Peterson Institute for International Economics

AI and economic transformation: Growth, adoption, and who captures the gains

October 6, 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.

[00:00:24.12] - Anna Stansbury

Good morning, everyone. Welcome to the Peterson Institute for International Economics. Today, we're going to be hosting a panel on AI and economic transformation, growth, adoption, and who captures the gains. The reason we convened this panel is to think not only about how AI is advancing, reshaping productivity, economic growth, jobs, and living standards from the perspective of one country, but to extend that across different countries. AI gains are unlikely to be distributed evenly across countries, and how quickly the technology is developed, how quickly economies adopt it, And whether they have the infrastructure, the skills, the investment, and the international access to benefit is going to shape AI's impact worldwide. So I'm moderating a discussion today. I'm Anna Stansbury. I'm a nonresident senior fellow here at the Peterson Institute and also an assistant professor at MIT. And we have today 3 people who are absolute must-listens on the topic of AI, economic growth, and who captures the gains. So I'm delighted to welcome them here today. First, joining us for the discussion, we will have Martin Chorzempa. Martin is Dennis Weatherston Senior Fellow at the Peterson Institute. He's been at the Peterson Institute since 2017 and is one of the leading experts on financial technology, digital currency, technology and national security issues like export controls and foreign investment screening, as well as bringing particular insight to issues of technology and China and the US-China relationship.

[00:01:59.19]

His 2022 book, The Cashless Revolution: China's Reinvention of Money, was named one of the best economics books of the year by the FT. And earlier in his career, Martin worked in China at the China Finance 40 Forum in Beijing and was also a Luce Scholar at Peking University. Martin, we're delighted to have you with us.

[00:02:19.10] - Martin Chorzempa

Thanks for having me. I'm going to share some slides.

[00:02:23.04] - Anna Stansbury

Before you present, Martin, I'm going to bring Filiz, introduce Anton and Filiz, and then we'll do the presentations. So we're also going to have with us Anton Korinek, who is Head of Transformative AI Economic Studies at the Anthropic Institute, and a non-resident senior fellow at the Peterson Institute with us. He's also currently on leave as Professor of Economics and Faculty Director of the Economics of Transformative AI Initiative at the University of Virginia. Anton's research analyzes how transformative advances in AI affect the economy and society, AI's implications for economic growth, labor markets, inequality, market concentration and how to steer progress in AI to lead to shared prosperity. Anton was on Time's 2025 Time 100 AI list of the most influential people in AI and serves on the Anthropic Economic Advisory Council, was a member of a G7 high-level panel of experts on AI, and has testified on AI's impact to the US Senate. So we're delighted to have you with us as well, Anton.

[00:03:20.00] - Anton Korinek

Great to be on this webinar with you.

[00:03:22.15] - Anna Stansbury

And last but absolutely not least, we will have Filiz Unsal, who is Deputy Director of the Policy and Research Branch in the Economics Department of the OECD. She is a central leadership role in the OECD's analysis on a wide range of macroeconomic and structural policy issues, as well as in coordinating some of the OECD's major flagship publications like the Economic Outlook and Going for Growth. Earlier in her career at the OECD, she was head of the Structural Policy and Research Division, and prior to joining the OECD, she worked at the IMF where she led key policy and research initiatives as well as country missions. And Felice's economics research has focused on macroeconomics and international trade and finance and has been widely published and informed policymaking across both advanced and emerging economies. So thank you for joining us, Felice.

[00:04:05.21] - Filiz Unsal

Thank you so much. Thank you. It's a pleasure to be joining you.

[00:04:08.23] - Anna Stansbury

Fantastic. So the run of show today will be that Martin, Anton, and then Felice are going to each present some remarks to us, and then we will have a discussion between the 4 of us, and we will wrap up on the hour at 10:00 AM Eastern time. So Martin, please take it away.

[00:04:23.18] - Martin Chorzempa

Great, thank you. And I think I was trying to share early to make sure that nothing went out of date before I spoke, because it's moving so quickly. I'm going to talk about China and what I call formidable fast followers and China's challenge to US AI leadership. This is based on a forthcoming Peterson Institute policy brief, also coming out with the Hoover Institution, which generally generously supported this project. So, I'm going to look at China's ambition and strategy first. What does China want out of AI? And then I'll get into the competition with the United States and how they're doing it. So the ambition, first of all, for China is to build an indigenous, self-reliant AI ecosystem. And that includes everything all the way down from the chip layer to cloud computing to the models and applications. They don't want to depend at all on the United States, in part because, of course, the US has put export controls on, on China and realize the fears of being cut off. But also, this is a longstanding goal for China to not be reliant on others. The other 2 goals are domestic diffusion and international diffusion.

[00:05:33.07]

If you look at the mentions of AI in China's 5-year plans, which are some of the most important documents giving us an idea of their strategy, the mentions of AI in the plan goes from about 5 in 2021 to 52 in 2026. So you can see that AI now permeates China's planning in many different areas. And in terms of domestic diffusion, they have this AI+ plan as well, which is designed to implement AI where useful all throughout the economy and boost productivity. And then they also want international diffusion. And this is really where things get the most interesting in terms of competition with the United States, both in the United States and globally. for whose AI systems at many different layers are going to get adopted, whose values are going to be implemented through that. In terms of the tools that they're using, it's quite a few. As usual, China has, you know, many, many plans and many, many different policy instruments that they use. One of them is that they have these government guidance funds, which mingle together public and private money to invest in AI companies and also chip suppliers. So most of the recent rounds of fundraising for Chinese AI labs have at least some role for one of these government guidance funds that's giving them the capital.

[00:06:48.17]

The second is input cost subsidies. Obviously, one of the most important inputs is computing. And of course, to make that computing work, you need a lot of electricity. If you are using domestic chips, you can sometimes get 50% off of your electricity bill every single month thanks to government subsidies. And that's designed to make up for the fact that Chinese chips are

significantly less efficient. Power efficient than Western chips designed by companies like Nvidia. Then also coordination. So they're using government procurement and trying to influence global governance through especially the United Nations system, trying to get Chinese standards adopted and convening the world at the World AI Conference and now the World AI Cooperation Organization, trying to get global buy-in to China's idea of You know, we are the open player. AI should be an open public good, not something that U.S. companies are able to hold on to and, and control. That's how they're trying to position themselves. And then, of course, protectionism. If you try to open up the Anthropic Claude website or ChatGPT in China, you will be blocked by the Great Firewall. But they try to have strong domestic competition behind that protectionism layer.

[00:08:01.12]

And then, of course, support for open weights, which is a different way. We'll get into detail about what open weights means, but it's a bit of a different way than the leading AI labs, the way they release those models. Yeah. Well, how have the results been so far? So, so far, interestingly, on the self-reliance point, the chip production in China is improving. They're able to produce many more chips than before. Their chips are becoming higher quality than before, but they still remain far behind the United States. The DeepSeek founder, Liang Wenfeng, in some leaked remarks that were reported on recently, said that the gap is fourfold plus 2 years, meaning If you get 4 Chinese chips together, you have a chip that's equivalent, you know, a super chip that's equivalent to a U.S. chip from, from 2 years ago. And supply is still heavily, heavily constrained. And now if we look at the capability side, the Chinese models, it depends on which benchmark you look at. It's more of an art than a science at this point, but are something like on average 6 to 8 months behind the United States, certainly less of a gap than people were in the U.S. were hoping would result from export controls that stopped them from accessing these advanced chips.

[00:09:11.05]

And we can get into the reasons why that is. Now, one of the things that narratives about China that's really common is that China is all about diffusion and the US is all about some moonshot race to artificial general intelligence, and China doesn't care about AGI. And one of the things I always put in my talks is every single major Chinese lab talks regularly about artificial general intelligence, and all of the US plans that I have seen related to AI also mention diffusion. But people might say, well, China has all these policy instruments and these plans, so therefore they're much better at diffusing AI. And a lot of people go to China and see these anecdotes about, wow, there are robots everywhere. China is adopting so much more. Well, if you look at the data, you find there is no evidence whatsoever supporting greater AI diffusion overall across

the Chinese economy than in the United States. If you look at the end of last year, which is the last time we have data, the official Chinese sources estimate about 600 million people in China use AI. That's 43%. of the population. If you look at Pew in the United States, something like 49% of adults are using AI.

[00:10:16.28]

So the US number is larger. Of course, we don't have the same denominator. It's a little bit tricky, but I actually think that the gap would be larger if you had the same denominator. And then in terms of firms, we have some surveys of digitally active small firms in China, and you might be surprised that only a quarter of small firms in China have adopted AI, whereas people just generally think Wow, everyone in China has adopted AI. Not actually true. And then in the US, it's maybe a little bit lower, pretty similar if you look at census. But of course, again, different denominators here. The firms involved in this survey are actually already digitally active small firms, so they're more likely to have already adopted AI. So no gap on domestic diffusion, pretty similar. It's something that's very different than the narratives out there. In terms of international diffusion, this is also very hard to measure. If you look at the revenue side, US companies are just dwarfing Chinese firms. They are making so much more revenue all around the world than Chinese labs, in part because companies around the world have historically been able to use Chinese AI models without paying their developers anything.

[00:11:25.02]

That's what— downstream of the open weights discussion. Chinese models are far more downloaded than US models. That's in part because they post them online like a recipe. You can download it and then run it on your own computing power. You can run it on a cloud provider, but it's really hard to tell what exactly the downloads mean because a lot of the downloads are the really small models and might not be used. And if you and I download one of the models once and then never use it, that's counted the same as a big application developer that downloads the model and then serves it to, you know, tunes it and then serves it to maybe hundreds of millions of users. That's all going to count as one download. So it's hard to interpret this. But generally, my sense is US is leading on international diffusion. And by the way, any, any consumer pretty much anywhere in the world can use US AI models for free. So the Chinese pricing advantage is less than people often assume. So one of the things to think about is how much of the AI model usage is going to go to US model providers versus Chinese.

[00:12:25.03]

Generally, people assume that the Chinese model producers are much cheaper, in part because they actually give the models away for free. Sort of. And that comes down to, first, the frontier gap and what tasks that is available. So if the US continues to be ahead of China, what are the tasks that only US frontier models can reliably perform? Those are going to be areas that sustain pricing power. And that, of course, is this strong competitive imperative to keep going very quickly in AI. The second is where the Chinese models can match, where you could use a Chinese model or a US model for this task. how much is this market going to commoditize and lead to effectively no economic return on providing that activity? If you look at these non— that gets into how much are non-price and capability factors going to matter. If we look at the trust and packaging package deals, you know, you get a subscription to Anthropic or ChatGPT, you might be able to use all these different models, including the downstream models, and you might just have picked one as an enterprise. That favors generally US providers that have a major lead and more trust.

[00:13:31.04]

But then in terms of control of the models, data, and access, this favors Chinese open-weight models because the access to those models cannot be revoked, unlike US closed models, and they can be fine-tuned and also run without sending any data to the model developer. And that means some companies that are particularly concerned about the security of their data might not want to send them. So let's look at what's happening in the market in terms of competition. Really interestingly, for recently, the total spending, at least measured by Ramp— these are imperfect data looking at kind of businesses and credit cards, what they— what the businesses spend— showing some decline in spending. And it attributes the decline in the spending entirely to price competition between OpenAI and Anthropic. and attributes almost none of this to competitive pressure from Chinese models because Chinese model spend is something like less than 5% of the total. But you watch token volumes continue to expand. So we continue to see this trend of more capabilities, more tokens, not necessarily more money. And the really interesting thing that has happened quite recently, actually, even since I wrote the initial paper, is that if you look at the frontier of cost to capabilities, so this chart is showing you the x-axis is the cost per task, and the left side is the artificial analysis intelligence index.

[00:14:50.19]

It's kind of a composite of different measures of how intelligent the model is. And the dotted line there is the frontier where you have, you know, what is the most intelligent model at any cost per task? And what you find is that except for one area in the middle with MEMO, which almost nobody has ever heard of and pretty much nobody uses, the US models right now dominate the cost capabilities frontier. At any level of capabilities, you're likely to get a cheaper model from

the United States than you are from China. And I think in part that's because, I mean, this is a major narrative violation. Everybody assumes Chinese AI is much cheaper. But in fact, US has— US companies have the best computing power, the most efficient compute, and they have a lot of that compute that they've used to optimize the inference side. So it actually shouldn't be that surprising that US companies can be really competitive here. China's main advantage may be shifting instead of cost to this control side, the fact that users can inspect and change the model, fine-tune it for their own purposes, and it can't be taken away from them.

[00:16:00.05]

But it's really unclear at this point how much will move towards open weight. So I'll conclude here so I don't go over. China's fast followers, they do not need to compete with U.S. companies all the way at the frontier in order to shape global AI and the economics of U.S. AI. There are some implications of Chinese competition is that, as we've seen now with this intense competition, that will tend to transfer the gains of AI to the users instead of the developers of AI. That lower pricing power does create a risk for the ability to recover the enormous costs involved in developing frontier AI models. And I want to bring distillation into the discussion a little bit as well. That's where, you know, Chinese companies have been credibly accused of trying to extract the capabilities that were hard-won in top US AI models and getting it with a lot less computing power. If US labs fear that their models that are publicly released are going to be quickly distilled and those capabilities kind of taken away nearly for free, by Chinese companies, they might do less public releases, and it will internal— it will incentivize them to keep their best models internal.

[00:17:10.08]

This is a challenge to their economics, of course, because they really need to diffuse those models around the world and make money from serving them, because these models are very quickly depreciating kind of assets. But it also could lead them to branch more into applications themselves and use those really powerful internal models to release products rather than just releasing the model to anyone. And that then creates really interesting, complicated concerns around competition. You know, if they have this super special internal model that no one else can use, they might be able to beat companies in many other areas. But one interesting relationship to the slowdown debate of whether AI needs to slow down and whether China would need to agree to slow down its AI development if the US slowed down in order to focus more on safety. Well, if distillation is a major source of gains for China, then US slowing down would also tend to mechanically slow down, at least to some extent, the Chinese labs, which would not be able to benefit from distilling those models. So my recommendations to the US

side is to compete across all across the cost capabilities curve, something that's happening right now, and then restore faith in the reliability of access.

[00:18:17.17]

I think there's— we're still dealing with the fallout of The fact that Mythos and Fable were forced to be taken down in a very abrupt way by the US government. And we need our partners to know that they can rely on our models. Otherwise, they'll move to Chinese open weight models. And then we need much more rigorous safety testing of adoption of Chinese models. I think too many of them are being adopted now without really understanding what the risks are. But if we don't have concrete risks and we're just kind of waving our hands about potential risks, there's no smoking gun, then that's not going to be credible to firms or to other governments. Thank you.

[00:18:57.13] - Anna Stansbury

Thank you very much, Martin. That was fantastic. And I'm hoping we can pick up on some of those questions in the discussion, including whether the gains accrue to the frontier model developers or the users, and under what assumptions that's going to happen. Now we're going to hand over to Anton. So Anton, take it away.

[00:19:13.02] - Anton Korinek

Thank you very much. And indeed, it was a fascinating presentation, Martin. Thank you very much. So I will pivot a little bit to the economics, and I will give a presentation based on my recent work at the Anthropic Institute on economic scenarios for transformative AI. And if there's like one takeaway line, then it is the one we've put here, which is prepare for a range. Yeah. not for a point forecast. So why do we say this? There is an enormous amount of debate about how AI will evolve in the coming years. Some people view it as a normal technology, but in Silicon Valley, the talk is more that it is going to be a very transformative technology. There are now conversations that AI may reach the point where it's smart enough to develop its own successor systems. This is called recursive self-improvement within the next couple of years. And if something like this turns true, then we should expect the economic impacts of AI to be much greater than what we can see in today's data. So in this model that we developed, we attempt to capture the wide range of uncertainty in a small number of measurable parameters.

[00:20:38.29]

The first one is how capable these AI systems will actually become, meaning how many of the tasks in an economy it could possibly reach. The second one, how fast the adoption proceeds.

This is something that Martin paid a lot of attention to. Basically, if we have really powerful AI systems, but they are not diffused through the economy, we shouldn't expect significant gains from that. The third parameter is the magnitude of gains. So how much more efficient AI systems allow us to do a given task. And then a fourth parameter that is very important for the labor market impacts is how much AI systems automate work versus augmenting work. There's a fifth parameter not listed here, which has to do with how quickly the job market can adjust. I will come to that. in a couple slides. But based on this set of parameters, what we do is we develop a framework that lays out paths for GDP growth, for wages, for job numbers, and unemployment between now and 2030. We focus on 3 specific illustrative scenarios. We intentionally don't assign probabilities to them, Because we believe that at this point, the range of uncertainty and the degree of unpredictability is too large to really make informed point forecasts.

[00:22:20.19]

And another point I want to emphasize is, so in today's data, we can't see very large impacts of AI yet. We have at best modest GDP effects in 2026. We have a very healthy or reasonably healthy job market, and that is consistent with 2 possible readings. So first of all, it is consistent with a world in which AI may actually not make that big of a difference for economic growth and jobs and so on. But the second reason— reading is that it may actually have quite extreme impacts, but that we are just at the very beginning of an exponential. So distinguishing between these two over the coming years is important because they call for very different policies. So here on this slide, I want to show you our GDP forecasts in the 3 scenarios that we consider. The scenarios range from modest to extreme, and you can see in the modest scenario here on the left-hand panel, overall GDP is not even 2% above the path where it would be without AI by 2030. On the other hand, in the extreme scenario, GDP is more than 32% above its no AI path. So that means it would really have very significant impacts on the economy.

[00:24:03.05]

In this extreme scenario, the growth rate in 2030 is approaching 15%. Now let's turn to what this implies for the distribution of the gains. between labor and capital. So, I'm depicting here the labor share of national income in the different scenarios. Without AI, we pin this to a rough estimate of what it has been in recent years, about 60%. In the modest scenarios, almost no impact on the labor share. But if you look at the extreme scenario, the labor share has gone down significantly. So, I'll combine this again with a number from the previous slide. Overall GDP has gone up by 32%, but labor income is almost flat, whereas most of the gains in this extreme scenario go to capital. So I guess the positive part of that is that there are very large aggregate gains, but the question is how do we make sure that everybody in the economy can

share in some of these gains. So let me put a specific spotlight on what happens to jobs in our different scenarios. Right now, 62% of the economy roughly are knowledge workers, and the remaining 38% are in non-knowledge work. That means physical jobs or interpersonal jobs.

[00:25:45.29]

And if AI develops rapidly, it will be most capable at performing cognitive jobs and performing the knowledge worker jobs. So that means across the scenarios, but in particular in the most extreme scenario, the economy needs to reallocate jobs from the pool of knowledge work to the other sectors of the economy. In the other sectors, there's actually going to be a significant shortage of jobs, of workers, because the economy is growing so fast and it really needs workers in these other sectors. So on the right-hand side of this table, I'm listing what this all would imply for unemployment, for the number of knowledge work jobs, and for wages. So, we can see in both the modest and substantial scenarios, unemployment numbers would not move out of the range of where they typically fluctuate in a normal economy. But in the extreme scenario, there's a possibility that unemployment really spikes, and that would be driven by knowledge work jobs because in the knowledge sector about the number of jobs would decline by about 22%. Wages in the sector would also experience heavy pressure. But on the other hand, in the remaining parts of the job market, there's going to be shortages and the wages of other workers would actually rise by about a third, and that would basically pull workers from knowledge work into the rest of the economy where they would be urgently needed.

[00:27:42.06]

So to conclude, there are many possible scenarios, and at this point we can't rule out any of them. In the first scenarios, The hard problem is not whether we have gains, but how we distribute these gains. My third point was that there is a small set of measurable parameters which will tell us where we are headed, how fast the technology evolves, how fast it diffuses, and so on. And then my final pitch is we need to prepare for the whole range of possible scenarios. because it's always better to be prepared than to be caught flat-footed. Thank you. Back to you, Anna.

[00:28:36.15] - Anna Stansbury

Thank you for that fascinating presentation, Anton, and I think using those scenarios to frame our conversation and how we think about implications and responses will be hugely valuable. Filiz, please take it away.

[00:28:48.25] - Filiz Unsal

Thank you. Thank you so much. So now I'm going to add one more parameter to To those that Anton mentioned. Sorry, let me just— this one. So I cannot see my— I don't know if there's something wrong, but I cannot see my— Yeah, we can't see your slides either

[00:29:19.06] - Anna Stansbury

right now.

[00:29:19.23] - Filiz Unsal

Yeah. Cannot see that window. Okay, maybe I'll do that. And it's very strange, but the PPT doesn't appear when I want to share the screen.

[00:29:44.21] - Anna Stansbury

Are you able to just share the screen and not the window?

[00:29:47.28] - Filiz Unsal

I'm trying, but let me do it again. Yes, maybe now it's gonna work, I think.

[00:29:54.18] - Anna Stansbury

Brilliant.

[00:29:56.17] - Filiz Unsal

All right. Well, thank you so much. It's a pleasure to be here. So, as I mentioned before, I'm going to add another layer to these different possibilities by focusing on global linkages, the role of trade, and what type of distributional differences across countries we can expect based on our work at the OECD. So as Antoine mentioned, there are a lot of different scenarios and people have been working on it. As you see here, the estimates vary across countries quite a bit. At the OECD, in our earlier work, where you see in red, we focused on the US and our expectation is around 0.4 to to 1.5%, 1.4% with some faster adoption scenario. Of course, there is a lot going in here. Many people are using relatively similar exercises based on some of the task-level gains from the earlier booms and some of the data that we get from the more recent work. The productivity gains for the US is around 0.4 to 0.5%. But of course, as Anton mentioned, there are a lot of differences and heterogeneities in these assumptions that we can vary and see where things are going. So in our approach, we go from micro to macro, and in some ways it's very similar to what Anton mentioned in the sectoral level.

[00:31:33.27]

You see here that we combine the micro-level productivity gains with sectoral AI exposure and to come up with certain scenarios over the next 10 years, what type of productivity gains we can expect from AI. And these are translated into welfare gains, and we look at both in terms of the domestic impact and foreign. But next, we also add a global dimension to it. So as you see in the middle— sorry, this is a bit too— yeah. In the middle, we have the global macro model where we take these domestic gains and put it in a model with 76 countries and 45 sectors to be able to take into account all these global production networks, both for final goods and for intermediate goods, to see what type of global differences we can expect across countries and what type of parameters we can identify which drives countries' gains from AI. So, here you see that the AI's— the gains across countries differ quite a bit, and there are large differences. Of course, this— one of the main parameters of these differences is the adoption across countries. In this work, we were— we did a lot of work in harmonizing across countries what type of AI gains we can expect.

[00:33:00.23]

However, of course, these are— this harmonized data on adoption is not great. And we are actually collaborating with Google and Ipsos on a cross-country adoption survey to get a bit of a sense on what we can say across countries on these macroeconomic gains. Here, I should also mention that, you know, again, our— the model, when you do this type of model-based assumptions, you are taking a stand on how people use AI in the sense that they are actually using AI for work and they are using it to do core work, which is very different than just taking and using AI for the purpose of asking a question that you can ask to Google. So therefore, I agree with Anton that there's a lot of heterogeneity, but this type of global surveys are helping us to understand a bit more what's happening in that space. And hopefully I will be able to talk more on that in the, in the, in the coming weeks. As you see here, the United States and several smaller European economies are among those with the largest productivity gains. And the reason is that they combine relatively strong adoption with a concentration in knowledge-intensive sectors where the current AI applications have considerable potential.

[00:34:22.26]

The projected gains, however, are much smaller for some of the Latin American economies. And among the non-OECD G20 economies, we see that China and India have relatively large gains, while Indonesia and South Africa have so much smaller ones. The central point here is that adoption is a major driver of these differences, and sectoral specialization matters too. But the technological potential really needs to translate into actual use. So, of course, these are just scenario results, and they are not fixed ranking of countries' features. And as the countries are

connected, these gains will reflect, of course, how they are connected to the world and the developments in trade. So let me unpack some of these numbers. So in our work, we are unpacking this international impact or the impacts through trade with 3 channels. The first one is more about benefiting from cheaper intermediate and final goods from abroad. So the idea is simple in a way that as the other countries that a country is trading with are benefiting from productivity gains, their import prices are going to be lower. Therefore, the country, domestic country, will also benefit from these productivity gains abroad. The second channel is more of an export channel because, as of course, as countries go through structural transformation adopting AI by their firms and companies, their productivity competitiveness is going to change.

[00:36:05.23]

Their terms of trade is going to change. And therefore, this is a second channel through which we will see how countries' impact are going to be affected by AI. And the third one is international knowledge spillovers, because we all know that trade also brings some of these positive spillovers. For example, if a country managing inventory is trading a lot with other countries that are using AI, there could be some potential benefits. So I will be going through all these 3 channels now separately. So let me focus on the first one. Again, the cheaper imports transmit substantial benefits across countries. If foreign suppliers become more productive, the benefits reach domestic firms through cheaper imports and households through cheaper final goods and services. This includes also indirect effects along the production networks. A country need not import directly from the, from the original AI innovator to benefit from this because we model the entire production network. Therefore, the productivity improvements to across these production network can transmit to countries. In this decomposition, we find that the foreign productivity improvements account about 20 to 30% in gains on average, and more than half of this occur to Latin American economies.

[00:37:29.04]

So this is an important source of gain, particularly where domestic adoption starts from a very low level. It also illustrates why we need a global framework to assess AI's economic effects. But it raises also a tempting question, which is that could countries simply rely on adoption abroad and import benefits? And this question, this decomposition that I show you here cannot answer this question. So we need another experiment where we— to capture this type of mental exercise to see if countries do not adopt themselves, how much they can gain or lose from trade. So here the main idea is that basically the foreign adoption and cheaper imports just cannot substitute for domestic adoption. And here the main channel, which is I think is a bit less discussed is the loss in terms of trade if the domestic country do not adopt AI and their firms do not benefit from this productivity. So here we take each country in turn and assume that it

doesn't adopt AI while the rest of the world does. We then ask how much it gains through these international connections. So here the benefits from cheaper imports remains, but is largely offset by larger changes in competitiveness and demand.

[00:38:56.02]

The net gains, as you can see here with diamonds, are generally small and in some cases even negative. So the negative effects are particularly more pronounced in economies specializing in highly AI-exposed activities such as United Kingdom and Israel. These are also the sectors with substantial upside from adoption, but also substantial competitive pressures if the domestic firms fail to adopt. This is a result about falling behind adoption, not a case for reducing trade, obviously. Trade remains the source of gains, and in the model we show what type of welfare gains we are talking about in the much more— in the order of magnitude. However, the implication here, the lesson here is that the openness and domestic adoption need to go together. So, the trade support, of course, can support the benefits too. But there's a third channel, which works through the knowledge spillovers. Because, of course, with these trade linkages, we shouldn't treat them as static. We should treat them as also benefiting on both sides, maybe in terms of the technology spillovers and adoption spillovers. So, we modeled these spillovers through countries' exposure to AI-adapting suppliers, in intermediate input trade, and we examine a standard scenario informed by the innovation literature and illustrate a stronger spillover scenario.

[00:40:28.11]

The potential uplift is particularly large in several Latin American and Central and Eastern European economies, and these combine relatively low initial adoption with strong links to more advanced adopters. So, obviously, the magnitudes are uncertain, depends on our assumptions, but Here, I would like to emphasize the mechanism behind this, which is that international integration can support convergence when firms can absorb and apply the knowledge reaching them. And in all these scenarios, obviously, we are— there's one prerequisite or assumption, which is that the firms have access to leading AI models and infrastructure that supports them. That access cannot be simply taken for granted. In fact, in this final chart, which comes from our preliminary work, I show substantial cross-country differences in data center capacity and investment in AI, with the United States standing out even relative to the size of its economy. This distinction, therefore, I would make is between developing frontier AI and being able to use it proactively. Obviously, countries can benefit without developing their own frontier models, but they also need reliable technology and domestic capabilities to use this. So therefore, if we would like to go to the policy from what all this implies, that obviously preserving international connections matters a lot, and also, but this is not enough, countries need to strengthen their

domestic adoption capacity, and also pay attention to countries and firms at risk of falling behind.

[00:42:12.04]

I stop here. Thank you.

[00:42:14.27] - Anna Stansbury

Fantastic. Thank you for such an interesting presentation, Filiz. And I think, you know, quite a striking and sobering takeaway that essentially whatever assumptions you make, because this is a massive productivity shock to knowledge-intensive activities, it's going to be something that boosts the rich countries more than it boosts less rich countries. So widening an already wide global income distribution gap. So we have a lot that we might discuss. And I wonder if I could kick us off with a question about assumptions about the market structure of the production of AI tools. And particularly, I'm interested in, you know, starting where Martin left off about the question about whether gains accrue to frontier model makers or to adopting companies and consumers of AI, and whether and how US-China competition and the different models of the US labs versus the Chinese labs affect that. I'm also interested in then hearing from Anton and Felice about how different assumptions about whether the labs capture the models versus adopters versus consumers play into your scenarios, Anton, about distributional effects, and Felice, about how we think about whether that affects the cross-country distribution of gains and losses.

[00:43:32.12] - Martin Chorzempa

Should we jump in on those?

[00:43:33.26] - Anna Stansbury

Yeah, yeah. Martin, can you go first? Yeah, absolutely.

[00:43:35.21] - Martin Chorzempa

Sure. Yeah. So one of the really interesting things we've seen recently is that firms around the world, especially US firms, are really— have seen since the adoption of agentic AI in particular, their token bills just exploding. And there's been then obviously, people are doing that because they see this positive benefit. But then also they're thinking, how do we make sure we control the cost. And, and so the Wall Street Journal quoted an executive that's very funny saying, I, I don't need a Lamborghini to go to the grocery store. Like, you don't necessarily need Fable to

like check your calendar, that kind of thing. And, and, and so this cost competitiveness leads to, leads firms to both select less advanced models, which by the way can do everything that the most advanced models could do not that long ago. But then also to consider the Chinese, the Chinese models. And so that is really where some of this competitive pressure comes from. Now, there's something really, really interesting that's happening right now, which is that the Chinese companies that develop these open weight models, yes, they give them away for free. But one of the reasons that they've been giving them away for free is that they don't lose any money from giving them away for free.

[00:44:49.07]

That sounds bizarre. But the reason for that is because they're heavily compute constrained. they can only serve a certain amount of demand for their models. So if you look at Moonshot, which releases Kimmy, for example, they had to quickly cap and say no new users for Kimmy K3, not allowed, because they didn't have the computing power to serve them. However, you have even companies like Harvey in the United States, which is the main AI tool that lawyers use, now has the main model that lawyers all around the world are using, mainly in US firms, are using to understand the law is a fine-tune built on top of the KIMI model. Now, they don't necessarily have to have any relationship with KIMI for that, for that to happen with Moonshot. To do that, they don't need to send any data that way. But that's just an interesting way to think about the adoption. Now, though, most firms do not have their own computing power. So they have to access these models through cloud providers. And the really interesting shift here is that Moonshot and other companies are beginning to say, we want to make revenue if you're cooking based on our recipe, if you're a cloud provider.

[00:46:02.07]

They don't want to see Amazon, Google, and Microsoft making giant profits from charging other firms to serve their model and then not making any revenue from it themselves. So Reuters has reported that Kimmy is at— that Moonshot is asking for about a 30% revenue share. And so what's interesting here is that the era of being able to access these cheap Chinese models for free might be ending, and that you might have to start paying more. Obviously, Amazon and Microsoft and Google are not just going to absorb that themselves. They're going to pass some of that on to the consumer. So, so we're no longer maybe in a world in which unless you have your own spare NVIDIA GPUs sitting around, which is not a common activity in the economy today, you're going to need to pay more for the Chinese models. And then, as I established in the, in the talk, the cost of using US models is declining. And another thing is that the efficiency of Chinese models is actually usually less than US models. This is another narrative violation that Chinese firms charge less per token, per word that they, that they process.

[00:47:11.06]

But US models tend to actually be more efficient with tokens. than Chinese. And so that's why you see that on that chart I showed, the cost per task being lower in the US one. So you have to really drill down past the cost levels. But overall, the fact that this competition is there does tend to reduce the cost of models because the US companies or Mexican companies, Kyrgyz companies can say, hey, you know, Anthropic, you want to charge us all this for the model, but hey, there's a Chinese model available that can do a lot of that thing. available at this lower cost, can you make me a deal? And so that will tend to lower those returns. Now, the question is, are they high enough to retain the level of spending that US firms have had on computing power? I would generally tend to think yes.

[00:48:00.25] - Anna Stansbury

Really interesting. Thanks, Martin. Yeah, I'm curious, Anton, about how this plays into your scenarios and how these different assumptions affect especially the capital labor assumptions about distribution.

[00:48:10.02] - Anton Korinek

Yeah. Yeah. So the way that we model the economic gains from a macro perspective is we observe that in almost any task that AI touches, if you do like a kind of case study, you find that the cost of AI compared to the cost of how it was previously done is usually an order of magnitude smaller. So in other words, if AI can do something, there are huge productivity gains, and the vast majority of these gains goes ultimately to the economy at large. So, as Martin emphasized, there is considerable competition. Competition makes sure that the gains are spread, and in our scenarios, that is also the kind of working assumption. Now, I think from a distributive perspective, what I consider the most important factor is whether something can be done by AI alone or whether something can be done by AI together with the worker. So automation versus augmentation. And from a macro perspective, that really seems to be the main driver of how the gains are distributed.

[00:49:36.26] - Anna Stansbury

And in your, just if I can just follow up on that one quickly before we go to you, Filiz, in your main scenarios, Anton, about the very transformative scenario versus the moderate ones, are there different assumptions about automation versus augmentation baked into those?

[00:49:51.06] - Anton Korinek

Yeah, that's actually a good point. So in the modest and substantial scenario, we assume that there is considerable augmentation. So it means the role for workers is still very significant even in tasks that the AI can already help with. Whereas in the extreme scenario, the assumption is that most of what AI touches will be automation, and then the outcome for labor is considerably worse.

[00:50:19.29] - Anna Stansbury

Right. And I think that's a really important insight that complete automation is much more effective in generating high productivity growth because augmentation is a you still need the human in the loop. And so there's only so much you can increase the productivity of the task to some extent. And so you've got automation leading to higher, the highest scenario for productivity growth, but also the most scenario for labor displacement.

[00:50:40.01] - Anton Korinek

Yeah. Yeah. So from a perspective of making sure that workers share in some of the gains, the augmentation is actually the one that creates less disruption.

[00:50:50.11] - Anna Stansbury

Yeah. Really interesting. And Filiz, I wonder, you know, you said in your scenarios that it depends very much on assumptions about access to AI technologies and also the physical infrastructure piece. And so how does this question of what kind of models are available and where the gains accrue, and if they are, you know, if they end up as monopoly profits for the frontier labs versus distributed by a competition, does that change your outcomes?

[00:51:12.27] - Filiz Unsal

I mean, I would like to make a couple of points. First is that, of course, in the model it's a TFP shock, so we don't really take a stand on the prices. But, but obviously, you know, we are talking about scenarios next, next 10 years, right? So we are assuming that during that 10-year period, the firms still will have access to the frontier models, the, the sort of the adoption of the technology is going to go at a relatively high speed, at the, you know, a little more than maybe internet, but comparable. So all these requires that actually the adoption continues. And I think adoption itself should be endogenous in a way that it's a function of the prices, obviously, and trust as well. But also, I mean, the prices play a big role there. And in some ways, you know, you may think that, of course, we are presenting these, you know, we are picking certain

parameters and making scenarios. But even the exposure, you know, the sectoral exposure is also even endogenous to some of these, you know, the different elements, exposure itself is country-specific. You know, being a restaurant worker in France is very different than working in a restaurant in the US.

[00:52:22.06]

So, you know, there are different parameters. And I think the prices and the trust has a big impact on there. But there's another point that I would like to make, which is that, you know, there's also the reallocation within the economy. So we are focusing on the headline numbers, but there's a what we call boomhole effect. I mean, of course, I don't need to explain to this audience that much. But maybe for our viewers, you know, at the end of the day, productivity gains in one sector, and if it is not distributed to the rest of the economy, it's going to bring the prices in that sector and increase the prices in other sectors. And over time, we experienced this before as well in different technological transformations, where that other sectors may over time increase the share their share in the economy. Therefore, aggregate productivity actually comes down. So, you know, these— of course, the improvement in AI prices is great, and it allows access to, at least so far, increasing the access of firms. However, we don't know, you know, going forward whether these gains are going to be distributed across sectors. And I just— I think this is an important policy issue.

[00:53:31.03]

And even within the firms, we have done some work on how the frontier firms in which sectors adopt AI versus not. And in fact, you know, the firms that have much more— already have higher markups and already in ICT sector are adopting AI much more quicker. So in a way, currently, who is adopting AI, who is actually changing the organization in a way that our models in some ways envisage is the frontier firms that already have money, already have resources invest in this. Of course, I mean, the rest, the rest of the economy too. But that reallocation, the changes in the relative demand versus— even for the AI, ICT product as well, I think is an important direction. So it's not— I wouldn't take it as given that these gains are going to be distributed. Maybe Anton had a bit more optimistic view on that than I have.

[00:54:25.13] - Anna Stansbury

Yeah, that's really interesting. I mean, I wonder if we can build on that, because I think one, you know, one theme I'm taking away from these these presentations today is if we are thinking about who won from past technological changes and transformative technologies versus who might win disproportionately from the AI boom, you know, you look at computerization and you

see that within countries who won from computerization and AI, it was largely a technological change that benefited highly educated people who were working in knowledge-intensive sectors. So it was people that were already well off. And so you saw these big widenings of inequality within labor income and somewhat of a widening of inequality from labor to capital in many countries. And that was very consistent. In this case, Anton's scenarios would suggest it's the opposite. It's displacing knowledge workers within countries and sort of narrowing labor income inequality, although it might widen labor capital inequality. On the other hand, when you're looking at the international scenarios that you presented, Filiz, you know, what's looking like is happening is those that have already won are winning more, right? The countries that are already ahead, that are already specialized in the highly profitable, highly productive knowledge-intensive services are going to win, you know, advance more and see bigger increments of their productivity growth.

[00:55:42.09]

So is our takeaway that, you know, we flip the within-country effects from the last technological revolution, but we just build on the disparities across countries, or how do you all think about that?

[00:55:57.28] - Martin Chorzempa

Do you want to go on, Anton, or others? I can get—

[00:56:00.20] - Anton Korinek

Yeah, happy to go first. Yeah, Anna, you're exactly right. In our scenarios, that's what happens. In some sense, there is basically an undoing of the skill-biased technological change that we've seen in the past couple of decades, and the non-cognitive workers are actually the ones that reap the larger gains. And at the same time, income shifts from labor to capital to some extent. Now, I should also add a footnote. So our scenarios in what I presented, they go to 2030, and we assume explicitly that there are no big advances in robotics that are macroeconomically relevant. But if we also experience a takeoff in robotics, then that scrambles everything again. And then non-cognitive jobs may also become quite significantly affected.

[00:56:59.07] - Martin Chorzempa

I mean, one thing, Anton, in your scenario that's, I think, driving some of this is that there is, there is no differentiation between different knowledge workers. So just one interesting example I'll give you is talking to a lawyer recently. They're no longer hiring paralegals. Uh, but the partners are significantly more productive. And so a lot of this is gonna come down to something

what Alex Imus has called, you know, we should have a Manhattan Project for measuring elasticity of demand in different sectors. And the reason for that is the idea of, okay, uh, now you no longer need, if you're hiring a lawyer, you don't need a person to spend thousands of hours going over document review. Uh, AI can do that part quickly. And that frees up the partners to be focused on building relationships and discussing with clients strategies and all that kind of thing. If that— if the lower cost of legal services, because, you know, you're paying billable hours, but if the lawyer is 10 times more productive in that billable hour, you might demand a lot more billable hours, actually. So the hope is that in many of these sectors, we'll have this what's called Jevons paradox, where becoming more efficient and lower effective cost drives more demand, and you could actually have more employment.

[00:58:13.17]

But the question is, if the paralegals disappear, what happens there? Do we only have the people who already have the skills to benefit and to drive AI and be augmented? Does that outweigh the people who are unaugmented? So I think one interesting thing to think about, I'd be curious your thought, Anton, is kind of Below that overall surface of a decline, where you have some pretty significant decline on average in knowledge worker wages and in employment, do you see some of those workers really, really benefiting? And is it actually what I described where the paralegals are unemployed, or is it that we have a lot of middle management unemployed, or do you expect it to be more of the senior partners who are really expensive? Do they get you know, laid off and more tech-savvy young people who can use AI better? Do they use— what do you think?

[00:59:04.23] - Anna Stansbury

I'm afraid we won't have— sorry, Anton, we won't have time for that conversation to continue because we're almost at the hour. So I'm just going to invite Filiz for any parting thoughts on this question and everything else, and then we're going to have to wrap up and continue this further another time.

[00:59:17.01] - Filiz Unsal

Maybe that also goes a bit to Martin's question. In fact, in Anton's framework, there's also reallocation and boom-all effect. But not across the sectors, but across tasks. So, you know, it may increase the— decrease the wages in sectors where AI is very adopted and exposed. And, you know, maybe in some others that wages increases because, you know, they are not exposed to AI. And or maybe there are other things going on, like, you know, we have all these arguments about messy jobs, etc., that some of them cannot be done. And therefore, it's going

to be very, you know, not very much affected. One thing that I would like to mention here, and maybe leave the floor to Anna to also conclude, and also given her expertise, is that, of course, the labor share, how labor share affected in the short term and in the long term, depends on many things, including on, you know, different labor institutions and how things worked out. I think when you look at the history there was both sides that, you know, the laborers increased their power and at the same time they managed to keep their labor share maybe over time a bit more than what the technological advance can bring.

[01:00:27.15]

And obviously the policy also play a role. So I think that there is a lot of thinking has to be done, you know, how we can make labor adjust better in the both in the short term and maybe even prepare for scenarios if AI hits us even quicker. But I stop here. Thank you.

[01:00:48.17] - Anna Stansbury

Thanks so much, Filiz. And yeah, I think there's always much more we could discuss for many more hours, but luckily we'll all be continuing to think and write and speak about this. And so thank you all very much for joining us. We'll be continuing to look at the economic indicators that tell us exactly what scenario we're in. As Martin said at the beginning of the talk, they're changing weekly. So thank you all for joining us. Thank you, panelists for joining us today, and, um, have a great rest of your day, everybody.

[01:01:16.05] - Filiz Unsal

Thank you. Thanks.

[01:01:17.26] - Anton Korinek

Thank you.