proficient in understanding cross-sectional risks, it is unclear whether it can capture general equilibrium effects and how this would impact financial cycles.

On macro stability, Prof. Jason Furman believes that in the short term, the application of AI will not increase output but will only increase labor input, thus overall productivity is declining. This also means that managing inflation has not become easier in the short term, and may in fact increase, requiring higher interest rates to stabilize the macro-economy. However, in the long run, the development of AI allows for significant growth in productivity, which will translate into increased income and higher interest rates.

Regarding the commonly concerned issue of labor replacement, both Adrian and Furman are optimistic because historical experience shows that with technological progress, new types of jobs will emerge, and people's overall income levels will rise, leading to more service industry positions. Moreover, technology can only replace some jobs, not all. However, Furman also emphasizes the risk of increasing inequality, such as people lowering their wages to compete with robots. He believes that wide-ranging measures must be taken to address this challenge, one focus being the investment of more educational resources.

Regarding the regulatory challenges of artificial intelligence, Adrian believes that a major challenge currently facing financial regulators is that the data needed to understand these risks are not necessarily the same as those needed in the past, and where and what kind of data to collect may differ from before. Therefore, increasing transparency may be the top priority. Meanwhile, regulators need to try to use AI to provide information for regulatory visibility and system stability judgment.

Furman believes that waiting for regulatory measures to catch up with the development of AI is not a wise

¹ Jason Furman is professor at Harvard University and former Chairman of the White House Council of Economic Advisor. Tobias Adrian is Financial Counselor and Director of Monetary and Capital Market Department of IMF.

* This article is a transcript of a dialogue between Tobias Adrian and Jason Furman at the 6th Bund Summit on September 6, 2024. The dialogue was moderated by Shen Yan, CF40 guest speaker, Professor of the National School of Development at Peking University, and Deputy Director at Institute of Digital Finance of Peking University.

move. Slowing down the deployment of AI also carries many risks. When considering AI regulation, the first priority is to balance benefits and risks. Second, attention should be paid to AI biases and not to view it as omnipotent; third, regulation should not become a moat to protect existing businesses. Finally, he points out that many solutions to AI issues are not related to regulating AI but involve independent strategies concerning the labor market, tax systems, etc.

Shen Yan (Moderator): So today we are going to talk about the impact of AI, as Charles Dickens' famous quote says, it's the best of the time, and it's the worst of the time, probably that is very appropriate when we think about the impact of AI. For enthusiasts, AI signifies the beginning of the fourth industrial revolution, and for those who are pessimistic, we may think that AI may present a fundamental risk to human civilization.

So here today, we have two distinguished guests to discuss these issues. Okay, so let's open the discussion with the first question, to clarify the concept. When we talk about AI, what do we exactly mean?

Tobias Adrian: Let me give you two sets of answers. So the first one is, what do you do when you work with AI, and you know, I don't know about the audience, but when I learned how to interact with computers, I learned to code MSDOS and C++ and, you know, it's quite, quite difficult. It's like learning a language, you know, the next generation, like, what my daughter learned was symbolic coding, right? So you don't learn a language, but you still have to, sort of, like, understand how it all works.

AI nowadays is based on large language models, and you interact by talking to the machine, right? So you still have to sort of figure out how it all works. But it's not by learning a language or learning symbols. It's by talking to the machine. So that, I think, is one way in which AI is very different from the way in which you know previous computing worked.

I think the second observation I would make is that, you know, generative AI and large language models are based on calibrated models that are using billions, and

nowadays hundreds of billions of parameters. So it's extremely complex, beyond anything that was available, you know, 10 years or 15 years ago or so and. So it's really a game changer in terms of how computers are interacting with reality. The order of magnitude of complexity that is captured by the models is totally different. So it is revolutionary.

Shen Yan: Okay, revolutionary from two perspectives. One is a large language model that never exists before, and second is the number of parameters and the way it works.

Jason Furman: Yeah, so there's an old joke attributed to the computer scientist Larry Tesler, that AI is whatever hasn't been done yet, the rest of it is engineering. And that gets at people said, oh, you know, only humans can play chess. What if you had some artificial intelligence, and this artificial intelligence could play chess? Well, the problem of chess was solved reasonably well a while ago, extremely well, a couple of years ago, and now we almost don't think of it as AI anymore.

Speech recognition, the same thing, you don't say "powered by AI". That's just engineering. That's something that's been done. So the AI is whatever hasn't been done yet.

Now that's changing a bit now, partly for marketing reasons. If you call anything AI, your stock price will be higher. Your customers will buy more of it.

And all of that just says there's a little bit of ambiguity about exactly what it means. The version we're using now so much and that device was talking about is generative AI, which is basically looking at things

that humans have done and remixing them in some way that maybe answers a novel question, or maybe answers a question that was asked hundreds of times before,

In a lot of ways, that does seem to mimic what a human does. And you know, there's the old Turing test. You can have a conversation with ChatGPT, Claude, or any of these things, and they're enormously human like, but it's still not close to, you know, what's called general AI, where it basically would be a human in lots and lots of different dimensions, including creativity.

Shen Yan: Okay, great, so you two already gave different definitions. Okay, so given that we clarify the concept. So the next question will be, we believe that AI can bring us brighter futures, then why and how? How it can boost the economic growth and then the finance industry?

Jason Furman: Yeah. Well, there are, broadly speaking, two types of ways.

One is taking things that humans already do and doing those same exact things, but doing them potentially faster, potentially better, and certainly doing them more inexpensively. So you look at, you know, a certain amount of marketing and design, you don't need the person to draw the thing anymore. The AI can do it now. But the other is doing new things that we couldn't have imagined, we couldn't have thought of doing ourselves, and that aren't replacing a human, but complementing and working with them.

So far, it looks like there's an awful lot more in that space. I mean, there's a reason that Microsoft calls its AI agent Copilot, because it's about working together with humans, sort of like spell check as you're writing corrects your spelling.

But we'll see both of these types of things, replacing humans, but also augmenting them. And as we come back and hopefully we'll be able to discuss the labor market, a big issue is how quickly this happens and how much of these different types happen.

Shen Yan: Okay, so two aspects, one is increasing efficiency, and the other is expanding the role that humans can have their activities, and then maybe Dr. Adrian can give us more details in the finance industry.

Tobias Adrian: Yeah, so let me zoom in on the finance industry in three steps.

So first of all, we have seen a revolution in trading already, which is not based on generative AI, but on machine learning over the past decade or a little bit more. You know, trading has really changed fundamentally, particularly in the most liquid markets, by becoming very much algorithm based, and arguably that is improving market efficiency and is benefiting households and corporations and governments in terms of lowering funding costs and making markets more efficient.

The second example I would give is credit scoring, right? And I think you've done some research in this area, and there's other academic research certainly that has looked at how artificial intelligence can improve credit allocation, and that can ultimately benefit the population at large via increased financial inclusion and better allocation of credit across agents, across people and across firms.

And thirdly, and that is really the frontier, what Jason was referring to as the next step of artificial intelligence. So general intelligence, I think, in the financial sector, this is the question about fully autonomous decision making in financial markets or financial institutions. There's a great degree of skepticism at the moment as to whether we will get there, and there's certainly we have not talked to any market participant at the moment that is willing to let the machine run without human intervention. But there's this sort of like possibility, theoretically, that there could be fully autonomous agents, and presumably that would be achieved once there would be, you know, a more general form of artificial intelligence.

Whether we will ever get there, we don't know, but conceptually, it's possible.

Shen Yan: Okay, given that you've said three parts, algorithmic trading, and then using AI for more of financial inclusion, and then also automatic decision, even though we don't know how far, but even for what you said, If we flip the coin, there's always another side, and that I want to bring up right now about the stability, when we are using more of our algorithmic trading, will it make the financial market more volatile? That's one thing.

And the other thing is the automatic decision, a recent talk by Peter Schmidt even mentioned that for large language models, there are risks that a human may not even be aware of. So what would AI mean to the stability of finance, and then to Dr. Furman, to the whole macro economic development?

Tobias Adrian: Yeah. So let me start with the firm level, and then go to the economy wide level, and then, so like the financial cycle.

So at the firm level, when you think about the risk manager in a financial institution, right? So risk managers actually spend a lot of time trying to figure out what people are doing. So, for example, a risk manager that is overseeing an options trading book, you know, has to sort of like, understand what people are doing, so like, theoretically, that should be pretty straightforward, but in practice, it's actually very hard, because these are very, very complex, multi dimensional trading positions. Once you allow traders to use artificial intelligence, the problem for the risk manager in terms of understanding what kind of risks are being taken can be that much more complex. So, you know, on the one hand, every financial firm is hiring a great amount of expertise on artificial intelligence, but on the other side, you know, managing the additional dimensionality of risk taking, I think this is sort of like the number one issue to think about.

Secondly, you know, artificial intelligence can be used to do good and they can be used to do bad, right? So cyber risks are already potentially micro-critical, you know, financial institutions today are under constant attack, and that could be scaled up and made more

efficient by artificial intelligence. So there's a kind of tug of war going on where, you know, artificial intelligence is used to increase efficiency, but it's also used to attack institutions and then to defend against the attacks. So on net, where we come out in terms of stability from a cyber resilience point of view, I think, is a second thing.

The third thing, I would argue, is like aggregate stability and aggregate fluctuations, right? So artificial intelligence, such as generative AI or large language models may be very good in terms of understanding the cross section of risk, in terms of using past data to tell me how to allocate credit or liquidity across the economy, but is it good in understanding the general equilibrium effect of all of us using those kind of tools and what that is doing to the financial cycle as a whole.

So the analogy here is, you know, the run up to the subprime mortgage crisis where financial innovation and credit risk modeling was being used, and that might have been effective cross-sectionally in terms of allocating credit, but in some, it generated a huge amount of instability. And so like the general equilibrium effect, and how it all adds up, whether that is captured by the models, is very unclear at the moment.

Shen Yan: Okay. So you think it's very unclear, and it's about whether you can do good or bad, and also the algorithmic herding that given that we use fundamentally similar algorithm. So what's going to happen to the aggregate level individually?

This aggregate level is good, but every level is unclear. Dr. Furman, what is your opinion on macro stability?

Jason Furman: Yeah. So in the aggregate level, let's divide the question into the next two and a half years, which is the time horizon that central banks are operating on. On that time horizon, you're getting a lot of demand, and you're not getting very much supply.

The demand is coming from building data centers, building the electricity to power the data centers and the like. The reason you're not getting a lot of supply is

you're hiring a lot of people in businesses throughout the economy to do everything Tobias is talking about, but they haven't figured out how to do a lot of the things Tobias is talking about. They will eventually.

But in the short run, it means you don't increase your output, but you increase your labor input, and so your productivity goes down. That's a J curve. That means, in the short run, if anything, this is going to make it a little bit more difficult to manage inflation. Require somewhat higher interest rates to stabilize the macro economy.

When we look past, I don't know exactly when, two and a half years, five years, 10 years, there's a lot of upside for productivity growth here. That is the far and away, the most important thing for China's economy, for the US economy, for any economy around the world, that, broadly speaking, will translate into rising incomes, also higher interest rates. And the big question most people have, and I don't know if you want me to address it now or later, is, you know, what the impact of that on unemployment.

Shen Yan: Okay, so when you talk about employment, so let's move naturally to the next question, because today, here we are not only talking about finance, about economics, we are also talking about the fate of human workers. A large scale of concern is about how much AI will replace the human workers. I think there's a statistic saying that there is a study by Alfred and off spawn, in the year of 2017, it shows 47% of US employment is at the high risk of automation. So overall, it's still a concern, both for US and for countries like China that we have a huge labor force.

So what do you both think about the impact of AI on human employment? In the sense that what type of jobs will disappear and will AI be dominated by this substituting effect, for Dr. Furman I would like you to give the overall picture, and then for Dr. Adrian, can you give us a more focused picture about the finance industry, please?

Jason Furman: Yeah, if we were holding this panel

200 years ago, I wouldn't be able to join you from Cambridge, Massachusetts in this manner. But we would find it unimaginable that the majority of our population wouldn't be engaged in agriculture. In fact, only a tiny fraction of our population would be engaged in agriculture. And that one, we'd still have enough food, and two, that we'd still have jobs. Well, what happened?

First of all, we got new types of jobs. All you know, everything Adrian's been talking about, couldn't possibly have imagined any of those jobs in the finance industry 200 years ago.

The second is, as people become richer, they consume more services. So there's a lot more people working in hotels and restaurants today. Those are jobs you would have understood 200 years ago, you just would have been shocked that people were rich enough to eat out so much and travel so much.

Third, a lot of these technologies replace parts of jobs, not all of jobs. As we have more spreadsheets, we don't have fewer accountants, we have more accountants, and they're doing more stuff. But finally, this can be a dark side to it, which is you can get more inequality as people lower their wages in order to compete with the robots. Broadly, that story is why I'm mostly sanguine. But one, there is this downside of inequality. Two, the pace at which this happens matters a lot.

And then the last thing I'd say is we did not sit passively as the transition happened, as meant as agriculture went away, we made high school universal. As manufacturing diminished, we dramatically expanded college, and so something at that scale to meet this challenge, rather than assuming it's always happened automatically, because it hasn't quite happened automatically, I think is an important part of the answer here.

Some of that actually, is education. I think the more education makes you a little bit more impervious to these shocks, although with the latest AI, that's raising some questions about that as well.

Shen Yan: Okay, more education, then Adrian.

Tobias Adrian: Yeah, so I fully concur that, you know.

We oftentimes start to take today's economic of invented and created which we may not be able to imagine today, right? I mean, smartphones have completely changed our lives and the lives of everybody around the globe. And you know, when they first appear, they were mainly dismissed as something that was not needed, and it's only once it's out there that people understand how to use it.

So it's going to be difficult for us today to understand how all of this is going to be used in the future, and how that can create jobs in the financial sector and in the economy. So while the displacement of workers by technology has been a theme for hundreds of years, right, these are debates that are going back to the 19th and 18th century, right, where people are arguing, well, technological progress will make people irrelevant. You know that has never happened, because new ways of using people has continuously been invented. So in the financial industry in particular, while it is true that there could be substitution, there could also be a great amount of inclusion.

So new technology and artificial intelligence could, in principle, be an opportunity for countries to be part of the global financial system that previously were not. So, you know, there is potentially a great equalizer here, in that, you know, having a smartphone, being able to interact with these large language models, you know, could allow people that are not based in New York or Shanghai to participate in the financial sector.

And you know, finally I would also agree that, you know, the way in which humans are using the technology may increase output tremendously. So for example, in scientific research, technology has certainly allowed scientists to generate new papers and new results more quickly. So the amount of research and patterns that is put out there is increasing very steadily, and that is certainly true in terms of financial innovation as well. To what extent that is going to translate into...

we'll find ways to deploy the technologies for new output.

Shen Yan: Okay, both of you are quite optimistic about the future. And Dr. Furman gave us a situation about the changes, update of industrial structure. And Adrian also mentioned the positive side of AI, like equalizing, inclusion, type of effect, and also its help of research and so on. But in the short one, there still can be a disruptive force.

So my last question, because the time is very limited, my last question for both of you, is, what can the regulators and policy makers do to help in this structural transformation situation that we may still face a large scale of unemployment for some of the workers? And for Dr Furman, do you think that we need AI department? Do you think that we need universal basic income or other types of subsidy that proposal should go on? And then for Adrian, what do you think that the policymakers and regulators can do to stabilize, to foresee the potential threats of AI, so that we get prepared to stabilize the financial system.

Tobias Adrian: So, you know, in early August, we saw huge market swings. So you know, in Japan, the Nikkei lost 12% in one day. In the US, implied equity market volatility shot up from the mid 20s to levels above 60 intraday. These are the kind of levels that are usually observed in severe crisis, and one thing that market commentators pointed to is potentially the role of correlated trading, where algorithms are using potentially similar strategies that trigger sales. So to some degree, systemic risk could be increased because of correlations across signals that are amplifying downside moves.

So we saw some of that already back in 2010 in the stock market, which was called the flash crash, and then later in the treasury market in 2014, and at that time, there were adjustments at the microstructure level that really mitigated those risks. So these are stopping rules and the way in which trading venues are organized.

But of course, there could be broader macro

consequences as well, that could amplify, you know, systemic risk in principle. So this is certainly an area that regulators are studying, and a big challenge is that the kind of data that is needed by policymakers to understand those risks is not necessarily the data that was needed previously to assess risk, so what entities to collect data from, and what kind of data to collect may be different. And so you know having additional transparency, maybe first order here.

Finally, I would also come back to what I alluded to earlier. You know, there could be efficiency from a micro point of view, but from the macro point of view, risks could be increasing, so the general equilibrium effects may not be taken into account by these models. So, you know, the policy makers, such as systemic risk regulators, are certainly also trying to figure out how to use artificial intelligence to inform their own visibility and their own judgment on how stable the system is, specifically from an aggregate and systemic risk point of view.

Shen Yan: *Okay, so there's still a lot to do and to be fully aware of, and AI has created many new challenges. Dr. Furman?*

Jason Furman: Great, let me suggest five principles that should guide us as we think about regulation. The first is balanced benefits and risks. This sounds really obvious. Who could be against that? Well, a lot of the AI regulatory discussions are against that. They say, don't deploy this until we're sure there are no risks. Get rid of all the risks.

Well, there's a lot of risks of slowing AI down and not solving car accidents and climate change and digital tutors and all the other things. And so you have to ask every time you regulate, not just what are the risks that the AI does something bad, but what are the risks that you regulate it too much and you don't get all the good things that you want. And you balance those.

Second, compare AI to humans or to the alternative, don't compare them to the Almighty. If your autonomous car crashes, that's not a reason not to

have autonomous cars. You want to ask how much it crashes compared to how much humans crash. People write endless papers about AI end bias, for example, that's an important thing to write a paper on. But you know, when you're making a regulatory choice, humans are supervised too. Is the AI better than them? Is it worse than them?

Third, this gets to something in your initial question. Whenever possible, use domain specific regulation. Maybe we need a super regulator. I'm pretty skeptical what we need, just like I'd be skeptical of a regulator for linear algebra. This is a tool that's used in lots of different sectors. It's the domain specific regulators. So Adrian was talking about correlated strategies across financial markets. An AI super regulator, they don't understand financial markets. They don't understand financial stability. They couldn't do that.

You want your financial regulators looking at that issue. But you need more AI expertise in your regulators. Same thing, the Food and Drug Administration in the United States needs to know more AI, the Highway Administration that regulates auto safety, they do. So get all your the regulators you already have to know more about AI.

Fourth, regulations should not be a moat protecting incumbents. A lot of the big companies have welcomed regulations. Some of that, I think, is public spiritedness on their part, and should be welcomed and applauded. Some of that is that they know they can comply with the regulations, the smaller upstarts can't, we should be very skeptical of that.

And then finally, not every problem caused by AI can be solved by changing AI. You talked about Wuhan, if taxi drivers are losing their jobs due to AI, the answer is not to reprogram the AI and ban AI that has driverless cars, and China's not doing that. The United States is not doing that. The answer is maybe a training program for those taxi drivers. Do we need UBI? I don't think we need it yet, but if we have larger scale increases in productivity and replacement of jobs, then we would, and by the way, then we'd be able to



afford to have it as well. So a lot of the solutions to AI's problems are not about regulating AI. They're about separate programs in labor markets, the tax system and the like.

Shen Yan: Okay, great. So pretty much as closing

remarks, you've given five principles, the balance between risk and the benefits, compared to human, no super regulator, and also, especially the last one that I enjoy a lot, problems caused by AI may need to solve with forces and understanding that is outside of AI. 🙏