Transcript

Conversations on Central Banking: Macroeconomy in the Age of Artificial Intelligence
Federal Reserve Bank of Cleveland
November 17, 2026

Panelists

- **Nicholas Bloom**, William Eberle Professor in Economics at Stanford University; Senior Fellow of SIEPR; Co-Director of the Productivity, Innovation and Entrepreneurship program at the National Bureau of Economic Research | slides
- **John Fernald**, Schroders Chair in European Competitiveness and Reform and a Professor of Economics at INSEAD Business School | slides

Moderator

- **Anne Artrip**, Senior Financial Economist, Federal Reserve Bank of Richmond

Ina Hajdini:

Welcome everyone to the Cleveland Fed Conversations on Central Banking. I'm **Ina Hajdini** with the Cleveland Fed Center for Inflation Research, and thank you all for joining us today. This conversation is part of the work that we are doing at our Center for Inflation Research to help researchers, policymakers, and the public to understand inflation and monetary policy. We typically hold two or three of these events per year. The agenda and the link to all the previous conversations in this series are accessible in the Zoom resources tab.

Please subscribe to our newsletter to be informed about future editions of conversations and other news about inflation and the center. Today, our topic is macroeconomy in the age of artificial intelligence. Before we start, this conversation is being recorded and the recording should be up on the center's website shortly after the event. As a disclaimer, the views expressed during this program are those of the panelists and not necessarily those of the Federal Reserve Bank of Cleveland or the Federal Reserve System. I'll now turn the virtual floor to the moderator of this session, **Anne Artrip**.

Anne Artrip:

Thank you so much, Ina. I want to thank the Cleveland Fed for hosting this event and for inviting me to moderate the discussion today. Today's conversation will be led by two distinguished researchers... both of them, before we'll hand it over for the opening remarks. Sorry, can you still hear me?

Ina Hajdini:

Yes, we can hear you.

John Fernald:

This document is a transcript generated by a third party and may inadvertently include errors or inaccuracies. The opinions expressed are those of the participants and do not necessarily represent the views of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.

We lost you for a moment.

Nick Bloom:

We lost you for a moment there.

John Fernald:

Then you came back.

Anne Artrip:

I will continue.

Ina Hajdini:

Anne, are you there?

Anne Artrip:

I'm here. I'm sorry. Zoom is deciding to restart on me. Hopefully it works now. Okay. Today's conversation will be led by two distinguished researchers, and I'm happy to introduce both of them before I'll hand it over for the opening remarks. First, let me introduce **John Fernald**. John is a professor of economics at INSEAD Business School and Economist Emeritus at the San Francisco Fed. John is one of the foremost economists studying productivity and long run economic growth, and his work has helped us understand the productivity boom associated with the IT revolution of the '90s and early 2000s.

We also have **Nick Bloom** with us today. Nick is a professor of economics at Stanford University. He's also co-director of the Productivity Innovation and Entrepreneurship Program at the NBER. Nick is one of the leading economists studying how firms respond to technological and organizational change. He has been at the forefront of measuring the adoption and economic impact of AI. I'll hand it over to each of them for opening remarks. John, let's start with you.

John Fernald:

Great. Thank you very much. I'm very happy to be here and it really is a pleasure. You should be seeing my slides at this point. Okay, great. I guess my starting point is, we see AI just about everywhere. We all have experience of seeing many tasks that are clearly more productive that we do all the time, often because of AI. But paying homage to Robert Solo, my narrow question for today is, do we see AI yet in the aggregate productivity statistics? If so, how? My answer to that is going to be a qualified yes.

The qualification is, it's not necessarily in the way that we want or expect. Now looking ahead, I am cautiously optimistic, but I am so cautious you might well mistake it for pessimism. Let me turn to some data. Labor productivity is ahead of its pre-pandemic trend. What this is showing is the log level of labor productivity in the US business sector starting at the end of 1995 when dot-coms, other ICT related gains first started showing up in the productivity statistic. Now there's a dashed blue line, which is an

This document is a transcript generated by a third party and may inadvertently include errors or inaccuracies. The opinions expressed are those of the participants and do not necessarily represent the views of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.

underlying trend, which grows fast in this trend through 2004, then slows from about three and a quarter percent to about 1.5%.

I estimated that trend through 2019 right before COVID, and then I continued the blue dotted line thereafter. Of course, you see during COVID and immediately after, productivity spiked well above the trend line and then returned right back to the trend line. But since the end of 2022, productivity has grown almost 3% a year and is now about 6% above its pre-COVID or its pre-2022 trend. Now, productivity rarely fits a trend line neatly, and one aspect of that that you often have to worry about is the business cycle.

That's going to hit labor productivity through labor composition, capital deepening, variations in factory utilization. One very simple reduced form relationship that I did with a paper with Huiyu Li was to progress labor productivity growth on the change in unemployment. It's very simple, but the red line here gives you the fitted value from that. It's going to fit this surge in the great recession and the surge in COVID as well as the slow return after the great recession to the trend and the fast return after COVID.

Now, this is getting the general pattern that goes back to the 1980s that when unemployment goes up, productivity tends to go up. Well, why is that? Well, younger, less educated, less productive workers lose jobs, so labor composition boosts labor productivity. Second, each remaining worker has more capital, which also boosts labor productivity. In both cases, these are cyclical and they unwind as the labor market recovers. Now, as you can see, that reduced form cyclical regression explains none of the pickup after 2022.

Let's dig more into the growth accounting. Well, what these bars do is show average labor productivity growth over four periods going back to the 1990s. Each bar is split into the standard growth accounting components. The height of the bar is labor productivity growth in the period, capital deepening, that is capital per hour. Labor composition mainly increases in educational attainment and somewhat experience, and the residual, which is TFP. The swings across periods are mostly in TFP, in this blue part.

Over longer periods of time, TFP is a broad measure of efficiency, and that's what you want to see. Well, since 2024, TFP growth has grown about one and a quarter percent a year versus about half a percent from 2005 through 2022. The thing is, as I said, TFP is a residual. It's everything about output and productivity that we can't explain with observables. In the short run, it moves with the intensity of factor use, mainly with labor effort and the work week of capital.

My quarterly TFP data that I post every quarter or twice a quarter on my website has a model-based proxy for utilization, which is based on Basu, Fernald and Kimball estimate. It's a particular model, particular estimate, but if I just apply that, take that utilization measure, I can split measured TFP growth into the contribution of factor utilization, that is workers working harder, capital, working more hours, that's in gray. Then the residual, which is in black, utilization adjusted TFP. If we start with 2023, utilization adjusted TFP surged.

Okay, I don't have any fully satisfying explanation, but it doesn't seem like it's LLMs or AI more generally. It was too early for LLMs. It's not clear why earlier machine learning things would suddenly

This document is a transcript generated by a third party and may inadvertently include errors or inaccuracies. The opinions expressed are those of the participants and do not necessarily represent the views of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.

show up then, you never know. But if I look since 2024, the picture's different. Yeah, we've had pretty good TFP growth, one and a quarter percent, but the vast majority of it in my estimates is factor utilization. Utilization adjusted TFP, the black bit, has grown no faster, maybe even more slowly than it did before the pandemic or before 2022. In these estimates, utilization explains all of the pickup in TFP.

Now, I could list caveats to my utilization measure. It's just one measure, for example, and list other caveats, but taking it as given, why is it intuitive that utilization growth has been so strong? Well, what we've seen in the past few years is pretty strong demand combined with a lot of uncertainty about business cycle, about tariffs, about AI itself. The AI has certainly contributed to the solid US demand in recent years. It spurred investment in chips and data centers and power. It's also created a wealth effect, the strong stock prices have driven consumption, but it's also contributed to the uncertainty.

Yesterday, as the New York Times summarized it, the Fed chairman highlights economic uncertainty around the AI boom. We think about what are businesses going to do in an environment which they have demand, but there's a lot of uncertainty about what the right size and mix of the labor force is. Well, lots of people have pointed to that leading to firms to being in something of a wait and see mode. I think this was MarketWatch a couple weeks ago said, "It's the strangest job market the US has ever seen. The low hire, low fire economy has gone on three years." In other words, to meet demand, what do firms do?

They don't want to make long-term commitments in workers, but they can work their existing people. They can work their existing capital harder. That's going to show up as faster measured productivity growth without it being a gain in efficiency. It's basically an endogenous response of unmeasured inputs to demand. That's my qualified yes. AI does appear in the productivity statistics, but so far it's through strong demand and heightened uncertainty leading to rising factor utilization, not yet very completely clearly through efficiency.

This kind of productivity growth, of course, is inherently temporary because you can only work people and machines so hard. If I look ahead, well, we can ask, are the broad efficiency gains going to arrive and persist? Well, quite possibly. Certainly, introspection tells me I use Claude to create the figures for this talk. It was orders of magnitude faster than anything I or my RA could have done previously, but there are plenty of caveats. One caveat is bottlenecks. Chad Jones has a couple of papers related to this with different co-authors, which basically growth is held back by the tasks that are essential, but also the hardest to automate and improve.

In fact, in the Jones and Tonetti paper, I think if the productivity of the software industry goes to infinity, GDP growth and their benchmark model rises 2%. That's a sense in which it's held back by all the other bottlenecks. Now same time, you can think about monopoly power, whether it's by AI firms or by AI users who use AI to build better moats around their business and allow them to charge higher, higher markups. Well, in the model of Aghion et al, the way that works is these monopolists then have reduced incentives to keep innovating as fast, because they'd be competing with other monopolists and it's harder for new entrants to break in.

Think about things like cybercrime, supermarket refrigerators being hacked and turned to defrost mode is just one recent high profile example, but AI superpower scams and misinformation as well. Actually, Jones and Tonetti, I think the latest version I saw, their final sentence was, "In a world of weak links,

This document is a transcript generated by a third party and may inadvertently include errors or inaccuracies. The opinions expressed are those of the participants and do not necessarily represent the views of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.

AI's benefits may arrive slowly while the dangers arrive quickly." To conclude, I'm optimistic about the future, but only cautiously so.

Anne Artrip:

Thank you, John. Over to you, Nick.

Nick Bloom:

Thanks very much. I always learn so much from listening to John, so we're going to be pretty aligned. I will share my slides. Here we go. I'm going to talk about what I call firm data on AI. It's with a large number of co-authors and many banks and central banks around the world. I should say we have an MBR working paper out on this and we just recently updated, so all this data is live online. For monetary policy, partly why many people are listening, it's pretty critical to know the impact of AI. I was actually at the event and spoke at the event John mentioned yesterday, the Chairman Warsh, and it's like AI is pretty much everything at this point in forecasting the future.

What do we do? What's the data I'm going to show you today? With a team of eight central banks, we surveyed around 10,000 CFOs, CEOs in the summer of this year, and I'll show you exactly what they told us. Before I show you the results, I just explained the survey process so that hopefully you trust the data. Then I show results on adoption, and then I finally show you results on the impact on employment and productivity. We use some central bank survey panels that already exist and add in AI questions. What are these panels?

On the left is one in the UK called the Decision Maker Panel, and on the right is one in the US called the Survey of Business Uncertainty. The left is the Bank of England, the right is at the Atlanta Fed. How these are structured is these are survey panels that we recruit people initially by phoning them up. If I'm recruiting Anne, I'd phone her up, get her on the phone, it takes about 10 minutes. At some point when she agrees to be in the panel, then her details goes into this online panel and she's sent a monthly survey. Why do we do this? It's expensive. It costs about a million dollars a year per survey.

The reason is you know you're getting humans. In the world of AI, you want to be sure that you're actually talking to a person, and so we know we get a person and we know their position, or as I'll show in a minute, we're primarily surveying CFOs. Who do we go for? We actually ask for the CFO. CEOs in some senses are great, but they are very hard to get hold of. CFO's our perfect candidate. They're very senior execs. They know what's going on. They can talk comfortably about their firm. I'm going to argue in a minute, if you want to look at technology impacts of AI, productivity impacts of AI in the next three years, you probably want to talk to folks like CFOs.

The reason is, as John mentioned, most of the productivity impact debate is about bottlenecks adoption. Will firms use it? If you want to look five, maybe 10 years out, you want to talk more to technologists, my friends in Silicon Valley, but I'm going to focus very much in the next three years, and I think it's all about adoption. We in a sense have the perfect sample. It's primarily senior managers, think CFOs. Before I show you the AI numbers, just to confirm the data looks reasonable, we've been running these panels now for more than 10 years. On the left, I'm just comparing what they tell us about sales growth to GDP.

This document is a transcript generated by a third party and may inadvertently include errors or inaccuracies. The opinions expressed are those of the participants and do not necessarily represent the views of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.

On the right, you can look at employment, top row is US, bottom row is UK. It basically matches very well. What does that say? It tells you our panels which are one to 2,000 firms in each country effectively match the aggregate. They're representative by age, size, and industry. Overall, they do a pretty good job. We have also, for the last 10 years, been asking them to forecast, but not AI, forecast their own sales and employment. The nice thing about our panel is, if I interview Anne, I'm going to interview her a year later. If I ask her a forecast looking ahead of you, if I talk to her a year later, I know how accurate that's been.

What you can see is these firms are pretty accurate about forecasting their own sales and employment. It's not perfect. Shocks happen, including COVID, which happened during our period, but they're pretty good. With that in mind, I don't want to claim this is perfect, but it is a pretty good survey process. Certainly, if you want to look at three years ahead for AI, I think it's one of the most informative ways to get forecasts. What do we see? The first question we ask is just an adoption question, and this is in a sense the easiest question to answer because it asks people about their own use.

It's hard to talk about others in your firm, but let's talk about your own. On average, how frequently do you personally use AI technologies in a typical working week? We give them four answers, not at all, up to one hour, one to five hours a week, or more than five hours a week. When I asked our PhD students in Stanford, they said, we want more than five hours a day as an option that's missing, but they're answering honestly as you put it in the bottom bucket. If you look at that, I'm guessing most... Well, I don't know.

I'm guessing probably most listeners are probably in one of the bottom two buckets, not quite so with our CFO, CEOs. Step back, remember who these are. These are CFOs of moderate to pretty large companies. They're probably on half to $1 million a year. They're probably in their 50s. If you look at them, 20% roughly just don't use AI at all. If you live right in Silicon Valley, it seems surprising, but we see that. The most common response is using up to an hour a week. There are some folks using one to five and there are about another 15% using more than five hours a week.

Black bar on the left is all countries. Light blue is the US. I'm going to primarily focus on the US for the rest of the talk. We also ask about AI use in their firms, and if you look at the light blue bar, the US, more than 95% of US firms report using AI. This is 10 plus person firms. This doesn't surprise me. Honestly, it's hard to think of a firm now that's not using AI in some... The average size is about a hundred employees. It is very hard to find any firm not using AI, at least for one of their employees. AI adoption at this point is pretty universal.

This shows why AI is just incredibly rapid in adoption compared to pretty much any technology we've ever seen. In a paper with various co-authors, we have looked at the rate of adoption of different technologies, hundreds of technologies, by looking at how fast they show up in jobs postings. Generative AI has been adopted at something like three to five times faster than the average, and faster even than cloud smartphones, things that rolled out very quickly. Why is that? Because we're rapidly adopting it both at home and at work.

In some senses, this time is different just because the speed of adoption is unlike anything we've seen before. John and I worked on the productivity boom, as John mentioned, for computers and the internet. That rolled out, but that took one to two decades to see broad rollout, whereas AI is more

This document is a transcript generated by a third party and may inadvertently include errors or inaccuracies. The opinions expressed are those of the participants and do not necessarily represent the views of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.

counted in terms of years. We also ask about dollars spent. Exactly how many dollars do these firms spend? What we see is this spread to put one top line figure for the US. On average, firms are spending about $2,000 a year per employee. That is an enormous amount. That is an absolutely enormous amount.

It's about $300 billion a year or something like 1% of GDP at this point. It's pretty similar actually our numbers in your aggregate to other folks like Gartner's, CBO, we're coming up to similar numbers, but these are absolutely huge numbers. Perhaps more eye-opening is this, which is the skew. This AI spending is incredibly skewed, so 50% of AI spending is being accounted for by the top 2% of firms. Again, if you know the data, this isn't perhaps that surprising. If you look at Anthropic's S1 filing, two customers alone accounted for about 10% each. If you look at the ramp data, you see that.

In fact, if you just look in US firms and look at R&D and capital expenditure, it's similarly skewed. But just to be clear on thinking about rollout and fragility, AI spending is dominated by a few extremely large firms. Finally, what's the effects on employment and productivity? We ask this question, first productivity, which we're going to define here because we're serving execs as sales per employee. I'm aware that's not the same as John's total factor productivity, but this is something CFOs can easily answer.

We ask how has AI technologies affected sales per employee looking back over the last three years. We should think of '24, '25, '26 and looking over the next three years, which is seven, eight and nine. What do we find? Firstly, over the last three years, very aligned with John, we basically see close to nothing. About 85% of American firms say AI has not affected their productivity whatsoever over the last three years, and about 15% said it's small and the rest say nothing. If you look at the average, over the last three years, about half a percent, that's about 0.15% a year.

It's not nothing, when you've got an economy the size of the US 0.15% still is a big deal, but as John says, it's not really going to show up in aggregate figures and this productivity is measured so noisily, it's very hard to detect this. I would say looking back, we don't see much. Looking forward, the picture changes quite a bit. These CFOs in the US, the most common answer now is looking ahead for the next three years is a small positive impact. 38% of them notice, 22% are talking about large impacts, and around a third are talking about no impact. There's been a clear shift out.

CFOs are much more optimistic over the next three years. If you want a number on that, that averages out to about 0.8% a year, and that is a big deal. You saw John's numbers earlier, productivity growth in the US is something like 2%, 3%, so adding 0.8 is material. My takeaway from this is we are basically at a point of inflection. AI today has not had much measurable effect in productivity. We just don't see it in many of the data series at all. But looking ahead, it looks like it's starting to have a material effect, '27, '28, '29 in particular. That is pretty much what the stock market is capitalizing in now.

Finally, what about employment? We ask similar questions in the impact of employment. Looking back over the last three years, basically nothing. Almost 90% of American firms said it has not impacted employment at all. The aggregate is a big zero that matches many other datasets. We are just not really seeing any aggregate effect on employment. Perhaps tiny effects on new graduates, even that is unclear. What about looking forward over the next three years? Here, the net numbers are shown in a

This document is a transcript generated by a third party and may inadvertently include errors or inaccuracies. The opinions expressed are those of the participants and do not necessarily represent the views of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.

minute of a similar. They're not very big actually in the scheme of the US economy, but there's a lot of mass at the tails.

I think what it's telling us is looking forward over the next three years, AI is going to churn up the labor market. A lot of firms expanding, there's a bunch of firms contracting, and so we are looking ahead to more employment volatility at the individual level. Overall, it's negative, but remember these are existing firms. There's new firms coming in. My read, our read is the employment effects for the next three years are zero. AI uses here, pretty much all large American firms are using it, individuals, execs are talking using it two, three hours a week.

It's rising fast so far, and I think in the next three years, I don't see aggregate effect in employment, but I see a lot of churn. At this point, productivity is really starting to have an effect. We've not seen it yet, but I think we're entering the phase. This is the inflection point. With that, I will stop sharing now.

Anne Artrip:

Thank you so much, Nick. Before we open it up for the Q&A, let's start by getting some reactions from each of you on the others' remarks. John, would you like to respond to Nick's remarks?

John Fernald:

Yeah, I would just say first, thanks, Nick. This is terrific. You have helped illuminate so many of the great shifts we've seen in recent decades and different aspects of them. Asking firms is great as a way of getting information of what they expect. What's harder to interpret about this to me is, how much is it to a firm increasing markups and improvement in price discrimination is the same as a true cost reduction from an efficiency gain. It's hard to differentiate from this whether it's, they're going to have more market power or they're really going to be reducing their costs.

Nick Bloom:

Great point. I think John and I are very aligned. I'd say where I see it right now is employment effects, we just don't see them. Productivity effects, not yet, but we're starting to enter that. Of course, as John says, the stock market is looking ahead, so the stock market's capitalizing it all in. The data all lines up. As John says, who's going to be the winner? Imagine 10 years from now we're going to see some massive productivity effects. Where do these gains go? It's unclear. In fact, this was some of the discussion yesterday with Chairman Warsh. It depends honestly on competition.

If you look at tokens, token prices are falling by about 50% every few months. In that sense, the big Anthropic, Gemini, ChatGPT, et cetera, are competing with each other, prices are falling fast. If that is true across the economy, we all gain because our prices fall very rapidly and everything, particularly services get cheaper. If some sector gets monopolized, say maybe chip production, then we don't gain whoever owns that gains. My broad sense is the US is pretty competitive. I think in general, we are all going to gain, we're going to see prices fall.

One interesting thing is AI is very heterogeneous effects on the economy. There are some parts of it, particularly tech sector, that are seeing enormous productivity gains, and there are others, think of healthcare teaching, it's very hard to make use of it. One of the concerns is this two speed rollout. We

This document is a transcript generated by a third party and may inadvertently include errors or inaccuracies. The opinions expressed are those of the participants and do not necessarily represent the views of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.

have parts of the economy that are regulated very in-person that can't change and other parts that change. For the Fed system, it's tricky because you've got two measures of inflation, but that I think is a harder question to ask.

Anne Artrip:

Since you brought up, Nick, who's going to gain, I think it might make sense to start with the first question from the audience, which is pointing in this direction. It's asking how can we ensure that the efficiency of market gains once they start showing up in the data, how can we show that they're distributed to workers and don't just sit at the top? Any thoughts on that?

Nick Bloom:

I'm happy to go first and hand to John or John, whichever way you...

John Fernald:

Okay. Why don't you start?

Nick Bloom:

Yeah, I just thought, it builds on what I said. Competition is critical here. As long as we have honestly flexible markets and good competition, firms will compete. You see this exactly in token prices. Token prices are going down fast because there's a lot of competition. The issue is when you have monopolies or heavy regulation that prevents productivity gains being spread out. The US generally has a pretty good track record of being flexible and new technology. Think of computers or cell phones.

My cell phone, this is amazing technology and Apple invented it, but everyone, to be honest, there's a lot of replicas and we all have them and the prices in real terms have fallen a lot. I'm reasonably optimistic. I'm not so worried about a massive monopolistic sector. I'm grateful there are multiple large AI firms that are competing with each other, at least until now. I think it'll be broadly shared, but I'll pause, I'm interested to hear what John says.

John Fernald:

Yeah. Economic models, it can go different ways. I mentioned Jones and Tonetti because that's a nice one in which depending on the nature of technology, are there essential things that you can't automate that humans have to do? In which case, labor fare goes to 100% asymptotically because even if they're only a few things, they get all the rents or capital rents could go away, but if everything could be automated, capital share goes to 100% and then it's who owns the capital. We don't know. Yeah, I think we need appropriate guardrails, and that's a challenging thing on which I am no expert.

One is antitrust, which to make sure we keep this competitive edge. Part of that is if I think about... When I read The Big Short, the big takeaway to me was when I was naively reading economics models, what does a financial sector do? Well, it overcomes information asymmetry problems between savers and borrowers. Then I read The Big Short and it seemed like what the financial sector is doing is using information technology to exacerbate information asymmetry problems, to make sure they have

This document is a transcript generated by a third party and may inadvertently include errors or inaccuracies. The opinions expressed are those of the participants and do not necessarily represent the views of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.

more information than their competitors. Then they can charge a monopoly price on what is fundamentally an intermediate input of financial services.

It's going to depend, what are we making more productive? If we are boosting the monopoly power in some areas, and certainly, financial markets have a huge incentive to try to do that, then you could be making the economy less efficient in some ways. Even as you're making other aspects more efficient because we know that markups and intermediate goods are particularly damaging. But still, getting a bigger pie is certainly helpful however it ends up getting split. That's one of the many open questions to what degree we'll get that.

Anne Artrip:

Thank you both. I wanted to ask, given that we don't really see the impact on the current productivity measures, so even from the surveys, I was wondering how reliable are our traditional measures of output and productivity in the era of AI? I'm thinking AI impact knowledge workers, and we all know that the output of knowledge workers can be hard to measure. Is there risk that we are not capturing it right just because we don't have the right measurement tools?

John Fernald:

Yeah. Should I take that first, Nick, and then turn it to you?

Nick Bloom:

Yeah.

John Fernald:

First of all, it's a great question and always a question because data get revised. In the late '90s, earlier in my career, but I remember every data revision, things look a little bit better. We got a little more certain. It wasn't until '99 that financial markets were really sure that we'd gotten this internet boom that had started four years earlier. Then from 2004 to 2009, every national accounts revision made the post-2004 period look worse. The revisions may be larger when there are bigger and more one directional, when there really are big technology changes.

Beyond that, I can look at, say, well, how much of what we're seeing is learning and experimentation, investments in intangible capital that we hope you'll get a return on. Quality improvements, like if this makes quality better, or introduces new unpriced goods, new goods and quality improvements have always been hard to measure. That's been debated with the internet. Google Maps and the like substantially raise wellbeing in ways that weren't captured in GDP. Mostly conclude we're not missing enough because we've always been missing stuff.

Not enough to make up for the slower pace, but this time could be different. There's also not sure though, which direction do things even go in terms of broader measures. The social value of social media is debated and there are debates about, I think the technical term, the Cory Doctorow term is enshittification of internet platforms, if I can say that on a Fed panel, or even excessive standardization

This document is a transcript generated by a third party and may inadvertently include errors or inaccuracies. The opinions expressed are those of the participants and do not necessarily represent the views of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.

like optimization routines. Everyone uses the same one and therefore every downtown looks the same and could be anywhere. You lose some of the variety that you could have had.

We don't know. Or, a year ago, Harvard Business Review talked about work slop, when your colleague or worse, your boss sends you a beautifully written and beautifully formatted bit of meaningless slop.

Anne Artrip:

Any thoughts on that, Nick? Want to chime in?

Nick Bloom:

It's a very interesting question. It's always been hard to measure productivity and classically the one I often look at is just value added per hour worked. Value added is basically GDP and per hour worked. There are some challenges to measuring that. I think the hardest measurement thing actually, the impact of AI could be over measuring inflation. What I mean is there's a sense that it's going to heavily impact some parts of the economy. Take anything online, anything electronic related to software, a bunch of services, their price could easily fall incredibly rapidly by 20%, 30% a year potentially.

On the other hand, there are other sectors like healthcare where the effects are pretty limited or teaching because it's person to person. Why is that an issue? That's an issue actually if you are the Fed or any other central bank and you're trying to target monetary policy because Lord of the Rings, there's no longer one inflation measure to rule them all. Because take a 65-year-old, they're consuming a lot of healthcare and for them prices may be going up and take a 20 something and they're consuming a lot of online services. For them, prices may be going down pretty fast.

I think it's just generally going to make monetary policy and government policy just a lot harder if it has this bifurcation. That's one of the senses I get from folks around here, that there's winners and losers. Unlike other technologies, like cell phones everyone benefited, but AI is I think much more heterogeneous. That's a measurement issue as much as a policy issue in general. I don't know which way it goes. It obviously affects politics as well because these folks have different levels of political power.

John Fernald:

Yeah. Nick, on that, so of course the Fed says it wants an average 2% PCE inflation, but that's going to mean, as you're saying... but then relative inflation rates across different things in economic models can cause questions whether that's the same target. It would be a question of whether the 2% target in theory remains the right one, or would it be something different?

Nick Bloom:

No, you're right. They have targeted and officially are mandated to target in average, but if you discover that the 60 plus, they consume services like healthcare and in-person services and don't consume a lot of digital stuff, their prices may not be falling. They have a very high propensity to vote. They're a politically powerful group. That seems to be one challenge for policy in the Fed system in general

This document is a transcript generated by a third party and may inadvertently include errors or inaccuracies. The opinions expressed are those of the participants and do not necessarily represent the views of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.

about, this is five years out, it's not today, but I can see this. If we really get this bifurcation, what are you targeting?

There's been debates about what the level of the inflation target is, but this would be more a debate about what exactly is inflation, because it's no longer one concept.

John Fernald:

It would matter for other, it would affect other implications beyond central banks. Like where you're saying, it's for the indexing of something like Social Security.

Nick Bloom:

Exactly right.

John Fernald:

Index to an average may be completely unrepresentative of what people over 65 are buying.

Nick Bloom:

Exactly.

Anne Artrip:

Lots of sources of uncertainty going forward. Here's a related question from the audience. If productivity gains were not visible in the past three years at basically full adoption, what would drive productivity gains for the next three years?

Nick Bloom:

Happy to go on that. I think the adoption is a width and a depth, and what we're basically seeing is pretty much, it's a bit like email. Take email. John and I use these databases. There's a company called Harte Hanks we used to buy regular data on IT adoption by firms. One of the questions they were asking in the '90s is, "Does your firm use email?" Pretty rapidly that question became pointless because everyone said yes. But the question is how much do you use email and how intensively? I think it's similar to AI. If one person, a hundred person firm is using AI, it's a yes, but that's very different from entire departments.

I have people working alongside AI. If you take me when I teach, how much does AI impact that? I think what we're seeing now is very shallow adoption. Lots of people are using it, lots of firms are using it, but it hasn't replaced standardized systems because it's difficult to change processes. There's laws, there's regulations on changes. It's more the intensity. That's why just on a measurement issue, by the way, one thing to raise is I think it's important to collect data on how much firms are spending on it and how much they're using it.

Quite amazingly, there is no data on firms spending on AI because it's not in company accounts. It's not a line item in the SEC. It is not in administrative data. The census has not really been collecting it, so

This document is a transcript generated by a third party and may inadvertently include errors or inaccuracies. The opinions expressed are those of the participants and do not necessarily represent the views of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.

how do we know? We don't. It's partly why we got into the survey business, but that gives you a sense of depth. If you are spending a hundred thousand dollars a year or a hundred million dollars a year for the same firm. That's where I think the progression is depth of use and extent across the business process.

Anne Artrip:

On that, how can we or can we at all distinguish between firms that are generally transforming the processes and the way that they work around AI and those that are just dipping their toes, simply experimenting, helping it to write an email here and there, but not really transforming the business? Is there any way to detect that in your work?

Nick Bloom:

That's hard because that's looking inside firms. If I take my institution, Stanford, if you're looking from the outside in, you've no idea how much we're using it. It's not even clear. I was talking to our president, John Levin yesterday, he's fantastic. He's long been my colleague, but in a sense, it's hard for him to know in an organization of 10,000 employees because we have two hospitals, how much it's being used place per place. One metric would be expenditure. It's not a perfect metric. There's a five out of 10 metric.

But firms that tend to spend a lot on AI, you've got subscriptions, for example. You're paying Google for Gemini. There's consulting, there's internal roles. You can look at getting into the weeds a bit, but there's a company called Revelio. They scrape LinkedIn posts. Can you see thousands or dozens of employees with AI things in their titles? That's the way researchers are going. They're looking at many metrics. Just in our expenditure data, it is incredibly skewed. It's incredibly skewed. Most organizations, I think a third report spending nothing on it, and then you have 2% that account for half of it. In some ways that's a lot of promise.

Coming back to what I'm talking about earlier, that's why the impact is not very big on productivity because while there's a lot of firms that use a tiny bit, to be honest, extensive use is very narrow. I think the consensus right now is really the only sector seeing much of a productivity effect is tech, and that's just not big enough to move the needle on the whole economy.

John Fernald:

Yeah. I would just add that modern businesses tend to be very complex, and that's very hard to change because you change something here and it has ramifications that you can't really predict. You have to make sure your core business processes continue and don't get disrupted by a hallucination or by anything else. You need somebody to like, how do we do this? All of that can take time. [inaudible 00:41:58] I think famously talked about how hard it is for existing organizations to benefit from IT. Often you needed a new organization built from the bottom up, but that takes time. It takes time for them to be able to compete with the big incumbents often, but sometimes it happens faster than-

Nick Bloom:

This document is a transcript generated by a third party and may inadvertently include errors or inaccuracies. The opinions expressed are those of the participants and do not necessarily represent the views of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.

One very local example, if I look just there, there's a window to my left. I'm at home because it's 8:30 in the morning in California. But my next door neighbor for many years when I moved into this house was Sebastian Thrun, one of the pioneers of the driverless car. In 2012, which is now 14 years ago, he would have a driverless car parked in his driveway. I remember my daughter at the age of seven sat in it and I was thinking, "Well, she'll never learn to drive because when she takes her driving test nine years later, the driverless cars will be everywhere."

Clearly, that isn't the case. It's not that the technology isn't great at this point. Driverless cars are demonstrably safer than humans and they have many upsides and not that many downsides. It's just hard. It's expensive to roll it out. There's a lot of regulations inertia. Still the forecast is it's probably for another 10 years to be at least human. That could easily be the path of AI. Some places it spreads out. We have driverless cars in central San Francisco, but in other places, it just takes a long time to-

John Fernald:

Yeah, I thought one of the reasons for that delay with driverless cars is it can be great in 95% of the time, but you have to deal with the 5% of the time when it can't... edge cases when it has no idea what to do, because it's not obvious. You got to train it to do those and it takes a long time to see all of the edge cases.

Anne Artrip:

Right. The same with AI. We need to figure how to validate and what to do and ensure accountability, all of those things. I wanted to talk a little bit about capital expenditure and AI infrastructure and data centers, semiconductors, all of that. Do you expect that will become a major driver of US economic growth?

John Fernald:

Well, of course. Right now it's a question how much it is a driver because of the fact that, I don't know, Jason Furman estimates that about three quarters of data center investment is imports. It's a big driver for selected Asian economies. It's a modest driver for us. There's the domestic content, including the construction workers and such, and to some extent design. That's why I emphasize somewhat the investment, but also the wealth effects of the companies that are... Well, stock market gains that have been quite good in recent years as boosting demand, on the demand side.

Nick Bloom:

John, you have the perfect answer to that. I'm going to twist it around a little bit and say potentially if anything, the end of it could be a recession risk. There's something that my colleague, Bob Hall, that just recently won the Nobel Prize called the Anna Karenina Theory of Recessions. He said basically expansions are all the same, but every recession is different. It's like the opening line of happy families are alike, unhappy families are different. It's true, if you look back at recessions, they're all caused by completely different things, a Gulf War, a dot-com boom, the plane goes into the 9/11, the global financial crisis, the pandemic.

This document is a transcript generated by a third party and may inadvertently include errors or inaccuracies. The opinions expressed are those of the participants and do not necessarily represent the views of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.

I think there is a flavor of recession risk right now we're facing, of some negative AI shock. It would partly come to the end of the data center boom, potentially through a negative wealth effect. Stock markets are capitalizing a lot of future productivity growth. That isn't looking like it's going to come to pass, the stock market could crash. If anything, I worry about risks also on the downside of some flavor of AI crash. I don't know how. The Anna Karenina theory tells you if you knew how, policymakers would get in front of it and stop it.

Anne Artrip:

We are at time, but I really want to end on a positive note. John, you already alluded to you are cautiously optimistic on the future. Can we get some estimation from you, Nick, as well, on your outlook? Maybe we can get you to rate from one to five, one being most pessimistic and five the most optimistic.

John Fernald:

Well, I told you I'm cautiously optimistic and it could be confused with pessimism, so I'll give you a three.

Anne Artrip:

A three. Can we get slightly more optimistic from Nick?

Nick Bloom:

I'll give you a four, which is striking. I have this well-known paper at Stanford, including Chad Jones I saw yesterday called Ideas Are Getting Harder to Find, showing that productivity growth has been slowing for the last 70 years. At this point, I think there's a fair chance it's going to start heading up, which is a pretty optimistic. I give it more than 50-50 chance AI is going to tilt that curve.

Anne Artrip:

Great. Well, thank you both very much for your time this morning and for your insightful perspectives, and thanks to everyone for listening in.

Nick Bloom:

Thank you.

John Fernald:

Thanks so much.

This document is a transcript generated by a third party and may inadvertently include errors or inaccuracies. The opinions expressed are those of the participants and do not necessarily represent the views of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.