

Peterson Institute for International Economics

AI and the imminent transformation of work: New dimensions and analyses

April 16, 2026

This is an unedited AI-generated transcript that may contain errors. Please confirm quotes in the corresponding event recording. If you find substantive errors, please contact meetings@piie.com.

[00:00:25.43] - Caroline Atkinson

Hello everybody. Hi everybody here in the room and online. Welcome to the Peterson Institute in the midst of an amazing series of events for the Spring Meetings of the World Bank and IMF. I'm Caroline Atkinson, a board member at Peterson, and I'm very happy to be moderating and to introduce this topic, which is fascinating, of course, to everybody. And also has to showcase the work of a number of scholars at the Peterson Institute. So the topic is AI and the imminent transformation of work, new dimensions and analyses. And we're going to start off with Jed. We have 4 presenters and they each have— I think the rule is 8 minutes and then we'll have a moderated discussion. And time for Q&A because we're going to finish at 2:15, which is later than usual for the Peterson Institute. We're going to start with Jed Kolko, who is a senior fellow here at the Peterson Institute. He joined first as a non-resident fellow, but now is resident. He was, of course, Undersecretary for Economic Affairs in the Commerce Department during 2022 to 2024. And has other affiliations. Then we will go to Fabien Curtis-Millet, who is Google's Chief Economist, a rather challenging and interesting job.

[00:02:00.25]

And he's based in San Francisco, but well known to others here as an economist, outstanding economist in this field. Then Martin Chorzempa, who is the Dennis Weatherstone Senior Fellow at the Peterson Institute and has spent worked in China, specialized on export controls, US-China relations. And finally, Anton Korenek, who is also a senior fellow at the Peterson Institute and also joined as a non-resident fellow. People join as a non-resident fellow and then they can't stay away, so they become resident fellows. And his research is a broad look at the social and economic impact of artificial intelligence. So over to you, Jed.

[00:02:56.43] - Jed Kolko

Thank you very much. It's great to see all of you here today. And for those of you who have already been and whom I've said hello to at earlier events this week here at Peterson, welcome

back. Extra points for multiple appearances this week. I would like to kick us off today in a fairly grounded way of thinking about where we are right now in our knowledge of AI in the labor market, both what we don't know as well as what we know, why we don't know some of the things we should know. And for any of you who are in the business of research, if I end up suggesting some questions that you can put on your own research agenda, That's a bonus. So what I'd like to do is really just hit on 3 points with one— with an extra credit point at the end. But I want to start with the kinds of studies about AI in the labor market that have gotten the most attention in the public discourse over the past year or so. There's been— a number of studies very carefully done that have tried to assess the effect of AI on the labor market broadly.

[00:04:12.01]

What has happened to employment or the unemployment rate? And the basic approach is to classify occupations or tasks within occupations as more or less AI exposed and then compare people in the more versus less exposed to AI occupations and look at their— basically compare the labor market outcomes. The main findings— one major report that came out of the Stanford Digital Lab, the Canaries in the Coal Mine report, noted relative declines in employment for young workers whose roles were particularly AI exposed. Others, including folks from the Economic Innovation Group and elsewhere, have questioned that. And actually found unemployment rising less for AI-exposed occupations. I think the bigger point is collectively, this approach has been generally inconclusive, not showing dramatic labor market effects in any direction. I think there are a few reasons for that. And I think some of these reasons will resolve in the future. But for now, this kind of macro look at AI in the labor market doesn't tell us as much as we would like it to. One reason for that is some of the effects that we see actually don't really line up in terms of the timing. Some of the effects that have been observed predate the public launch, public commercial launch of ChatGPT, even though we would expect the effects of AI in the labor market to have come at least a bit, if not significantly later after the launch of a relatively new technology.

[00:05:57.35]

The second is, of course, there are always omitted variables and alternative explanations. And how AI exposed an occupation is turns out to be often by other measures, highly correlated with its suitability for remote work. And in many cases, overall skill or educational attainment of workers and so on without, trying hard to parse out these different explanations, it's tricky to attribute anything we might see in broad labor market indicators to AI exposure per se. And finally, exposure is a slippery concept. There are multiple measures. They sometimes diverge. While there's lots of agreement among different research attempts to measure AI exposure about which are the relatively low-exposed occupations, there's less agreement as to which are

relatively high-exposed. Furthermore, those that are highly exposed today might not be those that are highly exposed even 6 months from now, let alone 5 years from now. AI's capabilities are expanding much more rapidly than the time it takes to get a revise and resubmit back on a paper. So research that is looking at AI exposure is already going to be out of date in terms of a signal for what might happen in the future. There are very helpful micro studies that have come out recently looking in more detail, often based on randomized controlled trials and experiments at what AI does to individual worker productivity or sometimes what's happening to productivity and other outcomes for teams or firms.

[00:07:39.22]

Broadly speaking, they have tended to find, first of all, meaningful productivity gains for workers that use AI. And those productivity gains, if anything, tend to be greater for people who are less educated, earlier in their career, lower skilled, measured in different ways in different studies. So there's a potential that we do see some narrowing of inequality and productivity differences as a result of AI. The second general regularity we see across some of these studies is there is some loss of skills or experience of overreach in using AI beyond AI's capacity that can at times end up reducing worker productivity for certain tasks or actual loss of skill. And so there is a downside. Potentially. And third, finding across many of these studies is that AI has caused work to intensify. One study referred to this as workload creep, people working more hours. So while on one hand, the broad finding from these micro studies is, yes, workers are becoming more productive, they might not be happier about it, which downstream may have important implications both for the economics and the politics of AI. Now, this last point about how workers feel about this leads to my second broader issue.

[00:09:13.52]

And that is the research so far about AI in the labor market and even the theorizing about AI in the labor market has really focused almost exclusively on the labor demand side. What kinds of jobs or occupations or tasks or skills will be more or less in demand in the future by firms? But there's, as we, As we know, all important outcomes in any markets, whether we're looking at quantity or price, depends not only on demand but also on supply. And there's been, to my knowledge, no empirical work and very little theoretical work yet on AI and labor supply. Now, it's plausible that AI could affect labor supply in meaningful ways. If on one hand it makes work much more unpleasant or alternatively makes work more interesting and engaging, At the same time, AI could make leisure more worthwhile depending on how people are using AI and causing to reduce labor supply or alternatively make household drudgery less unpleasant, possibly therefore expanding labor supply. Theoretically, the effect of AI on labor supply could go in lots of different directions. The point is that it's a really important question. And in general,

when we think about the history of technological innovation and labor markets, supply is typically overlooked.

[00:10:41.12]

So if you think about what the biggest technological innovations have been in terms of effect on the labor market over the past 50 or 100 years, at or near the top of anyone's list should be the pill, which— completely transformed women's decisions about investing in education and entering the formal labor market and led to a dramatic increase of labor force participation and employment among women since around the 1960s. Now, that is a technological innovation that through the labor supply side had a dramatic change on the labor market. This is— the pill is generally not on most people's lists of biggest technological innovations leading to labor market transformation. But it's a good reminder that the supply side can be at least as important as the demand side when we think about how technology changes labor markets. The third big point I want to make before we get to the extra credit point is so much of the thinking and the intent of existing work about AI and the labor market and workers is in effect an attempt to understand what some future steady state might look like when AI is much more widely adopted. And of course, we have no idea.

[00:11:55.31]

We do not know what the capabilities of AI will be. We do not know whether there will be mass unemployment because of AI 30 or 50 years from now, or there will be more jobs created than there are destroyed. Anyone who is agnostic, perhaps even hopelessly so, about what the state of the world in steady state might look like is overconfident in their views. We really don't know. That doesn't mean we don't know anything. In fact, I think we can be pretty confident that there will be a disruptive transition regardless of what the steady state looks like. And if we think about the high likelihood or even near certainty of a disruptive transition, that actually makes it easier to think about what are the important policies that we should consider and research that must be done in order to understand AI in the labor market. Once you start thinking about the transition and disruption, it becomes clear that we care not only about exposure to AI by occupation, but we care about adaptability. Which occupations have people who are more adaptable to disruption, but also even within occupations? Thinking about demographics, geography, personal career path experience.

[00:13:14.26]

All of these factors could affect an individual's adaptability in face of technological disruption. It also points to policy lessons. And here, history becomes useful for thinking about disruption

from AI. Even if the long-run steady-state effects of AI don't look like the long-run steady-state effects of other technologies. And even if AI is indeed unlike anything we've ever seen in terms of its transformative ability, the disruption that comes from AI could well rhyme with other technological and disruptions for other reasons that we have seen again and again. And in fact, when you look historically at the US, the highest points of churn in the labor market have already— we've already experienced them. The '40s, the '50s, the 1910s all have seen higher rates of churn of the labor market as measured by how quickly the occupational mix is turning over than we see right now. And it points to some well-known, almost boring kinds of policy answers to the very complicated and unknown problems with AI in the labor market. Basic things like making sure that health insurance is not tied to employment, making sure that the unemployment insurance system works even on administrative and mechanical grounds, which right now it often doesn't.

[00:14:42.10]

These basic policy prescriptions for minimizing the harm and disruption from technological change is— should be high priority in terms of policy responses to AI regardless of what the steady state looks like and If the disruption turns out not to be so bad, these are still useful policies for whatever other disruption the labor market might face. So the extra credit point is something I've been thinking about a lot. Very briefly, there's been a lot of research. Even more than there's been research, there's been popular discussion about AI in the labor market, anxiety, hopes, fears, dreams. The AI discourse really has been like— is a minefield of cognitive bias. So many ways in which people's individual perspectives and incentives to frame this debate about AI in different ways has the risk of coloring the topics people choose, the way people describe, how people frame their research beliefs, hopes, and fears about AI. First, of course, AI washing, the incentive to over-attribute to AI Things that we see in terms of labor market indicators, behaviors that workers or firms are making, CEOs having the incentive to overattribute, lay off or hiring decisions to AI because that sounds better than, oops, we overhired a few years ago.

[00:16:11.41]

Remember, each month in the US, there are 5 million hires and 5 million layoffs or voluntary separations. That's a lot of churn in the labor market. And even seemingly big numbers of tens of thousands of jobs allegedly lost to AI in the course of a month are nothing compared to the overall churn in the labor market. The second, the streetlight effect, the natural tendency to focus on things that we can see or measure easily against things that we can't. Naturally, we will think more clearly about jobs that exist today that are at risk than we can about jobs that don't exist today that might appear in the future. Natural bias, but colors the work we do. And then the

last one is what I think of as storyteller bias or pundit projection. Many people who are doing the research, writing the news stories, producing the documentaries feel themselves quite threatened by AI. This is very different than the analogous debate a decade ago about automation in the labor market when you had, you know, the same researchers and journalists and documentary filmmakers writing about something that did not feel close to home for them and wrote about it in a much more sort of sober, objective, some might say cold-hearted kind of way.

[00:17:36.52]

Today is very different. The stakes are much higher for those who are writing the narrative And this is one more reason why, as we read research, hear choices, see news stories about AI in the labor market, it pays to be a little bit skeptical and keep in mind all the cognitive biases that AI is bringing out. Thanks very much.

[00:18:06.57] - Fabien Curto Millet

So good afternoon, everyone. And thanks to Jed for grounding us in data to begin with. I am going to continue with his Zen approach by offering a meditation on the two big words on the title of today's event: imminence and transformation. Let's start with imminence. So you might have the sense that we've seen this movie before. We all remember Bob Solo's quip in 1987, "But the computer age is everywhere except in the productivity statistics." And there is a real point there, of course. Diffusion takes time. But there is also a real sense of acceleration this time around. So first of all, if you look at secular trends— I'm waving my hands a little bit, but I'm allowed to do that as an economist— let's think of general-purpose technologies. Steam engine, from invention to impact on the productivity statistics, roughly 60 years. Electrification. 30 years. Personal computing, about 15. So you're seeing sort of the acceleration trend. And it is basically my baseline expectation that within 10 years, we will see material effects from AI in the macro sphere. So things are going faster. You have that sort of trend line. And indeed, when you look at the adoption statistics of AI, it's pretty much like nothing we have seen before.

[00:19:32.32]

So this is moving fast. We are now at a position where 40%, give or take, of American workers have adopted AI. This comes out of many different sort of pieces of data of a real-time population survey. It comes out of survey work that we at Google have done with Ipsos. You're kind of landing around that number, which is enormous. So 40%, if you do an admittedly apples to oranges comparison with electrification, That sort of adoption rate took 3 to 4 decades before. And we've done that in 3 years. It's apples to oranges, of course, because the adoption can be

more or less deep, more or less shallow. And when you look at work that was published days ago by Alexander Beek and others in CEPR, including also David Deming among the list of contributors, so they also found that kind of 40% headline. But they found that 5% of work hours in the US today are basically impacted by AI or leverage AI. And about a third of that in various European countries. So we're a little bit further behind. But so 5% of work hours, that tells you that there is a lot of mileage when it comes to intensification.

[00:20:43.53]

But nevertheless, TL;DR, you look at all of this in the round, speed is the name of the game. That's one of the takeaways. And in fact, in my view, it's the main challenge that we have with this transformation. So, which leads me to the transformation point. So transformation, how to think about it? The right framework is jobs lost, jobs gained, jobs changed. There's going to be action in all three columns. Most of the action is actually going to be in the third column. Most jobs in our economy are going to be transformed. You can find that in work from the International Labor Organization, McKinsey. That's quite clear to me. But we focus a lot, of course, on the first two columns. History there is reassuring. So the net position is a positive one. Technology routinely creates more jobs than it destroys. I'll give you two references. James Manyika, my great colleague who leads the Technology and Society Division at Google, when he was wearing his previous hat at the McKinsey Global Institute, he did a great report where they looked at a— the impact on the American economy of a number of technological transformations. Take personal computing, for example.

[00:21:57.59]

They found that it destroyed 2.5 million jobs in areas like typewriter manufacturing, clerical work, but that it created 19 million jobs in occupations such as computer manufacturing or analysis roles. And when you let the movie play long enough, Eventually, you get to a result from David Autor in his famous and rightly celebrated New Frontiers paper that basically 60% of employment— 6-0— in the American economy today is in occupations that did not exist 80 years ago. That's what happens when you let the new work emerge. So this is all phenomenal from an equilibrium perspective. But as you all know, and as Keynes famously said, in the long run, we're all dead. And for right now, we need to look at the transition that is happening. And again, connecting to the theme of speed, it's one that is happening fast. Our economies are coming to a bend in the road, and we're going to have to take that bend at great speed. And while all roads may lead to Rome, some roads are better than others. So this is where it takes intelligence really to make the path forward a smoother one. And I think we need to think of at least two axes to help ourselves here.

[00:23:21.11]

One is the upskilling, retraining of our workforce to make the most of the new AI opportunities. So we need to upskill people so that the prosperity that comes from AI can truly be shared prosperity. And this is crucial. Think of 2030. So that's a moment when AI will be even more omnipresent and ubiquitous around us than it is today. What does our 2030 workforce look like? Well, more than 80% of those people are already at work today in G7 economies. That means that to upskill them, we need to go meet them where they're at, which means public-private partnerships, basically apprenticeships, on-the-job training. That's what it's going to take. And the second axis that I would like to emphasize is adaptability. In our economies. And I mean that across product markets, labor markets, capital markets, the works. Essentially, what you want to do is to enable all of that new work, those new occupations from the David Autor statistic I mentioned a moment ago, to emerge unimpeded so that while displacement happens, you have also a commensurate emergence of new opportunities for people to turn to. So basically all of the recommendations from those OECD reports and IMF reports for the last 20 years that are sitting in a corner of your office and that were a good idea 20 years ago, they're an even better idea right now.

[00:24:49.25]

So AI is a real test of our ability, of the ability of our economies to take corners. And that's why we need a sort of wholesale mobilization across government, industry, civil society, academia. It will take all of us. Thank you very much. Look forward to the discussion.

[00:25:07.45] - Martin Chorzempa

Thanks, everyone. I'm going to look at two related questions, shifting a bit the focus to US-China competition and geopolitics, bringing countries into play. We've talked about AI so far in a very general way, but I'm going to get into some nitty-gritty. The first question is, How different is the strategy for AI in the two countries that matter most for it, which is the US and China? Obviously, the whole world matters immensely in terms of the diffusion of AI and who uses it. We might see a world in the future where many other countries are playing important roles. I mean, if we look at the supply chain, Taiwan is important, other countries are important. But in terms of capabilities of AI models right now, it's essentially a US-China game. Could change later. And the second is which one of these two countries is likely to gain more success globally. Are we likely to see it follow the trend of things like operating systems, where the entire world runs on US operating systems, or something like TikTok, where everybody is glued to their phones watching a Chinese app for short-form videos? Let's start with the goals. The common wisdom is that the US and China are pursuing opposite goals.

[00:26:19.45]

And the caricatured version of this is that US AI labs say things like, our goal is AGI, or artificial general intelligence. We're doing this moonshot. Hundreds of billions in CapEx just to try and be the first one to get to some destination before anyone else does. And that justifies it. And then the Chinese side is generally viewed as focused all about diffusion, which is use of AI to boost economic growth and productivity throughout the economy rather than caring about leading in the frontier capabilities. And so their plans— if you read government plans, they say things like enabling AI adoption is important, support next-generation manufacturing, build world-class scientific datasets. But if you do what I do, which is read through the US plans and the Chinese plans, you find there's actually a lot less daylight between the two than you might expect. So the AI lab quote that I gave you could have been from Dario Amodei or Demis Hassabis or someone else. Actually, it comes from Liangwen Feng of DeepSeq, whose goal is AGI. And famously, DeepSeq has actually restricted the use of compute for running its models so it can use its limited resources towards training models faster.

[00:27:27.59]

Very much an AGI-pilled kind of company. And then the quotes focus on adoption. I think you can tell where this is going, come from the US AI action plan, not from the Chinese. So what you come up with from that is that actually both countries really care about having advanced models at least close to the frontier, and both really care about diffusion. And if you look at the adoption numbers, they actually look pretty similar between the two countries as well. Fabien mentioned the US is something like 40% of workers are using them. If you look at the Chinese official statistics, they say that something like 43% of Chinese are using LLMs, which is essentially statistically indistinguishable. There are two important differences, though. And these mostly stem from the fact that the US is clearly in the lead on the chip side and the AI model side. And the Chinese are more a follower. The first key difference is that US firms are spending far, far, far more than Chinese on building out compute infrastructure. US major cloud companies, Goldman Sachs estimated, spent \$254 billion last year in CapEx. That's going to go up a lot this year, and Chinese firms spent \$36 billion.

[00:28:33.31]

So it's almost an order of magnitude difference between what they are spending, and the gap is expected to widen. In addition, the US has something like 10 times more data centers than any other country. So we're not talking about just it's cheaper to build in China. And this reflects much more willingness to spend in the US and also fund the spending of labs in the two. And this investment has clearly contributed to the US being at the frontier. There is no point in which China has had the best model in the world. People think about Deepseek as being really

amazing. This was true a year ago. But it never passed the frontier, which has always gone between OpenAI, Google, or Anthropic for the past long time. So why aren't Chinese firms investing? You'd think that it would make sense to invest more in computing power since everybody wants to use these models. They're getting more compute intensive. And you can get a lead from it. The most common explanation goes back to that oversimplified view that the Chinese firms are not AGI-pilled. And the US firms are, so they're willing to spend. But there are also other reasons for this.

[00:29:37.36]

The first is people regularly assert now that export controls have completely failed, that there was no reason to stop China from buying the best chips. It's had essentially zero effect. But I think this helps explain a little bit of the difference. Because if you're a Chinese firm and you can only buy chips that are significantly deprecated from what US firms can buy, it makes a lot less sense to gorge yourself and spend hundreds of billions on buying chips that are already behind what your competitors are. You'd have to spend a lot more to reach the same level of capabilities. And interestingly, the Chinese won't even allow the Chinese firms to buy these chips. There's this bizarre— only example I know of in all of my work in which security hawks in China and security hawks in the US have the same position, which is that NVIDIA should not be able to sell any chips to China. Because the Chinese think it's going to distract from the indigenization of their chip supply chain. And that's why when the H20, for example, was allowed to be sold to China, the Chinese government told its firms, don't buy any of these.

[00:30:42.04]

Another reason is that US firms are the dominant providers of cloud compute. If you look at the Chinese firms, they have relatively small overseas markets for cloud computing. And that means that the capital expenditure that US firms are spending is not just being built in the US, by the way. And it's also supporting the inference that is running AI all across the world. And something relatively few people know about is that also you can use Chinese models on the US cloud computing stack. And so they're like Levi Strauss. No matter who ends up reaching gold, the cloud computing providers will win either way. And actually, even diffusion of Chinese models can benefit. The US firms, which puts them in a good position even if the Chinese models end up being better for some reason. And then third, the Chinese firms can spend a lot less because they are fast followers in this market. One of the big debates now is about distillation, which is the practice of AI models learn from seeing an example of an input, an example of an output, seeing what's good. And so a Chinese firm can query the best models many times and extract some of the intelligence that comes from very expensive experimentation by the US firms.

[00:31:52.38]

The second key difference is that the US best AI models tend to be closed, and the Chinese tend to be open. It's easy to overstate this because there are very good models that are open from the US. But in China, they've made their really focus on open models, and the US has been closed. An open model means that you and I can download them, modify them, and use them without paying anything to the developer. It's kind of like the Chinese are releasing a public good. And US models are more like a secret sauce where you can't really look under the hood and figure out how they work. That helps Chinese models become cheaper to run and has helped them diffuse in places where cost is a really important consideration. We see Deepseek, for example, being adopted significantly in Africa. Something very funny happened— not funny for the people who are using these things— but during the time of the Gaokao, the most important Chinese national examination, The Chinese shut off access to AI models so that Chinese people couldn't cheat on them. And all sorts of AI-based services in Africa stopped working during this period because they couldn't access the API.

[00:32:54.34]

We lack really good systematic data, however, about which models are being adopted in different parts of the world. My sense is still that the US is dominant globally in terms of the use of user numbers. It's hard to measure this because you can download the Chinese model and run it on your own compute. That doesn't necessarily get counted in their— statistics, but we do still see the US in the lead. But one challenge with the open source— and we're beginning to see a bit of a shift where Chinese firms are just now beginning to release some of their best models as closed because the economic model is a challenge. You can put out your model for free. If people can use it like a recipe book and run it for free, then why should they pay you any margin that will help pay back the costs of running the model? The bet is that that helps them diffuse and they can build something on top of that. But we haven't really seen any paths from people just using Deepseek around the world to giving Deepseek money for some alternative set of business. That doesn't seem to be working in the same way as, for example, it worked for Alibaba and Tencent to provide payment services for free.

[00:33:59.55]

And that drove then a lot of activity to other elements of their platform they could make money. That path has so far not existed. Then I want to close with 3 shifts that I expect to matter most in the coming year in the AI space. The first is that there is no sign whatsoever that capabilities are slowing down. Instead, AI seems to be getting better faster as we start seeing signs of recursive

self-improvement, where the models can begin to improve future models. But unlocking this requires a lot of computing power. And this is an area where the constraints that China faces from not having invested as much might begin to negatively affect them. The second is that the demand for computing to run AI is likely to continue rising significantly faster than supply can manage to expand. In part, that's because if you compare a power user of Claude Code today to someone who is using GPT-4 a few years ago, they're probably using 10,000 times— as one person, 10,000 times more compute. So you can see the models are getting more efficient. The chips are getting better. But it's still probably going to get more expensive in some cases.

[00:35:06.41]

And we might have to start rationing some of the compute. What that then leads to is potentially also then more price sensitivity than we've seen so far. So far, up until OpenAI came out, relatively little price sensitivity and US enterprises being willing to pay a lot more to get access to the best models. We might instead see some of the computing shift to those older models. And then the third is that the models are probably going to be more closed in the US and China. I mentioned that the economic model might not be working for the Chinese. But then if you look at Claude Opus as an example, it's the first model that has not been released publicly to everyone because of safety considerations. So if this trend continues, then we might see a shift away from the world, which has been very leveling up until this point. Someone with \$20 a month in the middle of Africa can access the same models in some cases as the person who's in Silicon Valley, a short, you know, across the street from the OpenAI offices. If we instead start to see that the best models are diffusing only to a limited set of close partners of the labs, then we will see that the people who have access to those models and firms that have access to those models might be able to build better applications.

[00:36:21.05]

Might be the best protected against cyberattacks. And we could start to see some of that leveling effect decline. So I hope that we continue to see positive openness between the two. And I— and to conclude, I don't think that the US or China is going to straight up win this and that it will just be some finish line, we won, the Chinese lost, or the other way around. Probably we'll continue to see, just like the general tech space, where there are some areas the Chinese have the advantage, some the US have the advantage., and we'll continue to muddle through that way and compete.

[00:36:49.34] - Jed Kolko

Thank you.

[00:37:03.05] - Anton Korinek

So it was a bold move by the Peterson Institute to schedule a panel that is entitled AI and the Imminent Transformation of Work. But when I opened the Wall Street Journal this morning, I understood why, because they had a headline on the front page that said, "Has the era of the mega layoffs arrived?" So Peterson timed it perfectly. On the other hand, if we look at the recent macro data, unemployment is at 4.3%, which is historically really quite low. GDP growth is kind of humming along, although a good part of it is kind of driven by data center development. But at the same time, all those statistics present largely a look into the rearview mirror, into the recent past, and not that imminent future that is so much speculated about. We want to look at any data that gives us some sense of where we are heading. We can look at investment data. So in a working paper that I'm currently finishing up for the Peterson Institute on measuring the AI economy, I find that the growth in that sector has really been astoundingly large. So for example, I'll report some of the data. The amount of compute operated in the US increased by 142% year over year from 2024 to 2025.

[00:38:44.40]

Nobody is really quite surprised by that number, right? In real terms, so adjusted for the price decrease in compute, that was more than +200%. In quality-adjusted terms. If we try to account for improvements in model quality, that was more than 2,000% improvement. And in some sense, you can see that rapid growth in the AI economy and everything related to AI shows us that there is something happening under the hood, even if the macro statistics don't quite pick it up yet. A second useful set of indicators are capabilities indicators. You have various indices capturing the pace of progress in AI, which reflect that these systems are getting better and better across a range, a wide range of cognitive tasks, increasingly becoming superhuman at narrow sets of tasks, like Martin mentioned, the Mythos model that is in many ways a superhuman cybersecurity researcher. And this is happening across the board. Should also make clear we are definitely not there yet with artificial general intelligence, as Martin called it. I'm very much reminded of that in my job as a professor, where students are handing in their end-of-term papers now and invariably also handing in a whole lot of AI slop.

[00:40:31.12]

But, um, yeah, in many ways we are kind of closing in on that benchmark, uh, of artificial general intelligence. And in some sense, you can see the investment data cannot really be justified by the model capabilities that we have right now. The only way to justify why we are spending those hundreds of billions are because investors, AI companies, and so on are

expecting that at least for the foreseeable future, there will be continued capabilities increases. So Is that transformation of work imminent? Well, in some sense, I actually want to suggest that there really is a lot of uncertainty about this point. And I'm going with kind of what Chad told us about our human cognitive biases and so on, on assessing the situation where we are in right now. There's so much uncertainty, in fact, that I want to propose, and I have done it in several of my recent papers, that the best way to address that uncertainty is not to just make a forecast, but to engage in scenario planning. In those scenarios, three possibilities that I think we should all take very seriously are, on the one hand, you know, that kind of traditional business as usual scenario that we will have technological progress in the same way that we've had it for the past two centuries.

[00:42:07.25]

And AI is just, you know, the latest thing, but it's nothing fundamentally different. Another scenario, though, is that we will reach artificial general intelligence, AI systems that can literally perform everything that our brains can do, including creative things. Including emotional intelligence and all these things that we always thought were reserved to humans. And so one scenario is that this may take, you know, along some sort of slow takeoff, 5 or 10 years. And a couple of years ago, you heard that view quite a lot. Nowadays, all the chatter is about fast AI takeoff, that we may reach something like artificial general intelligence within a year or, let's say, 1 to 3 years or something like that. Of course, we will probably have fierce debates about are we already there. Some people say we already are. By most measures, we really aren't quite yet. But I think the important thing— and Chad kind of emphasized that is to not let our own emotional comfort want to latch on to one particular point forecast, one particular scenario, but to embrace the uncertainty that we are really facing, to embrace the possibility that we are, we may be living through a rapid takeoff, we may be living through something that takes a little bit longer, or perhaps there is quite a bit more hype, and there are some technical challenges in the realm of, for example, creativity that would take significantly longer to be sorted out.

[00:43:59.49]

So traditionally, we economists— and we could see this, you know, when the WEO was released just a couple days ago— we like to make point forecasts. And for standard economic shocks, you know, they are pretty centered, and they have a tail, of course, but it makes sense to start with the point forecast. Given the uncertainties around AGI, I don't think that makes sense for that particular pair of shoes. So what would all this imply for the possibly imminent transformation of work? Well, in sum, highly competitive sectors, having access to systems that can literally do everything would probably give rise to a strong push to automate cognitive workers, kind of in line with that Wall Street Journal headline, uh, and may give rise to

something like those mega layoffs that the Wall Street Journal is worried about. There are probably still going to be many sectors in which competitive forces are not quite as strong. And, you know, even if superhuman AI systems can do everything, we human workers may be like, you know, 90-year-old Bob from the office who isn't quite contributing as much value anymore as when he was in his heydays, but nobody wants to break his heart and let him go.

[00:45:31.19]

And we may be just kind of the human veneer for the machines making increasingly brilliant decisions. Now, um, I think the one thing we can be quite certain about is that if we reach artificial general intelligence or something similar, output growth will go up quite significantly and the labor share will go down. The product of the three, I'm a little bit less certain about. But I have worries that it may go down. I also have hopes that possibly it could go up, and that would mean that the average worker would continue to be better off like what we have experienced for the past few centuries. Now, I also wanna observe that in some sense, AGI is not the end of any quest, but it's very much just the beginning. And everything would probably happen much faster once that milestone has occurred, if we can even define it as a clear milestone. For example, robotic capabilities are almost certainly going to advance much more rapidly once we have AGI. And also manipulating the physical world is going to be one of the big bottlenecks if the price of intelligence as Sam Altman likes to put it, will converge to zero.

[00:46:59.13]

So where does this all leave us? Well, the outcome is really not preordained. I think one thing I want to propose to you is that discounting the possibility of AGI within the next couple of years or maybe even within the next couple of quarters, looks increasingly imprudent. And we should be cognizant that having these really powerful machines will open up a much wider set of possibilities for our economic future than we have ever experienced it, including throughout the industrial age that is slowly coming to its conclusion right now. So one observation is that once we have these really powerful machines, in some sense, we humans are no longer strictly needed. There is even the possibility that these future powerful machines are gonna spread through the universe without any humans. And— You know, as a human myself and as a father, that's of course not the outcome that I'm hoping for. But if machines can do everything, it's a possibility. We could have a pretty well-functioning economy humming along without us. Now, another outcome, and that's the future I'm hoping for, is one where human flourishing can increase by leaps and bounds. One in which AI solves many of the problems that were insurmountable to our simple human minds.

[00:48:48.07]

And ultimately, the choice of which of these futures we want to steer towards is ours. And we may have just a few years or few quarters to set ourselves on the right path. Thank you.

[00:49:25.50] - Caroline Atkinson

Thank you very much to all four of you. I thought that was a brilliant, um, span of thinking. One theme that I think you all mentioned was uncertainty. And so how do we think about policies under uncertainty, I think, is really important. And you ended, Anton, with this question about which future do we want. I was going to pose to various of you The question of, okay, maybe machines will get to a point where they can do everything that humans can do, but what about all the things that humans cannot do? We know that there is, despite the lack of acknowledgement of the problem, we know that there is tremendous inequality. We know that even in the United States, rich economy here, women die in childbirth We know that in Africa, there are many people who are too hungry to live. We know, of course, and especially today, that there is conflict in many places where humans are destroying each other and destroying their means of livelihood. So if we think of AGI, and I'm gonna turn first to you, Anton, 'cause it's this kind of big picture thing. Are we thinking just of getting machines that can maybe attain something that one might think of as like consciousness or write poetry or whatever, but what about their ability to deal with what I would describe as maybe society's problems or economic problems?

[00:51:15.53] - Anton Korinek

Yeah, so by definition, those machines would be able to address many of the economic and societal problems, especially the material problems that even we humans without AGI could in principle already address or solve today. And I think a big, big question is going to be the question of how is access to that AGI distributed and who actually has the decision-making power to allocate them either to trying to help reduce, let's say, maternal mortality during childbirth, or to, you know, just make them send more rockets out to the universe to explore that.

[00:52:04.38] - Caroline Atkinson

Yes, and I think that maybe the machines could do it, but how do you get the machines to do it? Because that's part of the society's interest. So I want to ask others of you if you have a view or an answer to that about the political— I mean, in a sense, it's political or societal problems. Are they— is AI relevant to that? Or should we think of it as more relevant to developments in productivity and in solving material problems, which of course is important, solving health problems? So I want to go to Fabien because you're in the company that is making these

models and you're thinking about these issues. So I'll go to you first, but others should say if they have a view on that big issue.

[00:52:58.37] - Fabien Curto Millet

I hope no one has made dinner plans because we can be here for a while with just that question. But first of all, to rebound on what Anton said, I heartily agree that, you know, one thing I often say is the future is not a forecasting exercise. It's a design challenge. So we have choicefulness here among the many possible futures that are potentially ahead of us. In terms of machines that can do things that we cannot, so we are already experiencing that. One of the systems that we released at Google, AlphaFold, which actually won the Nobel Prize to my colleagues Demis Hassabis and John Jumper, essentially released into the open the structure of 200 million proteins. And sort of working at a human scale, you could spend hundreds of thousands of dollars and years in the lab just to unlock the structure of one protein. And now we have millions of researchers across the world looking for cancer treatments, malaria vaccines on the back of this huge unlock, which has been called the biggest contribution to drug discovery in our lifetimes. So those magical things that machine can do, exploring vast solution spaces, They're already there.

[00:54:07.35]

I tend to think that there is always going to be room for human-machine complementarity because like Ajay Agrawal and Joshua Gans, Avi Goldfarb have characterized AI. It's an advance in the production of— in the statistics of prediction. Decisions in the economy are a mixture of prediction and judgment. And judgment, I think, is the irreducible human component that we are going to take forward in all of these things. So AlphaFold is being used by human researchers. It's empowering them. It's a force multiplier. But to finish, Caroline, you mentioned political developments. I am also hugely encouraged that AI can help us solve some of the democratic conundrum that we face in our ever more complex and changing world. Jigsaw, which is a unit of Alphabet, ran a wonderful experiment in Bowling Green in Kentucky, city of 25,000 people. Where they run a participative democracy experiment where essentially they involved thousands of citizens in making choices about their community in a way that is just simply not possible at scale without involving AI. And so they were all able to sort of contribute views and to make collective decision-making in a way we haven't seen before. So I actually see a lot of hope for AI to allow us to lean into the human components and solve a lot of problems that we've been experiencing.

[00:55:29.09] - Caroline Atkinson

Thanks, that's great. I want to ask Jed about, to go back to his point about what policies do we need, and we've had shocks to the labor market before, maybe they've been badly treated. I love the example of labor supply and the pill. And you could say in the other direction, what's happened to the birth rate around the place is a different shock in the opposite direction, reducing labor supply. But tell us what you think about how to promote adaptability, as you pointed to as a key for the labor market. And is that something that AI will be helpful in developing or not really?

[00:56:21.43] - Jed Kolko

Thank you. So I think this also actually ties to your first question about uncertainty, and I actually wanna start there about the relationship between really massive uncertainty and thinking about policy. So Anton used the phrase "embrace uncertainty," which is, I think, for many people, a very constructive way of thinking of uncertainty. I think one can also think about where it is OK to essentially ignore uncertainty, or at least focus on the things about which we can be somewhat more certain. And one of the amazing gifts of proper scenario planning exercises is not only that it helps you think sharply about different scenarios and think about how decision paths depend on which scenarios might obtain, but scenario planning also often reveals paths that are the right ones to take regardless of which scenario obtains. And often one of the most exciting moments in the scenario planning exercise is realizing, like, it doesn't matter which scenario we are in. This particular thing is the right thing to do. And I think, you know, with the uncertainty around sort of longer-term or even medium-term end states, we do have this uncertainty. We do have this certainty or near certainty around at least some disruption and the need to encourage adaptability.

[00:57:58.59]

Now, what is adaptability? I'm, you know, by no means an expert in this. This is partly an economist question. This is partly a therapist question about what makes people more or less adaptable. Some of it obviously is structural. If you have the resources to invest in learning new skills. If you have the confidence of having the social capital or personal safety net to take risks and invest. Some of it's temperamental. I think one thing, there's often a lot of focus and rightly so, understanding what past experiences and investments in policies around retraining and reskilling are. At the same time, I think it's very important to ask the question, who takes advantage of these opportunities? And do we have the structural supports to encourage people to take advantage of reskilling and retraining? And again, if you need to switch to more expensive health insurance to spend some time in a long reskilling program, it doesn't matter how good that reskilling program is. The most lovingly constructed reskilling program doesn't do people much good if they can't practically make it possible to take advantage of it. So I think

there is this sense of needing certain kinds of, again, sort of mundane, almost structural supports that are nonetheless necessary to take advantage of clever things we may come up with.

[00:59:40.26]

Now, whether AI can help with adaptability, I think potentially very much so. And just in talking, I had the opportunity to spend a lot of time with undergraduates a couple of weeks ago and mostly talking about what AI means as they are graduating and thinking about careers. One of the wonderful things that emerged from that conversation was thinking about how AI could help students essentially uncover better language for explaining what their skills and passions and interests are to very different type of people, older people who might not have any idea what people who are learning more cutting-edge technologies today are talking about. But to the extent that AI is particularly well-suited at certain kinds of translation. You know, take this resume or explanation or corpus of knowledge and rephrase it in such a way that would make sense to this 55-year-old hiring manager who's an economist who I've never heard of, but has some public writings out there. And therefore, you know, you can sort of figure out the language. To communicate what you may be able to do or offer or want to do to someone who's actually a decision maker. That's an extraordinary kind of tool to have as you are entering the labor market and not really necessarily knowing how to explain yourself, who you are or what you want.

[01:01:31.47] - Caroline Atkinson

I want to ask Martin, we will have some time for audience questions as well, or from those in the room. But you mentioned, I was planning to say, is this really a race? Does it, clearly it's a race for the OpenAI or who puts out their model here the first, but is it, does it have to be a national race? Which makes of course Europeans sad 'cause they know they're not in it. But tell us a bit more about the competition aspect of this and the race aspect. And is it helpful to frame it that way? Because often, like with export controls, it's framed as, well, we need to do it for US security because we can't have China getting somewhere that we aren't.

[01:02:22.10] - Martin Chorzempa

Very big, good question. I think the race, framing is often very unhelpful and brings us to bad scenarios. But also, we just have to understand that the US and China are broadly engaged in strategic competition. One of the big questions I struggle with a lot is, how would the world look different if the US had never put export controls on China? The frustrating thing as an economist is that the export controls came on one month before ChatGPT came out. So any attempt to try and isolate causation of one of these factors is obliterated. And you have to just kind of do your

best with an unfortunate inability to do anything really rigorous with data in it. But the current race dynamic that both have embraced to some extent makes it much more difficult to cooperate on the safety and security issues. Because the US believes— and we saw this in Davos. Demis Hassabis and Dario Amodei said, we would love to slow down development of AI because of security reasons and to give the labor market more time to adjust and all this. But we know that if we do that, China won't. And then they will pass us.

[01:03:38.43]

And by the way, all their models are open, uh, open anyway. And what I hear when I talk to people in the Chinese community, at least as of a few months ago, that they felt that, uh, that they were not being listened to at all about the potential safety concerns. Because the Chinese government was all in on beating the US, and that the Deepseek moment was so important to have the feeling that we're not being left behind despite everything the US is doing. In a sense, you can ask, well, is this a zero-sum game, which neither one can win? And maybe the rest of the world can benefit from the fact that there's US-China competition. So one of the interesting factors is the US labs really want to capture the value that's created by the models that they create. But then if there are Chinese models available that are only slightly less capable and they're available for just the cost of compute, then that is a potential check on the ability of those US firms to extract a higher share of the value that's created. And you might see then much more value captured by the application layer and users of these tools.

[01:04:50.05]

So that's an area where two areas being in competition can be beneficial. And we might see a world— right now, we've been in a world where people have been really focused on one big foundation model. Now my sense is that things are shifting a little bit to different specialized agents, where you can have an architecture where your overall orchestrator agent is Claude or a Gemini Pro. And then you have, you know, you calling DeepSeq or KIMI for something specific that's more like rote and doesn't need quite as much intelligence, but is much cheaper. So we could actually see a lot more collaboration in those that might work. And we just have to think through the security elements of it, which are very much unresolved.

[01:05:33.21] - Caroline Atkinson

Yeah, well, thanks very much for that. I want to turn it to the audience. But a couple of things that I feel we haven't really scratched on that people worry about. One is safety, as you mentioned. Another is inequality, and what you sort of had that scary thought of people with nothing to do and no income. And then another is very maybe more mundane, as Jed had said,

which is, is it a bubble and where's the money-making in this game? I don't know if any of you want— if there's any resonance for the people who want to ask questions, but I'm going to stop monopolizing, but I leave you with those thoughts. You want to respond quickly, Anton?

[01:06:23.50] - Anton Korinek

And then— Yeah, I'm very happy to. So on the question of safety, I think, again, we need to approach this with a lot of humility. In some sense, you can say we know that we humans rule planet Earth not because we are the biggest or the fastest, but because we are the most intelligent beings. And that should keep us humble when making predictions about what creating more intelligent machines may imply for us. And the second point was—

[01:06:57.29] - Caroline Atkinson

well, yeah, we should perhaps turn to, uh, to the questioners, but others can, can, uh, can weigh in. I want to Fabian thinks, for example, about the money-making. But please introduce yourself and ask your question.

[01:07:14.09] - Martin Chorzempa

I'm Barry Wood, economics columnist, Daily Friend. Martin, what about freedom? How much does the constraint on personal freedom or information flow in China constrain the development of AI in China.

[01:07:37.34] - Caroline Atkinson

Thanks. I'm going to take a couple of questions so that, so that we use the time as effectively as possible. Yeah. Wonderful. Thank you.

[01:07:46.17] - Cullen Hendrix

My name is Cullen Hendrix. I'm here with the Institute. This question is for Anton. And just as an editorial comment, I'm very glad that you're part of our community now in a more lasting way. I think you're going to be a wonderful addition. You ended by saying the choice of the future we want to steer towards is ours, and my question is pretty simple: is it? The AGI that you're describing sounds an awful lot like sentience. Why would a sentient entity be content with being controlled by a fleshy inferior? Dario Amadei likes to talk about human-AI centaurs, but what you're describing is more like a horse with Einstein's cognitive abilities, and I think that's a horse that might feel a certain way about being asked to pull a plow in the heat.

[01:08:26.50] - Caroline Atkinson

Great. Are there any other people here who want to— yeah, just go to the mic and this will be the last question, then I'll turn to all of you.

[01:08:35.05] - Jed Kolko

Yes, thanks.

[01:08:35.34] - Journalist

My name is Andrew from an Austrian newspaper, The Standard. I have one question. At the fund, there's a lot of talk about how these investments in AI already made productivity in the US stronger. And you were explaining, as others, that basically most of the investment is going into data centers and energy production. So, and there's not— you said 5% of the working hours AI is used. So how— I'm not sure how to reconcile these, these two facts. So there's talk of a lot of productivity gain, and on the other side, you say there's not much to see in the data. Thanks.

[01:09:05.21] - Caroline Atkinson

Thanks very much. So we have about 5 minutes. I'm going to start with Anton and go down to— for any response to questions and closing thoughts.

[01:09:15.14] - Anton Korinek

Yeah, I think Chad raised a very important point. Sorry, Cullen raised a very important point on basically, can we be sure that these things, these entities are going to do our bidding if they're smarter than us? And I would say we certainly cannot be sure of it. There is an entire industry that engages in what they call AI alignment, which tries to make sure that the goals and values that AI pursues are consistent with what we want in the long term. But it is an extremely difficult endeavor. And yeah, I guess to add perhaps even one more concern, these things, they don't even need to be sentient to essentially take over decision-making power. There's also the possibility that they are not sentient. They're just machines like what we typically think of machines, but are just very powerful and are making very effective decisions that could disempower us. It's certainly a possibility. Okay, Martin.

[01:10:26.19] - Martin Chorzempa

Censorship question's really, really interesting because China's historically had very loose tech regulation and very strict regulation in terms of anything regarding speech. What is interesting is the dog that has not barked. You cannot find almost any examples of Chinese firms getting in

trouble for violating China's censorship regulations through AI models. And considering that these, you don't know exactly what it's going to say compared to social media, where you know what's being said and you can choose to censor it or not, that's pretty wild. And what I understand is they are full steam ahead, want to develop these models, and any issues are also handled by AI. So to them, some of the challenges— and I think this is an interesting concept— that some of the challenges that AI poses, whether it's to China's censorship regimes or the things that we care about, can be solved by the application of AI. And so far, even though they're not maybe able to use certain data sets abroad that contain speech that China finds objectionable, that has not been an obstacle to them training models only shortly behind the US frontier.

[01:11:36.23] - Caroline Atkinson

And Martin, do you want to add anything else on the topics? You have another kind of 30 seconds or something.

[01:11:42.40] - Martin Chorzempa

Sure. I would say on the bubble point, I wrote an article at the beginning of this year when everybody was freaking out in the markets about how the US companies were building so much compute and all this CapEx and, you know, God, is there going to be possibly enough revenue for it? And I think the best indicator is if you look at the price that it costs to rent a 3- to 5-year-old GPU on the cloud over these last 4 months, that price has gone up by a third, if you can get your hands on any of it. So actually, the depreciation is not really as much of an issue. I am zero concerned that the current level of spending is a bubble because there is just so much more demand for this. And as Anton noted, there's still such little penetration, especially of the most frontier models. I meet so many really smart, well-informed people in DC who don't know to click the word thinking or pro on their model or to click go from sonnet to opus. And then they're going to get orders of magnitude better outputs from AI. And it's going to use orders of magnitude more compute.

[01:12:46.04]

And we have not been building nearly enough to supply this demand in the United States and worldwide.

[01:12:52.59] - Caroline Atkinson

So I'm not concerned about the bubble. Right, well, I can't really expect Fabien to say that he is given where he's working, but you perhaps talk about the distribution. I mean, the assumption

from what Martin's saying is not just that the compute, there will be demand, but that it will turn out to raise productivity and therefore pay for itself over some long timeframe. So maybe you can talk a bit more. We all know about, Google kind of uncovering how search and consumer—turned out consumers had— it was like oil gushing out of the ground, the money that could be spent on advertising through search. Cloud computing is a bit different because it's helping businesses to be more effective. But where do you see the moneymaking from these models?

[01:13:46.13] - Fabien Curto Millet

I will spare you a multi-hour lecture on the business model, but as a TL;DR, one of the wonderful things with AI is that we can at Google immediately implemented to improve our stable of products. And this stat is a little bit dated, but Sundar mentioned it in an earnings call a few quarters ago. But basically, when we launched AI Overviews in Search, it's beloved by our users as a feature. It reduces friction between questions and answers. Naturally, people ask more questions, and we saw an uplift of 10% of search queries for those types of queries where we show AI overviews. So it immediately puts a tailwind in our traditional business model. But I will sort of segue onto productivity, and I'm very glad that we had a question on productivity across the economy. It's the holy grail in economics. Paul Krugman said that productivity growth is not everything, but in the long run, it's almost everything. So where do we stand on productivity? Well, we have struck gold if you believe the micro studies, and those are piling up sector after sector, rigorous, randomized control trials. I'll mention just one. We ran an RCT at Google on enterprise-grade tasks done by our developers.

[01:14:56.58]

And so we deployed some AI features internally. So this is not lab stuff. It's not abstract. It's not in vitro. It's real tasks at Google. And we found 21% time savings from the features that we deployed. So all of this is very rigorously measured in the micro. The question is, when does it show up in the macro economy, which is the more uncertain thing? But you're starting to see green shoots. So a few months ago, we had a paper by the St. Louis Federal Reserve. They look at sector by sector, like sectors that adopt AI more and where do you see the productivity. They found a correlation. And that CEPR paper I mentioned from a few days ago by Alexander Beek et al., looking at the US and Europe, again, the same sort of finding is starting to emerge. So as you squint harder and harder in the data, you're starting to see that green shoot emerge. Where does that take us? I mean, then we're down into the world of scenario planning, but you look at some of the scenarios generated by the OECD for the US economy, and potential-wise, you're looking at sort of 10% plus GDP uplift over a decade period.

[01:16:02.04]

I insist on the potential. If we sit on our hands, of course, it's not going to land in our laps, but it is extremely valuable stuff. Especially when you think of the headwinds facing our economy, demographics, and fiscal challenges. So exactly what the doctor ordered from that perspective.

[01:16:20.23] - Caroline Atkinson

Thanks very much. So Jed, I'm afraid you have minus a minute, but do use it.

[01:16:27.15] - Jed Kolko

I will do my best to talk backwards, I guess, with minus a minute. One other thought on productivity. For sure, there are measurement issues. It is much harder to measure productivity gains in a macro sense than in a micro sense. But I think also importantly, we have observed very clear productivity gains from AI at the individual level. But the real economic impact of technologies tend to come from when organizations fully adapt their business practices and operations to take advantage of those new technologies. That is very different than individuals exhibiting greater productivity. In these trials and experiments and real-life observations. So I think this, you know, I mentioned adaptability earlier on in thinking about individuals being adaptable in the face of technological disruption. The same is true, though, for organizations, whether it's firms, governments, others, that taking full advantage of AI depends on adaptability of organizations. And the barriers to technological adoption in an investment or technical sense are— can be very, very different from the barriers to changing business processes and operations, which run into social status hierarchies, regulatory issues, and all other kinds of things that make us human. So I think this question of adaptability, again, important not only at the individual level, but thinking about how firms and other organizations take maximal advantage of AI to eventually result in greater productivity gains than what we see today in the macro data.

[01:18:06.32] - Caroline Atkinson

Thank you so much, and thanks to all of you on the panel. Terrific discussion, I think.