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.

[00:05:27.55] - Antonio Spilimbergo

Hi everybody, thank you very much for showing up this early this morning. And we are about to start a very exciting session on technology, new jobs, and skills. We have the same format as yesterday. We have 3 presentations, and at the end, a discussant will discuss all 3 of them.

[00:05:50.41]

And at the end, we'll have questions. I, um, let's, uh, well, Let's start with the first speaker.

[00:05:59.44]

The first speaker is Tucker Smith from the Federal Reserve Bank of Dallas.

[00:06:24.35] - Tucker Smith

Hi, thank you to the organizing or the organizations that put this together and then the committee. Really excited to be here and presenting our work. I think it fits well with some of the discussion yesterday and some of the discussion at Chairman Warsh's Reserve Task Force on AI and Productivity on Monday. So, so very opportune time to be presenting this. We hope to have our working paper out next week. Our paper's about how AI is affecting labor demand and then outcomes for new college graduates, and then how are current students adjusting their educational decisions in response. Everything I say does not reflect any official or unofficial opinions of the Federal Reserve Bank of Dallas, the Federal Reserve System, or any of these Texas state education agencies. And so generative AI has the potential to improve productivity, drive economic growth, but there may be costs of the transition. And potentially, there's potential for significant labor market disruption. So far, there's— we'll show some evidence that this is falling on new labor market entrants. This really rapid deployment of generative AI has coincided with an unusual increase in unemployment rates for recent college graduates, which motivates a deeper examination if they're bearing some of the costs of this transition and how will new or individuals coming of age adjust.

[00:08:03.20]

So there are a lot of reasons that we particularly care about young adults in addition to that sort of motivating correlation. A lot of the early evidence in this literature suggests that gen AI substitutes for less experienced labor while actually complementing expertise or more experience. These sort of tacit skills that you learn on the job are what maybe allow you to use AI more effectively versus the more codified language that you would learn coming out of school, AI can replicate. And so Again, these young adults may be more affected. Given the uncertainty still about AI's ability, firms are, at least at first, more likely to reduce hiring than to lay workers off. This bears out in the data. And from past technological changes, layoffs often don't occur until there's a recession. But in the meantime, this reducing of hiring is going to affect people that are stuck on the outside of the labor market, such as new college graduates. We know from the literature, and as Gordon Hanson spoke about yesterday, that graduating into a weak labor market can have persistent— can result in persistent scarring. Your initial wage conditions matter a lot for your lifetime wage profile.

[00:09:16.46]

And finally, workers throughout the economy are going to need to adjust and develop new skills to potentially complement this technology. It's really hard to observe skill development responses among entrenched workers except for through formal retraining programs. But it's really easy to observe— it's easier to observe these among individuals attached to the formal education system. So in this paper, we examine how the Gini-Ashok is affecting earnings and employment outcomes for new graduates, then examining how these students are adjusting their educational decisions. And we'll have some suggestive evidence that we cautiously interpret as evidence of whether these adjustments appear to be paying off or likely to pay off for these early cohorts affected. We'll use data first on job postings. You've already seen this Lightcast job data already used in this conference, and then linked individual-level education earnings records for all students and graduates from Texas public universities, seeing their earnings if they work in Texas. So our research design is going to be similar to what you've already seen in some other papers here. We're going to be projecting these task-based AI exposure measures onto fields of study and then comparing changes in outcomes for more versus less exposed majors across cohorts relative to the release of ChatGPT and the AI boom.

[00:10:37.04]

Just to preview our results, first, we find that the job postings fell by about 5% more for occupations with a 10%age point greater share of automatable tasks. After the release of ChatGPT, this reached about 8%— sorry for that, shouldn't say 10%— by mid-2026. And this did translate to harm, actual harm when it came to new graduates. That first-year earnings for

new graduating cohorts fell by about $3,000 if you're including nonemployment as zeros. But if you look under the hood there, that's both new graduates from more exposed majors were both less likely to find jobs. and conditional on finding work, they were still earning less. This even persisted into their second year into the labor market for cohorts that we can observe that extended time. So how did students adjust? First, those that had graduated into a soft labor demand or a soft labor market were more likely to re-enroll in grad school. However, they largely stayed in the same field. And so, for example, 2/3 of computer science majors that came— or computer science graduates from 4-year universities or bachelor's degree holders that went back to grad school were enrolled in the computer science program.

[00:11:47.20]

There's not really this pivoting away from AI-exposed fields. And then among undergraduates, those that were still in college at the time of the shock continued along their existing trends. Those were a computer science major even entering in fall 2022 continued along like typical computer science majors. But those that came after both shifted their majors initially in college, and even some of those that had initially declared for computer science then shifted their course and major choice in subsequent semesters. And so this demand shock is seemingly salient enough to enter into students' decisions, which is encouraging because AI, in addition to affecting demand, also kind of affects the supply side of education. It can make it easier to get good grades in a computer science course. So we may be worried that that would counteract some of the demand signals. The demand side seemed to be winning out. Then, which— the final set of results, which we interpret with appropriate caution, are whether these adjustments appear likely to pay off. And so a few pieces of evidence that suggest the perhaps greater returns to what we'll call pivoting strategies or moving to to less exposed majors over sort of upskilling or trying to obtain additional expertise in exposed majors are one that new master's recipients from these exposed experience actually experience similar or maybe even larger declines in earnings as 4-year graduates did.

[00:13:14.25]

And so it's not that this additional formal expertise or expertise experience— expertise obtained formally through education doesn't appear to be more complementary to the technology. And we can show for undergrads that those that's from more exposed majors that had really specialized course portfolios in these exposed fields like computer science, they did even worse than the average computer science major. Versus if you had a more diversified course portfolio and took courses in other fields, then you offset up to about 25% of that earnings penalty. So how do we do this? As I mentioned, we have 2 data sources. First, job postings from Lightcast, which is the near universe of online job postings. Just because our main datasets, the education

records, are for Texas, I'll be showing results just using job postings for Texas. If we use the whole country, then the results look the same. And so we're going to be looking at sort of trends in job postings by occupation and then also by firm. And so we'll— I'll show in a minute how we're going to define AI exposure at the occupation level and then projecting that into the firm level as well. Our second data source comes from the University of Houston Education Research Center, which is linked individual-level education and earnings records, where we can see for college students both their college entrance exam scores, what courses they take, their major, any degrees they earn.

[00:14:41.49]

And then if they work in Texas in a UI-covered occupation, you can see their earnings, employment, and their industry. So we're going to construct panels of both 4-year or bachelor's and master's graduates and then also of college entrants. We adopt similar kind of— we adopt approaches from the literature on how to consider or how to define AI exposure to this new technology and potential automation. Kind of 2 camps of— or 2 main approaches from this literature are one would be defining kind of theoretical— what we call theoretical exposure, which comes from, say, asking either human annotators or an LLM itself, what tasks can you perform? Just what tasks can you perform more quickly than a human? Another approach, which is what we will primarily use, is based on the actual— observed usage of Claude users. And so Anthropic, for a subset of Claude users in fall 2024, could observe the tasks that people were using the technology for. And particularly, they even classify if they're using it in an automated capacity. We're going to lean on— we're going to primarily use the second measure given— although it's slightly more endogenous, it correlates better with usage and the kind of— LLM-based measures themselves vary a lot by what model you're asking it can perform.

[00:16:10.08]

And so they're a little less stable and reliable. We'll then project these occupation-based measures. And sorry, I should note that an occupation's exposure then is defined as the share of its tasks or its core tasks that the new technology could perform or people have— or new Claude users have been observed using in an automated capacity. We will project this onto college fields of study or majors based on what occupations in their job posting listed particular majors. So for, say, a computer science major, if 50% of the job postings that list you are software developers, 25%, say, teachers, and 25% others, then you're going to get a lot of weight from the computer or software developer exposure level. In a big clunky table here, the— it just shows that the field-level exposure measures are fairly intuitive. The majors that were most exposed to this new technology are computer science, computer engineering, and languages. While the least exposed include psychology and nursing. We utilize an event studies

specification with this continuous exposure treatment. And we're regressing an outcome for individual i from a particular major and cohort on their major-level exposure interacted with cohort dummies, and then controlling for cohort and year fixed effects and student characteristics.

[00:17:39.08]

We'll actually augment that to university-by-cohort fixed effects so that we're limiting our comparisons to students, graduates from the same university that took different— or that had different majors. This is— requires— a continuous diff-in-diff requires slightly stronger assumptions than our typical diff-in-diff. Everyone is exposed here, and so we can't use the standard parallel trends arguments of in the absence of treatment that outcomes would trend in parallel, because even the control group isn't in the absence of treatment. Everyone is affected. And so the actual assumption is closer to assuming homogeneous treatment effects. Essentially, that if less exposed majors, say nursing, only had— or excuse me, if, say, computer science were only treated by 5% automation instead of 30%, that their trends would look similar to what we actually see among psychology majors. Our estimates then that I'll show throughout are scaled to represent a 10%age point difference in automatable task share. And they'll be relative to this final cohort that we see in most cases graduate before the start of the shock or experience their first year in the labor market before the release of ChatGPT. It'll be the 2021 graduating cohort. We are very cognizant that there are a lot of potential confounders or other shocks happening in the labor market and on the education side.

[00:19:09.33]

For the most part, Our results are really robust to controlling for any of these issues. So what people have proposed in the literature are that this isn't about AI, it's really about a tech boom and bust or a differential bite of monetary tightening cycle. This is actually about remote work, that that makes younger workers less productive, and so then firms are naturally hiring fewer of them. we can control for all this and everything looks pretty robust. So getting into our results, first we're looking at in the job postings data, this is a PPML regression, slightly different than the specification I just showed, but this is essentially looking at trends in postings for occupations by their exposure level. And you can see before we released ChatGPT and coming out of the pandemic, there was a relative increase in the frequency of job postings for occupations that ended up being more exposed to AI. And this is kind of that tech boom story. And so just from an accounting perspective, we may think some of that would have unwound after the technology came out. And so these— you could shave off some of these actual estimates in the post-period.

[00:20:28.23]

But it's not enough to explain all of it. And so in the subsequent years, occupation or postings for more exposed occupations fell by about 5% within a year of the release of ChatGPT and then up to around 8% through midway through 2026. We can also look at how individual firms have changed their job posting behavior. And so if we look at firms that prior to the release of ChatGPT were posting more for these more exposed openings, assigning each firm the exposure level of their average pre-ChatGPT job opening, we see that these firms that had been hiring a lot of, say, computer programmers or software developers were the ones that were pulling back their hiring after its release. This helps tell us, is this a story about new firms or is this existing firms? We see it even among these existing firms in a balanced panel. It's not just that say, older firms are shutting down and AI-native firms are coming in and not hiring anyone. This is even existing incumbent firms reducing their hiring. Moving on to how this affected or how this is affecting new college graduates. Now, this is a cohort-level event study tracing out this relationship between major-level exposure and earnings in your first year.

[00:21:54.06]

So this is including— zeros for people who do not find a job in Texas. And we see that after trending similarly before the release of ChatGPT, earnings subsequently fell for graduates in the 2023 and 2024 graduating cohorts by up to about 10% relative to the pre-period mean. This occurs both on the extensive side of finding a job and on the intensive side of your earnings conditional on finding work. And for the 2023 cohort, we can— our data, we can observe them into their second year in the labor market. And it shows that these effects kind of persist. You would expect this given how relationships are— given evidence from the literature that initial wage conditions are persistent, that if you graduate into a recession, for example, you essentially have lower earnings for at least the next 10 years. We know less from our literature about actual kind of field-specific shocks, but that seems to still be holding true in this case. And so some of the mechanisms here from the scarring effects literature that likely apply here are that if you graduate during a period of reduced demand, that you're more likely to kind of have a skill mismatch with your actual position that you find.

[00:23:15.38]

And you're not obtaining the kind of occupation industry-specific human capital because of that mismatch. And so those likely apply here as well and potentially could be even more important in this setting if, as some evidence shows, that GenAI actually complements this on-the-job experience, these tacit skills that you acquire while working. In which case, kind of being behind the 8-ball early, not getting into a job that fits your skill set, is going to prevent you from acquiring those skills. And so you could have even these effects possibly even amplifying going

forward. So then how are students adjusting to this shock? We're thinking of our kind of mental model that we— in the background is that we're thinking of GenAI as affecting both the benefits and the costs of education. On the one hand, the results that we've shown so far are consistent with it imposing what we call an automation penalty on graduates with more exposed skill sets. As these tasks become cheaper to automate, the earnings premium for a computer science major seems to have fallen. However, it's possible, and some literature suggests, that the AI may complement expertise. And so we're interested then, or we are open to the notion potentially that advanced skills that you couldn't even learn through, say, additional coursework or a master's degree could raise the value of working with AI.

[00:24:44.23]

And on the cost side of education, this is not just affecting demand, but it actually can make courses— it can automate schoolwork itself and potentially automate schoolwork in the same fields that it's automating work in the labor market. And so if it makes a computer science major much easier to complete and to get good grades in, then that could counteract the incentives from the demand effect and actually pull students into these previously difficult fields. We'll be considering that a field will lose appeal if the demand side dominates, or this automation penalty dominates the reduction in effort costs. We consider 2 strategies, or 2 basic strategies, that students can adopt to adjusting to this automation. One is that they can pivot or shift to a less exposed field, or they can upskill and try and deepen their expertise through advanced coursework or graduate study. So this dominant strategy is going to depend on the existence— whether there is actually the existence of these complementarities and how students affect the persistence of this shock on— or the persistence of the automation penalty and the demand effects. in addition to these effort costs. We'll consider, though, that switching cannot happen instantaneously.

[00:26:04.20]

There's naturally some costs. If you're a junior computer science major when ChatGPT comes out, even if you have perfect foresight of this is going to reduce demand, I'm probably going to earn less, there's a path dependence on your track. And so we're going to consider switching costs increase in seniority, and that newer entrants or younger students are more likely to adjust. First results on the education side is that graduates from more exposed majors were more likely to go back to grad school. It's not the prettiest looking event study. And so the point estimate maybe shouldn't be taken super literally, but there does seem to be some evidence of an increase of about almost a 1%age point or 6% relative to the pre-AI mean. However, we don't see any evidence of them shifting fields. The actual exposure level of their field, if you put that on the left-hand side of the study, it's flat. And as I said before, in simple language, about

2/3 of computer science majors that went back to grad school chose a computer science program. It's essentially unchanged from the pre-AI relationship. So this is consistent with some students adopting an upskilling response.

[00:27:15.38]

They're developing They're trying to develop additional expertise and taking advantage of the— or in responding to the lower opportunity costs of graduate school. Looking at undergraduate enrollment here by actual classification or grade level, we see that driven by the increased popularity of computer science, these more exposed fields had greater enrollments or stronger enrollment growth. leading up to the shock. And that kind of continued for upperclassmen. So even if you're a senior in 2025, then you entered in fall 2024 and likely declared— in some cases, may have declared your major before even this shock happened. But the new college entrants are pivoting away. You see this nice kind of waterfall-type effect where for new freshmen, freshman enrollment actually fell by about 6%. from fall 2022 to fall 2025 for a 10%age point difference in automatable task share. This is similar to the delayed enrollment responses you see to prior field-specific or sectoral demand shocks, such as the dot-com boom. And it's encouraging, again, given that there is this effort-cost channel as well. And so it's not given that the demand side would dominate, that students would really internalize the shifts in demand, but they seem to be doing so.

[00:28:40.46]

And so finally, we have, again, which we interpret cautiously, evidence on if these adjustments are likely to pay off. First, this is a triple difference here of the difference in relative earnings of master's versus bachelor's graduates from more versus less exposed majors over time. You can see that in levels leading up to the shock that master's completers from these more exposed majors were earning even more and more. This master's premium was going up. But that has changed. That's flattened out. And so they're experiencing— if I just have master's on there, it's an absolute decline. And so they're experiencing similar reductions in demand as bachelor's degree holders. Or in other words, they're not benefiting from any additional complementarity. This could change, though, as schools adjust curriculum and try to specifically teach skills that are more complementary to AI. But as of now, even this additional degree does not seem to protect students from the automation penalty. And then finally, we can look at the course level and see even within undergraduates, our students that developed or that specialized more and attained deeper expertise of these exposed fields, how do they do in comparison to even students from the same major who maybe had more diversified course portfolios?

[00:30:12.49]

This presents the marginal effect of— and sorry, I should note, what we define then is we formally will define individual level or an idiosyncratic exposure term, which is the mean exposure of all of the courses across fields of the courses that you took relative to the average for your major. And so what we find is that if we then estimate a specification that interacts this with major-level specification, we find that the marginal effects of taking additional exposure— more exposed courses become more negative as your overall course or as your major itself is more exposed. In other words, if you are a computer science major and you really specialize in computer science, your earnings fell by even more. And there's no point in this distribution, just from this at least similar, pretty simple linear specification, where the additional expertise from individual courses paid off or improved your earnings. Instead, it was only a negative, that this kind of was always— in any form of investment in these skills development in these more exposed fields seems to not have paid off. On the flip side, if you frame this in reverse, that if you were a more exposed major and you took less exposed courses, well, then you could have— that would offset about a quarter of the initial earnings penalty.

[00:31:49.41]

And overall, along with this, this suggests potentially these benefits to this kind of pivoting or investing in other skill sets rather than formally upskilling through the education system. So in conclusion, as has been— everyone here is aware that the widespread labor market disruption has not arrived yet from GenAI. But we're providing some evidence of it causing kind of various degrees of harm across 3 groups. First, these recent graduates from exposed majors who had skill sets that could be automated have experienced worse initial labor market outcomes. And this seems to be— this is potentially persisting deeper into their careers. Students in college at the time of the shock, for the most part, did not shift fields. And so they may experience similar harm as these recent graduates. But those that entered after the college— after the shock and had more time for this to become salient and enter their decision-making do seem to have shifted fields. And so they may be relatively less exposed or maybe more buffered against the shock. Of course, there's a lot of open questions about AI and education system. Of, I think, first-order importance is what actually happens to human capital accumulation, not just on paper, but are students potentially learning less.

[00:33:13.42]

And more broadly, if we're democratizing access to expertise, I'll be interested to see what happens to the demand for schooling itself. But that's all I have. And thank you. Looking forward to discussion and questions and the rest of the presentations in this great session.

[00:33:34.47] - Antonio Spilimbergo

Excellent start of the day.

[00:33:36.37]

We go on with the second paper by Tarek Hassan of Boston University, Old Technology, New Technologies, and the College Premium.

[00:33:52.50] - Tarek Hassan

Yeah, thank you so much for—

[00:33:56.44]

sorry.

[00:33:59.09]

Yeah, is there a keyboard? Oh, click— I see. Oh, wow.

[00:34:10.12]

Okay.

[00:34:12.29]

Who knew? Okay, great. So thank you very much for inviting me to contribute to this session. I guess I'm going to back up the bus a little bit and ask the question of how in general new technologies have affected wage inequality in the United States. So in some sense, the question is both broader because I'm going to be talking about all new technologies and narrower because I'm just going to be talking about wage inequality. And in fact, I have nothing to say about unemployment and all these terrible other things that we just heard about. So, okay, so let me— so this is joint work with Akash Kalyani and Pasquale Restrepo. Akash is at the St. Louis Fed. Nothing that I'm going to say here represents any official positions by any of these bodies. Okay. So a big problem in economics has been to think about this graph, which is the college wage premium has moved pretty dramatically between 1980 and 2010 and then sort of flattened out. And this is a major driver inequality across skill groups, regions, and all kinds of other problems that we have currently to deal with. So the college premium increased dramatically, and this increase was steeper in the densest areas of the country.

[00:35:37.55]

So for example, in New York City, basically today a college grad makes almost twice the wage of a non-college grad, whereas in the rest of the country, it's 70% more. Okay. So existing theories, of course, people have studied this a lot. Existing theories have been focusing on sort of postulating basically a skill-biased technical change or thinking about capital skill complementarity, automation, IT, and so forth. What we do in this paper is we use new job-level data on technology creation and use. So we're going to have very granular data from text and we're going to use that to show that the skill premium depends in general on the pace of technology creation and standardization. It's going to be these 2 forces that battle with each other, which is how many new technologies are created and how quickly are these technologies standardized so that everybody can use them. And we're going to claim that this measured technology creation actually explains a surprising amount in the variation in the college wage premium across skill groups, space, and cohorts. So I'm going to briefly show you basically a one-slide model for how to think about this. And then we're going to go into the data.

[00:37:07.38]

So the purpose of this model is to basically help us think about how different drivers of technical change affect wage inequality. And we're going to— again, we're going to do this on one slide. So I want to introduce this M of B. M of B is going to be the star of the show. M of B is the mass of new technologies that are introduced at a birth at a given date with cohort B. So think of M of B as the number of new technologies that were born in 1980, for example. So at a given time, I can denote this mass of technologies as at a given time t, these technologies from 1980s, for example, if I looked at them in 1990, they might be 10 years old. So u is the age of the technology. Age is the date minus when they were born. Now we're going to write down a production function that's going to look very familiar and it's going to make one twist. So I'm going to say each technology vintage that is born at a given time is going to produce output according to this CES production function where I have 3 forces affecting this.

[00:38:31.17]

So the first is how productive are the technologies that were born at date B? And we're going to just assume that newer technologies in general have higher potential. So this A of B is going to go up over time mechanistically. And then we have 2 new important objects. Z of U is the productivity lifecycle of the technology. So for example, 1990 solar panels were hard to use and didn't produce all that much electricity and so forth. And then you have a lifecycle where they get better over time. So Z of U is going to be the productivity lifecycle of the technology. And alpha U is going to be the skill bias of this technology, meaning how much does this new technology complement a college degree? OK, so and then given this alpha U, we're going to

produce the technology with some high-skill labor using this technology of vintage U and some low-skill labor L. And then the other parameter is gamma. And that's sort of like an elasticity of substitution between high and low-skill labor. OK, so that's it. Now, OK, I already said this A is going to grow at some exogenous rate. And then how do I get to output for the whole economy?

[00:39:56.40]

Well, that's just going to be the integral of all the technologies that we have of different ages. And that integral introduces a 3rd object here, which is how many technologies are of age U. So if you look at this slide, it's a model. It's a model of the skill premium, which how much more do high-skilled people make than low-skilled people? And all the red things are the things that we need to know. So the point of the model is to basically say, well, if you want to know what the college wage premium is going to look like, you need to know all the red stuff. Okay, so if I now solve the model and, you know, clear the market— so every dollar made has to be paid to someone— clear the labor market for high-skilled labor and clear the labor market for low-skilled labor, I'm going to get my equilibrium wage premium. So if you look at the slide, basically now the research question is, what is M? What is alpha? And what is z? And the way I'm going to get those is using new data on technologies that we generate from text, from the full text of patents and job postings.

[00:41:08.41]

So let me just pause here. Does this make sense so far?

[00:41:11.36]

OK. So all right, so now with this one-slide model, you can show 2 important results. So the first is that in the long run, the skill premium doesn't depend on how much we invent. It doesn't depend on m, how many new technologies are invented on average. So consider a balanced growth path in which the pace of technology creation is constant. So let's say I'm going to fix this m of b at m. And I'm going to have no other changes. Now if I set— I have 2 parameters here, gamma and sigma. If you set them equal to each other, you can solve for this here. which is that the college wage premium— this is WH divided by WL. So this is how much more do high-skilled people make relative to low-skilled people. That's a function of this thing times H over L. And this is sort of a famous formula. This comes from Katz and Murphy. Okay, so basically what this says is that the college wage premium depends on what share of GDP is produced using high-skill versus low-skill labor. And in Katz and Murphy, this was sort of a reduced form thing. And they kind of backed it out from observing the wages in the aggregate economy.

[00:42:36.41]

We're going to have a microfounded model of how this thing— we're going to have these sort of z and alpha functions that tell us how this thing evolves.

[00:42:45.50]

OK.

[00:42:47.05]

So what you see in this sort of like— in this equation is that there's no m's. So in the long term, all that matters is a battle between alpha and z. So alpha of u is the skill bias of the technologies of age u, and z is basically their productivity lifecycle. So if you— if the skill bias of new technologies kind of erodes very fast, so meaning as a technology ages, it becomes easier to use and you no longer need a college degree, then the college wage premium is going to be lower. If— but if it takes a long time to basically get non-college workers comfortable with a new technology, then the skill premium is going to stay elevated forever. So it's basically a battle between these 2 forces.

[00:43:45.38]

All right.

[00:43:46.34]

Now what happens now in the long run? We're all dead, as we know. So maybe we care more about what happens in the short run.

[00:43:52.02]

OK.

[00:43:52.52]

So in the short run, let's think about something happening that increases the number of new technologies that are introduced. So we have a jump from M to M prime. So we're going to have this burst of new technologies introduced at a given date. Now you can see, well, now this M is going to be in there. And again, it just depends on what share of GDP is produced with high-skill labor versus low-skill labor. So what you can show here is that if you have a burst of new technologies arriving, then more of GDP is going to be produced with high-skill labor because

high-skill labor complements these new technologies. And therefore, the college wage premium is going to be temporarily higher. And just to cut to the chase, when I say temporarily, I mean for like 60 years. Okay. So now the big question is, what are these M's, Z's, and alphas? And I'm going to show you that in the data.

[00:44:53.24]

All right.

[00:44:53.38]

So we're going to show you 3 empirical facts, M, alpha, and Z in that order.

[00:45:01.22]

Okay.

[00:45:01.32]

How do we get them? So there's kind of a trick that we took from some of our prior work. So the key problem in this literature is that you don't— it's kind of hard to know how to generate data on new technologies in general. So what's the trick? So when inventors create a new technology, they also have to create the language used to refer to it. So think about some technical change at the extensive margin. So for example, somebody invents the mobile phone. OK, we didn't have mobile phones before. We did have phones before, but we didn't have mobile phones before. So when mobile phones arise, inventors start putting phrases to describe this new thing that didn't exist before. One of them is mobile phone, another might be cell phone, and other kind of like cell tower. And there's going to be a set of phrases that are associated with this new technology. And we need them because the technology didn't exist before. OK? So the key idea in this paper is that we're going to take all the language in USPTO patents and we're going to deduct all the phrases that were known before these patents were written.

[00:46:19.52]

And what we're left with is a set of phrases that inventors invented to refer to the things that they invent. And then we group them into— we use Wikipedia in kind of ways I'm not going to get into here to basically group these new phrases into what we call technologies. So mobile phone is going to be a technology that is described by a Wikipedia page and that's going to collect all the new phrases that relate to the mobile phone. So using this methodology, you get like for this paper 6,259 new technologies that emerged between 1941 and 2006. Now with this new— now

the key thing to realize is that once I have a list of phrases that are associated with a new technology, then I can generate data on that new technology. by looking first of all in patents, like who mentioned it first. So where did it come from? Who invented it? And when was it invented? And then later I can look in the job market, I can look at job postings and see who's using it. And that's the— okay, so that's the whole trick. Now once you have that, I'm going to use that to kind of tell you about m, alpha, and z.

[00:47:29.46]

And so once I identify the new language that people need to invent to refer to newly invented technologies. I can track them in all kinds of corpuses of text. Here I'm going to track them in patents and job postings to tell you about when did the technologies arise, where did they arise, and who's using them. Now, the first kind of surprising thing is if you plot M of B using this method, you get this picture. So what this is showing you is how many new technologies were invented in 1940, 1950, '60, Okay, nothing happening. And boom, we have this thing here happening between 1985 and 1995 where there's this massive increase in the number of new technologies arriving. Some of this is IT, some of this is biotech. You can look at this in whichever way that you want. You can imagine that we spent like many, many months like plotting this thing using different assumptions and different cuts of data and trying to figure out what's going on. The answer is we just invented a whole lot of new stuff. between 1985 and 1995, and then it went back down. Okay, so remember I said when m goes up temporarily, something's going to happen in the labor market.

[00:48:43.27]

Okay, so that's why I'm showing you this. So this is m, the first star of the show. The second star of the show is alpha. And alpha— this is actually also already in our prior paper, but now we have an interpretation for In terms of the model. So what does this show? So I'm going— so now I'm going to my corpus of job postings. And for each job, I make a list of the technologies that are mentioned in the job posting. So for example, it might say, your job is to stand there with a mobile phone and record your customers— with a smartphone and record your customers' responses. So job postings are not very long. But if an employer really needs you to be able to use a specific thing or produce a specific thing, then they are going to write it in there. And so for each job, I'm going to assign it the age of the technologies that are mentioned in the job posting. And then I'm going to look whether that job posting requires a college degree or not, 1 or 0. Does it require a college degree? And what this shows you is this is just showing you skill-biased technical change in a novel way.

[00:49:52.38]

So skill bias technical change here, you can see that newly born technologies, 57% of them— that's the intercept here. So at age 0, 57% of new technologies require a college degree on average. And this is remarkably stable throughout history. You can sort of see if I only used technologies older than 50 years and estimated the line, I would get the same one. So now the other thing I want you to look at is the axis here. Okay, so it ends at 25. So what does this say? It says at age 70, at age 70, job postings that mention 70-year-old technologies, on average 35% of them require a college degree. What's average college attainment in the United States? Well, like 32%, something like that. So what this graph is telling you is that new technologies disproportionately generate employment opportunities for high-skilled workers for something like 80 years after they were invented. So when we go through a period of new inventions, we're going to have to expect that for a long time these new inventions are going to generate employment opportunities for high-skilled workers. Now, I'm going to use this line here to estimate my alpha.

[00:51:16.17]

Remember, I wanted to get to the alpha. And it turns out you can solve the model. The model says h over L is a function of alpha and the wage rate. And you plug that in, you get the alpha. So that's the second star of the show, the alpha. So we have an empirical estimate of how the skill bias of new technologies changes with the age of the technology. So now the third star of the show, and this maybe is going to lower your blood pressure a little bit because we're all talking about everybody's adopting AI. So what is this? So this is now H plus L of a given age divided by H plus L. So this is what I'm going into job postings and I'm showing you the market share, so to speak. So of all the job postings, How— so what is the labor demand of technologies of different ages, basically? And so I'm basically kind of looking at what share of job postings— what share of job postings in the US labor market mention 20-year-old technologies versus 35-year-old technologies versus 70-year-old technologies? So this is sort of interesting in its own right.

[00:52:32.23]

So the first thing that might lower your blood pressure, it takes 35 years for the average technology to be fully deployed on the US labor market. And then it gets slowly replaced over time. That's what this says. So technical change in the labor market really takes decades. And I know everybody's freaking out about AI. But the place where you're really seeing this, in my opinion— this is now outside the paper— but where you're really seeing this outside of outside of sort of I have my chat window open, is big corporations trying to automate their marketing department away, right? And they switch on the— or their customer service department away. And they switch it on and nothing works. It's basically like trying to use SAP for the first time.

You switch it on and then you switch it right back off because it's really hard to implement these things at scale. And that's where I kind of see that in the AI way. OK, but now back to the paper. So if I have now the h of u plus l of u divided by h plus l, I can look— ask my model, can I back out the z?

[00:53:34.59]

And the answer is yes. And this is kind of like in IO papers, like market shares and productivity are always closely related. So essentially, that's the same trick here. So by looking at the market share of technologies of different ages in job postings, I can back out how productive are technologies of different ages. So that's the z of u object.

[00:53:56.56]

Okay. How much time do I have?

[00:54:00.06]

Okay.

[00:54:00.17]

So now let's— we've done all this work. Now let me give you the payoff. Okay. So I'm going to plug the alpha, the z, and the m. I'm going to plug that into my model and run it from 1940 onwards. And what you get— I'm going to ask the model, what is the skill premium? What is the college wage premium? And it gives me this. Okay. So that fits pretty well. In brackets, I'm holding skill supply constant here. It would look much worse if I didn't. But that's sort of well known about this problem in general. But look at this. I mean, this is fairly— I'm not targeting anything aside from this point here. And you see that the run-up in the skill premium and then its flattening is perfectly predicted by the model. Why? Well, because this z function is so delayed and the alpha function is so slow that essentially the IT revolution from 1990 really only fully had its effect on wage inequality by 2010. That's what this says.

[00:55:05.31]

Okay.

[00:55:06.09]

So now, because this is a model, we can now also predict the future. And this is like, if everything kind of stays the way it was, we're going to have a slow decline in wage inequality over the next kind of, I don't know, many decades. If you put an AI bump, it's going to do this, by the way. But it kind of doesn't matter because it's far in the future. All right, so now you might already be thinking, I can now write the same model for a city instead of the whole country. And there's interesting facts about the skill premium to try and understand across space. So a key fact is, like, we are— so if you now kind of plot the college wage premium over the density of the population, we're in a place like here. Washington, DC is fairly densely populated. The skill premium basically today is much higher than it was in 1980, but it also got much steeper. So this means that the skill premium rose much more in places that were more densely populated. OK, now we're going to add a small thing here to our model, which is the diffusion probabilities.

[00:56:22.29]

We're going to basically— it turns out in the data, if you look, where are the new technologies born? Think about it for 15 seconds. They're born in very densely populated places like San Francisco and Boston and New York. And then you can estimate using the job posting data again, how long does it take to get to Oshkosh, Wisconsin? Okay. And we can use that to estimate a diffusion probability, like how many years does it take for the technology to go from New York to Oshkosh? And here's kind of the answer. So just look at like, so this is now my kind of my z function plotted for for densely populated places red and less densely populated places blue. And you can go see there's sort of like time travel going on.

[00:57:10.39]

Yeah.

[00:57:11.24]

So in other words, if you look out the window here, like, you know, we're like, DC is sort of similar to Boston. So if you look out the window here and you look at the jobs that are being advertised, the average age of technologies mentioned in job postings here is like 36 years. If you go to Bismarck, North Dakota, the average age of technologies advertised is something like 48 years. So it's like time travel going from— and as a result of that, you're going to get differences in the skill premium across space that are commensurate. Long story short, I'm targeting this one dot and the model replicates this like a charm. OK. And then finally, There's this other kind of puzzle in labor economics, which is, you know, like when the skill premium went up, it first went up for young people. So now my model said, well, there's— or not my model, like I estimated in the data that college-educated workers have a comparative advantage

in using young technologies. That was the alpha function, right? The alpha function was downward sloping. Now you can say, well, modify the model and say, well, it's young educated workers that have a comparative advantage in using new technologies.

[00:58:28.08]

And that rings true to me. I was using Skype until like 2 months ago and my grad students were making fun of me and then they shut it off and now I have to use something else. So this is kind of like the graphical representation of this. The rise in the skill premium first happened for young workers. in the 1980s, you can add this to the model by just kind of saying, okay, now I'm gonna have not just the age of the technology but also the age of the worker. And then you get, just take a mental snapshot of this, you get that. Okay, so basically it looks like this idea that the skill, sorry, that young people have a comparative advantage of using new technologies, explains a lot about the data here. So let me conclude. So the invention of new technologies and their diffusion in the labor market offer a compelling explanation for wage inequality and how it's evolved both across skill groups and across space and across age groups in the United States. We use kind of this novel text-based approach to estimate the M, the alpha, and the z functions. We found this extraordinary surge of new technologies between 1980 and 1995.

[00:59:46.19]

A technology's share of jobs peaks 35 years after its emergence. And by the way, if you plot this for GenAI, GenAI is spreading much faster than the average technology in our sample.

[01:00:01.01]

But we'll see.

[01:00:06.30]

OK. A simple model in which young educated workers have a comparative advantage in using newer technologies explains variation in the skill premium across time, space, and age groups. That's it. Thank you so much.

[01:00:27.33] - Antonio Spilimbergo

Thank you very much.

[01:00:28.41]

This was very, very interesting as well.

[01:00:31.16]

We have a 3rd paper by Julian Jacobs. of Google DeepMind. He will be online, and unfortunately he will not be here for questions and answers. So, okay, the floor is yours.

[01:00:50.36]

I think the link—

[01:00:52.07]

the paper is titled, Do US Retraining Programs Protect Against Technological automation.

[01:01:02.09] - Julian Jacobs

Yeah, thank you. And so sorry that I can't be there in person today. But yeah, really appreciate doing this virtually. So the title of this paper is, Did US Worker Retraining Reduce Participant Automation Exposure? I'm an economist at DeepMind, also based at Oxford. And I did this paper with my co-author Jordan Kennedy. And really what we wanted to do was understand, looking at— taking a kind of backward-looking historical lens, how had the US's primary public worker retraining program supported upskilling for people and resilience in the face of digitally enabled automation. So that's the kind of core question that we're looking at. And so to do this, we looked at the US's primary active labor market program, which is WIOA, and specifically this question about the extent to which it insulated people from automation. So with this in mind, I would maybe flag 3 reasons why I think this is a kind of interesting question to study. So on the one hand, Worker retraining is quite an interesting theoretical concept when we think about potential policy responses to AI-enabled automation, future-looking AI-enabled automation, specifically because it might help smooth labor market frictions that might emerge if there's a sort of— there's particular demands for workers with a particular skill set, and those skill sets are not necessarily spontaneously prov ided.

[01:02:37.31]

Okay. Can you see my slides? Yeah. Okay. I think you can see my slides anyway. Yeah. So, you know, at the same time, there are also some barriers theoretically that we might expect to active labor market programs. So what might those be? Well, for one thing, there are participant barriers to actually joining these programs, and that involves costs that people incur. And I'm not talking about actually the cost of the program, but I'm talking about transportation to the

program, the opportunity cost, for instance. Childcare can be a problem as well. And another issue with it is there's a significant heterogeneity in program quality. especially across the United States, as we'll see. And third, there's a kind of more foundational issue, which is there may just be a structural mismatch. So, for instance, it may be that as technological change moves forward, it creates demand for certain skill sets or reduces demand for certain skill sets at a faster rate than new jobs are generated. And retraining is— it's not obvious that retraining would be well suited to create jobs where they might not otherwise exist. So we're looking at WIOA, but it's worth talking a little bit about, you know, what the United States' recent history has been.

[01:04:05.33]

So at a high level, across decades now and multiple iterations of active labor market programming, some of which place more of a role on retraining, some of which place less of a role on retraining, there has been a significant effort to basically show lasting positive effects of active labor market programming and retraining. Despite that, the evidence consistently shows mixed or, in some cases, negative results. The most, I think, rigorous attempt to do this was the Workforce Investment Act ran a national RCT. It found that counseling was most helpful in supporting outcomes, but the effects of actual worker retraining or null at 30 months. Now, with WIOA, we do have high-level results that show 70% of people are employed post-exit, but we don't have a control group. And so, you know, one thing that I'll say about the research that we did in this piece is we have to be very epistemically humble here. This is descriptive work showing what has happened. We can't necessarily say with any degree of certainty what the role of WIOA is specifically in inducing any of these particular outcomes. And indeed, there's plenty of reasons to think that there's incredibly strong non-random self-selection and variation in program quality that are confounding results and actually driving outcomes.

[01:05:31.32]

And the estimation strategy that we've taken here cannot credibly sort of get around those issues. So what does WIOA look like? Well, there's approximately 500,000 people that participate in WIOA's adult and dislocated worker streams in a given year. This amounts to a total funding of $500 million, which is actually quite small in total. So, you know, it is worth saying as well that when we think about, you know, the broader question of the United States' policy response to technological change and retraining as part of that, It's not as if this scheme currently is particularly well-funded. But furthermore, the actual program intervention is patchwork. It's highly, highly decentralized with state and local workforce development boards actually administering these programs. And there's kind of a 2-tier system here as well, where there's light-touch search assistance, which is what the vast majority of people receive. And then there's heavier-touch formal retraining, which is, again, a kind of small self-selecting

minority of people that participate in these programs. Now, just a preview of kind of, I would say, the 3 core findings of this PSR. So the first is that a person's starting point, so their prior occupation, their prior industry, as well as their local economy and geography, those seem to be the things that moderate outcomes most.

[01:06:59.13]

So success seems to be at least partially conditioned by, again, a person's— what a person did before they joined WIOA, and then their local workforce development board, and their age is also a factor as well. We find that incidence of success in employer-led programming is basically most positive compared to other groups. That is in line with other literature. But the thing to say here is that very few people actually participate in employer-led retraining. Only around 5— as we'll see, only around 5% of people receive formal retraining, period. And a small subset of those actually receive either apprenticeships or an employer-led training model. So we conclude that WIOA acts primarily as a kind of second chance program. It provides a wage floor, can help people bounce back. from periods of perhaps unemployment or particular vulnerability, it has not served as a means of consistent upskilling. So potentially not well suited to deal with large-scale productivity shocks through AI or digitalization, let's say. So to do this study, we looked at around 23 million administrative records. From 2017 to 2024, this is what all we could get our hands on using publicly available data. This comprises around 23 million unique completed participation records, and we exclude people with missing exit dates.

[01:08:33.25]

So that does exert a kind of another bias as well, which is, you know, we're not necessarily seeing cases where people just dropped out of the program entirely. We're seeing only instances in which people completed the entire program end to end. And then we look at the relationship between these outcomes, and we fit in a measure of routine task intensity, which I'll talk about momentarily. So the kind of core measure that we use in this piece is we use a concept of a retrainability index, which is basically decomposed of 2 parts. It's a very simple measure. It's wages comprising 50% of how it's calculated, and then the other 50% of it is decomposed of 2 pieces, which is routine cognitive intensity and routine manual intensity. Here we're considering a successful outcome any case in which routinization decreases. So there's less— a person exits into a less routine cognitive or less routine manual job. And again, you know, this is kind of weighted by wages. So did a person experience at the end of this program an increase or a decrease or kind of stagnation of their wages after participating in the program? So that's basically our measure of success.

[01:09:58.53]

We also, throughout the paper, you know, this is kind of a nice sort of tidy way for us to understand generally how we think about this program as a tool to support upskilling in reaction to technological change. We decompose— we look at each of these individual components throughout the entire paper separately, of course, as well. So one question that you might ask is, why are we looking at routine task intensity instead of AI exposure? And this is where I would say that our measure is backward-looking. We're looking at RTI primarily because we think it serves as a reasonably good proxy for the computerization-driven automation that occurred during 2017 to 2023, right? As other people today have spoken about, AI's effects are only beginning to separate within data. Certainly not something that we would have expected to manifest in data as early as 2023 or indeed before then. Now, we also think that AI exposure can't necessarily help us distinguish between complementarity and displacement. This is a kind of longstanding problem. with that as a measure. And also, you know, traditionally we've found that AI exposure and others have found that AI exposure is concentrated in highly paid sectors where WIOA rarely operates.

[01:11:19.44]

So it's not clear why that would be the right measure to use. That being said, I do note that a concurrent paper by Ben Hyman and others does look much more clearly at AI exposure. And I think they find some very interesting results. about how WIOA can sort of insulate people who, you know, are disproportionately exposed to AI today. And so it finds, like, slightly more auspicious results, I would say. So in total, we find that 83% of people, just over 83% of people, enter the program unemployed. And this level of, like, employment detachment deepens as people get older. So the share of people that enter unemployed rises as people get older. So I think this speaks to 2 things. One, the general vulnerability of people as they enter the program. And then number 2, the fact that this vulnerability deepens as people— as the age with which people participate rises. I mentioned earlier that just a small handful of people actually receive formal retraining. So that is, like, one of the, I think, most important things to note when we talk about WIOA is we often think about it as the United States' primary public active labor market program and retraining scheme, yet only a very, very small fraction of people actually receive the benefits of formal retraining for a number of reasons.

[01:12:48.31]

So here we find just about 5% of people. receive formal retraining. We also find that formal retraining recipients differ, and this will be a kind of recurring finding throughout this paper. We find that people that did receive formal retraining differ very significantly systemically from people that didn't receive formal retraining. So there's a strong— we think there's basically a strong self-selection bias here as well. So for instance, they tend to be like less educated. That's

what we find here, and this is true across core demographic categories. When we look at the specific movements of people across occupations, we find that the top subsectors that people started in primarily, you know, 16.9% of people in administrative and support services, followed by food services. Top occupations: customer service representatives, cashiers, laborers, and material movers. Post-training, we find a move— actually quite a marked move into healthcare, which is quite interesting, as well as trucking and transportation generally. So, you know, on the surface, when we initially did a kind of like a descriptive analysis of this piece, we thought that, okay, this might provide some evidence of potential upskilling. This is in line with prior literature that suggests that WIOA has moved a sizable share of people into healthcare and related industries.

[01:14:20.04]

Yet when we look at broader results, you know, I think that things become a little bit more complicated. So we used an XPBoost model here to look at the specific factors that seem to most correspond with programmatic success. We found that pre-program occupation is a single single biggest predictor of someone successfully— basically receiving a successful outcome using the retrainability index that we made. Now, that being said, this model only captures a small— there's substantial unobserved heterogeneity that our model doesn't capture. And so that's something that we absolutely need to note here. Meanwhile, at the same time, you know, industry as well, similarly, kind of second most important factor in predicting outcomes. And then I— we would also flag the local workforce development board as being significant as well. So in other words, we find again that what a person did prior to joining— to joining WIOA or participating in WIOA, as well as the constitution of their local environment, the possibility availability of jobs, for instance, as well as possibly the management of their local workforce development board and the nature of the retraining program that they participate in most significantly affects outcomes. But again, even within this, this can only explain, you know, a relatively small portion of the outcomes that we see.

[01:16:03.53]

Now, when we find successful outcomes, we find that a majority of people do see an increase in wages after participation in the program. So about 60% of people that did not receive formal retraining see an increase in their wages after participation. The number— it's 76% of people that receive— see an increase in wages who participated informal retraining. That being said, the story is very different for our measure of routine task intensity. When we look at movements in RTI, we see that the vast, vast majority of people are not seeing any change in RTI. They're seeing either no change or an increase in the routinization of their job. So we find basically that successful outcomes as measured by this index are driven primarily by wage gains, not

necessarily changes in occupation or industry, not necessarily changes in the level of routinization. or susceptibility to digitally enabled automation. Now, when we look more closely at what's driving this wage differential, one hypothesis that we wanted to test was basically whether or not people are simply experiencing a kind of bounce back. So often what we find in the literature is that people who enter WIOA do so from a period of very— of a high level of vulnerability.

[01:17:35.02]

And we wanted to know basically whether or not WIOA was primarily serving as a vehicle for wage bounce-back and rehabilitation as opposed to actually, like, boosting wages beyond what they might not have otherwise been had they not participated in the programming. We find that people that joined formal training programs had lower wages when they entered. When they exited, the gap there— they diverged from the rest of the group. So on the one hand, this might seem to suggest that formal training was disproportionately likely to boost wage outcomes. We explain this by a simple Ashenfelder dip, where actually it's just the most vulnerable people are the people that participate in formal training programs. And those programs help them bounce back and catch up. Now, I mentioned earlier that the most successful programs, at least that we saw in this study, were registered apprenticeships, customized training programs. Again, we can't say for certain at all whether or not this is driven by inherent qualities about these particular kinds of programs. or strong self-selection biases. It's probably a combination of the 2, but we note that this is roughly in line with other literature from people like Lawrence Katz, for example, that have suggested that programs that involve employers in the process of retraining do tend to experience higher incidence of success.

[01:19:06.07]

We also— I also will note here that we have subsequent work using polling data and conjoint experiments to understand preferences around retraining. And we found, looking at that, that the public in the United States overwhelmingly favor and demand employer-led retraining and apprenticeships— apprenticeship models, as opposed to classroom-based training. So there's an interesting marriage of what data suggests seems to correspond with the highest incidence of success and what the public seems to demand within polling data. Equally, we find some loose evidence of potential catch-up across geography as well, with incidence of success most pronounced in areas with consistently lower wages and generally worse-off economic conditions. I'm very happy to talk about that over questions as well, but won't spend a whole lot of time on this right now. And I mentioned earlier that age is a factor as well in predicting program outcomes. So as people get older, their likelihood of benefiting from these programs, as measured by this index, declines over time in essentially a linear fashion. Whether or not a

person is employed is also related. to their likelihood of succeeding in— through our measure. So people that go into the program employed with a preexisting level of labor attachment are more likely to experience programmatic success.

[01:20:44.10]

And I mentioned earlier that people that are coming from, you know, a lower-income background, which, to be fair, is the vast, vast majority of people that participate in this program, are similarly more likely to experience a successful outcome. And again, that is possibly driven by catch-up gains more than anything. And WIOA's status is kind of providing a baseline floor for people as opposed to really serving as a tool to generate significant upskilling capacity. And, you know, I think that this is also borne out in the education data as well. So people with advanced degrees that participate in the program, it's possibly on the surface a little bit surprising that the incidence of success there would be as low as it is, but I think that partially speaks to the kind of baseline floor function that WIOA provides, number one. And I also— and on the other hand, I also think it speaks to the self-selecting sample of people with advanced degrees that would be participating in WIOA. So briefly, I think that this work— again, I want to be super humble about what this says and what it doesn't say. This is just a kind of descriptive piece looking at what WIOA evidence suggests today.

[01:22:02.20]

There's not— we can't say anything causally about the exact effects of these programs, but we nonetheless think it's interesting to look at the characteristics, programmatically, individually, and environmentally that correspond with incidents of success. So we think that there's 4 implications of this work. One is we do take this research, along with work that others have done, as providing some further support for private sector partnerships and apprenticeships as being a model worth exploring further. A key question there is about the generalizability of these programs. So typically very small numbers of people. are actually joining these interventions. There's a broader question about how widespread could those programs actually be disseminated, how many people could benefit from them. The second is, I think that we clearly need better causal evidence and evidence that can be used for a rigorous causal identification strategy. So this could, of course, take the form of randomized controlled trials on retraining programming. I equally think that there's a great opportunity to look at cross-national data to better understand not just works— what works in the United States, but what works in Denmark, what works in Germany, and a whole host of other contexts to better understand perhaps core characteristics of effective retraining or active labor market programming.

[01:23:33.03]

And I also think that there's an opportunity for us to think about how to help participants navigate the over 75,000 programs that WIOA currently has on offering. People like David Deming have rightly, I think, pointed out that this system can be bewildering to a lot of people and impose barriers to people's ability to navigate that system. And so I think this is actually a potential opportunity for AI to support participant navigation of the WIOA system. And finally, I think that it is also, like, worth thinking about the limits of active labor market programs and where we might want to explore other sorts of interventions. Where might active labor markets fail? Where might other kinds of interventions be most helpful in supporting the AI transition? Thank you.

[01:24:30.34] - Antonio Spilimbergo

Excellent, Julian.

[01:24:31.47]

Thank you very much.

[01:24:33.22]

We have our discussant, Tito Boeri, Professor Tito Boeri, who will discuss 3 papers in 12 minutes.

[01:24:46.23] - Tito Boeri

So let me thank Peterson, IMF, OECD, and World Bank for inviting me to this very enriching conference whose, I think, broad purpose is to sign the exposure. So my slides are not yet on, so please put on the slides. Because I think that a lot of progress has been made in the last few years to measure exposure to AI, but policy needs more than that, needs to understand whether these technologies are complementary or substitute labor, so signing this exposure, and possibly also identify causal effects, because also these are very important for policy. And I think that the 3 papers I was given to read really do provide an important contribution in this respect. A nice thing about this paper is that they do cover different phases, different stages of a process. When a new technology arrives, there is creation of new firms, and then there is a standardization taking place. That's stage 1. Stage 2 is the fact that you have a change in the skill demand. So labor demand changes, the composition of labor demand, and there are clearly both people benefiting from this, gaining, and people losing from this.

[01:26:20.44]

And there is a time involved in this. And that's stage 2. Then stage 3 is when there are different educational choices. So students re-optimize their choices based on the labor demand modifying. So choices of majors, of courses, graduate schools. And that's stage 3. And stage 4 is when displaced workers adjust. and there is retraining and occupational. So I think the paper by— and I would discuss them in this sequence rather than in the one they were presented today— the paper by Hassan, Caliani, and Restrepo covers very much stages 1 and 2 and covers the period between 1970, but there are also predictions going up to 2080. Dodini, Smith, and Hsu instead cover stages 2 and 3. And finally, Jack, Jacobs, Kennedy cover stage 4. I think the 3 authors did a very good job in presenting, so I won't need to provide a summary of the papers in case— but, you know, in this case I would need another 5 minutes, but I think Antonio is not going to give me those time. So I will address 3 key issues that correspond broadly also to these different stages of a process. The first one is who is exposed?

[01:27:44.01]

And clearly, given that my role is a discussant, I will try to point out problematic issues in the 3 papers on measurement. The second is who benefits and who loses? And above all, you know, the problem of the sign and the causality. And thirdly, who— what is happening to the adjustment? You know, the potential losers are adjusting and policies help. And finally, I will say something a bit heroically in trying to draw some tentative policy implications. So who is exposed? The paper by H.K.R. uses Wikipedia to separate technology vintages and the bigrams, a sequence of 2 words, to measure the intensity of different technologies across patents as well as job posting. The paper by DSX looks at job posting again, but also major choices, and uses some given measure of exposure— I will say a few words about those— to measure what is the impact of these technologies. And finally, The paper by JNC instead uses a different measure. It's a measure of routine task intensity, which is more a metric for computerization than for AI. Clearly, the problems of measurement here are, one, first of all, that this time is different, as also stated clearly by Julian in his presentation.

[01:29:24.13]

Another simple feature of a new technology that is different from previous vintages is the fact that Here is a technology that not only requires a new language to be described, the problem is this is a technology that draws on language. So that's a kind of different story. Also in the DSX, there is some problematic aggregation from task to occupation to bundles of occupation and then also to course level. Clearly, all of these aggregations are somewhat problematic. And finally, in JAC, the fact that computerization is not AI. JAC did do a very interesting job in defining this retraining index, which is a performance metric. It's more or less like a scorecard about how participants fared. It's not so much informative about which displaced workers are

needing AI-intensive retraining. And they were extremely honest in stating that they don't do any causal effect. It is very descriptive work. Yet I think it would be nice to see and try and possibly go beyond that. And clearly, the strategy is to find a proper control group. There are methods to do that in matching with people who did not receive the program assistance. But that's something I will also talk about when we discuss about causality.

[01:30:59.06]

There is also another problem, which is the fact that the history of a shock is relatively short. Also on that, I think something can be done because some AI tools were already in use before ChatGPT. Think about Copilot, which was available in June '22. And, you know, it is indeed interesting to see that this is when DSX find in their paper that there was some decline that was taking place too early. Perhaps it was because of that. So in a way, one has to find this canary in the coal mine groups to prolong this series. Finally, an issue that was discussed also yesterday. Job postings are not hires. So Lai-Cast records provide some description of what is asked for, clearly not what is done. So I think it would be important to try as much as possible, and there has been attempts in this direction to link job posting data with administrative data on actual hirings. The second issue, who benefits and who In the past waves of vintages, in the H.A.K.R. paper shows that it was the young, educated workers who were the early adopters. And the premium, the skill premium rose first, and then, you know, for this, and also in dense cities.

[01:32:33.52]

cities. And then it was later on that older workers were catching up. And clearly technologies aged and there was a routinization at that stage. In the case of the AI, what we see from the DSX paper is that instead it is the young educated entrants who are the first to lose. And now there is this decline in earnings. which is quite large. And it is also consistent with some evidence that adopting firms, which is also provided in their paper, cut junior hiring while holding senior jobs. So these two may seem to be in contradiction. While they may not be contradictory findings, Because AI substitutes for codified knowledge, which is actually what the new entrants bring, may be instead complementary to tacit experience. And there is also some evidence in this respect. For instance, work by Björn von Sonn and others is pointing in this direction. So that, I think, is a very crucial issue. So I think further work should really devote attention to this, also because we know that there are these scarring effects. So this is an issue that I think it's important to address very thoroughly. So a number of questions that are raised by comparing these studies is that, do these AI-linked job postings and new occupational titles skew young and college-educated?

[01:34:18.08]

Do they carry a wage premium decaying with vintage as it gets standardized? Is the decline in entry-level hiring for tasks that are being automated by AI— is there a young worker advantage in using AI on the job once they have a job? Is it outweighed by the hiring decline in exposed entry-level roles? I think these are all very, very relevant important question for the future and future work. On causality, clearly calibration from an aggregate time series is not clearly causality, and it's not a test matching. There is also this 2/3 of the demand shift which is residual. I think it would be nice if they could tell us which parameters— you know, there is also this battle between z and alpha. Which parameters of a model would change with AI, and what would happen to the wage premium? In the case of the DSX paper, there are problems in using this shift-share instrument insofar as the exposure measure, which is based on cloud use, is clearly endogenous. It was done also in 2020 before. It is true that some of the results also with a human-rated exposure measure, but not always this is the case.

[01:35:47.01]

And also, as also the authors acknowledge, there are many confounders. Remote work is— some of them are discussed, but I think there are others as well, more recent, like the federal layoffs, you know, some people being fired during these layoffs may be competing with younger people. And also the Trump uncertainty, I think, is a factor that should be taken into account here. So here it would be nice to check the parallel trend assumption with longer series before COVID possibly. And fortunately, some of the series like the LIHCAS are available since 2010. There are also some further identification identification problem, this tech hiring bust in DSX. JC, as I already mentioned before, there is— one should try and find as much as possible some counterfactual. So perhaps matching to nearest neighbor similarly exposed to AI but not receiving the training would be a good strategy. I've seen some work by Hyman and others that is doing something along these lines. Otherwise, you know, there are all of these problems, also the Eschenfelder dip type of thing. So you get that you have this increase before and then decline afterward, which is simply a mean reversion effect.

[01:37:16.00]

If you do matching, that could help out in taking away this type of thing. of things.

[01:37:22.41]

Third issue, adjustment of people and then the adjustment of policies. So the question that comes out from reading the paper by DSX is, why don't students leave the exposure? And in particular, it was in the previous version, the one I read. I see that they have done some more

progress today in the presentation. I learned more about it. But the explanation they provide for why they keep on choosing majors that are subject to declines because they're exposed to AI is that there is a lower effort cost. Another possible explanation is that there are weak signals. Given that the shock is relatively recent, there may be some noisy signals that students have in choosing majors. And in some cases, it also may be rational to wait and just see what happens, you know, over time. So clearly, this is very important because we have to understand whether they're making a bad investment or is simply option value here. In terms of the policy adjustment, Clearly, WIO was developed for other shocks, and clearly this is a different one. And therefore, there will be some needs to adjust this program to the new situation. And some of the considerations that were made yesterday by Gordon Hanson, I think, apply also to this case.

[01:39:03.01]

So as I said, In any event, let me just see what are possible directions for policy discussion. First of all, one should provide good measure before targeting. So one should link education, earnings, and benefit records in real time, what the Texas data do. But, you know, and in a way, AI could be useful in this respect. So use AI to monitor and predict AI. Second, we really have to focus very much on protecting the entry ladder. So because early scarring is long-lasting, so it is now the time to worry and think about about how to do it and to address this. Training on AI-complementary skills as much as possible. So this paper by Hyman I was mentioning before is useful in this respect. And apprenticeship may be a good way to provide experience to entrants. Pair training with insurance, it's very important. So Social Security should be able to react and provide early warning systems and address the issue. And AI, again, can be very helpful in this respect. And finally, inform students. So publish outcomes by field, redesign grading and degrees to keep the signal value. That's very important, I think, for the university teaching, so the work all of us do.

[01:40:48.32]

So just to take away, past technologies took a decade to absorb. A distinctive feature of AI is the speed. And I think all the evidence brought about in this conference suggests that the adjustment of our institution is relatively slow. So that's something to worry about. I will leave here just the last slide for the authors. are suggesting for them specifically.

[01:41:32.40] - Antonio Spilimbergo

We start the questions, Q&A questions. I have 2 questions.

[01:41:45.56]

Yes.

[01:41:48.05]

I have 2 questions.

[01:41:49.48]

One for the first paper, for Tucker. I was thinking the period before the introduction of ChatGPT was a particular period where it was just the exit of the pandemic where people were hiring a lot of AI. experts because working from home was becoming normal.

[01:42:11.11]

And so you needed new technology, new colleagues were helping with this. So maybe companies hire or even overhire just before the introduction of ChatGPT. So the identification, looking before and after, could be a little bit problematic. This is a little point. And the second point, the second question is actually for Tarek. You talk about new technology, but is really new technologies or new products? In the sense that new products like mobile telephone go into the utility function, new technology go into the production function, and these have different implications for the way in which we think. And another thought, we focus a lot on labor market. We don't focus on product, not only because it's not in the title of the conference, but also because we have much less research and tools to deal with product market, AI and new products.

[01:43:18.51]

And maybe your database could help that.

[01:43:22.02]

Just a consideration. Well, we gather A few questions, but if you want to answer.

[01:43:28.57] - Tucker Smith

Oh, I'd rather—

[01:43:33.46] - Audience

Yeah. I think— ah, voilà. I have a question for each of the panelists. For Tucker, I found it super interesting, you know, how you tried to look at change of fields versus upskilling in the same field. One question I had is whether upskilling would work better in fields that are less exposed. Because when you think about the use of AI, I mean, you still need to be able to assess whether the outputs it's giving you, you know, are reasonable and make sense. You need to be able to design prompts that guide properly the AI. So I wonder if you would be able to test that with the data you have and whether that would give different results than what you find for upskilling within very exposed fields. For Tarek, you started extrapolating a little bit about how AI could be different, talking about the much more rapid diffusion of AI. And some evidence was presented yesterday that actually in the US, 5% of job postings already require an AI-related skill, which is much higher than what you showed for other technologies. But I was curious to hear what your take is on how skill-intensive you think AI is or will be, because that was an important parameter in determining the evolution of the skill premium.

[01:45:12.56]

And for Tito, I wanted to ask, because you are an expert of European labor markets and you also had a role in policy, could you tell us a little bit, you know, what's the evidence in Europe on the effectiveness of retraining programs for the labor force? Thank you.

[01:45:41.22]

Hi, so I have 2 questions. One for the first presenter, Tucker, whether he also looked at older workers looking for retraining. We really focus here more workers that are graduating and workers that decide to stay longer, but I feel that as technologies are changing faster, we need to also find opportunities for older workers to come back 10 years or 15 years if you see uptake of that, or even less, workers are looking less. And then for Tariq, I was curious to know whether you focus a lot on the demand side with job postings, but what is the supply side? Because we know as well that college education has increased dramatically in the US, and also the immigrants that are staying in the US. So how do you take these into consideration in your framework? And maybe you can elaborate about what would happen if the supply of these workers go down. Thanks.

[01:46:54.34]

Thank you for the presentations. They were super interesting. One question for Tucker and then one for Tarek. For Tucker, I was wondering about the Texas application that you provide because in line with what Antonio was saying, my reading is that there were a lot of things

happening in Texas around this period. So the movement of many companies to Texas from California on the production side and also from the educational point of view.

[01:47:25.46]

From reading anecdotally, there seems to be a higher demand for Texas and southern colleges. So, I was wondering if in light of this, you think that your results are some kind of bound. I don't know if upper or lower bound of the effects of these technologies and the retraining effects you find. And for Tarek, I was wondering about if you have any views or thoughts about the fact that in Lightcast, you're only looking at job posts for the past 15 years or so.

[01:47:56.22]

So, I don't know if this introduces some kind of biases in terms of, for example, for older technologies, do we think that these diffusion rates and these parameters that you're extracting are very specific to this time period or you think they're generalized more to other periods?

[01:48:17.13]

And one comment also on the model, or not comment, but a question is, Here, I see this as more of a— the theory is allowing you to back down some of the parameters you're using then to provide these facts. I was wondering if you have any predictions in your model about what drives these different diffusion rates and so on, à la Asimov et al. Thank you.

[01:48:48.22] - Tucker Smith

Great.

[01:48:49.38]

Thank you all. Those are all great questions. Thank you, Antonio. I'll try and quickly answer some of them. So in terms of upskilling for less exposed majors, that's a really good point. So what we did there, like with the coursework level exercise, it was only considering upskilling as in just taking more of these more exposed courses. And so from that, I expected that maybe less exposed majors that had more computer science experience experience may benefit. And I figured that that may complement some of knowing how to use these tools better may help. It doesn't seem like it has so far, but it could be that also just having more expertise in your field and lower at less explored fields could be more helpful. So there's definitely more we can do there with the course-level data. We haven't seen— we haven't investigated yet about older workers retraining. One exercise we have done is showing that older workers they're not that

old, but still fairly recent grads that were in the labor market before this came out and before the pandemic have had kind of suppressed wage trajectories. And relative to this setting of Texas, a lot of things, you're right, are going on.

[01:49:58.55]

And in Texas, I think the economy has done quite well and featured this kind of tech growth. And so from that sense, It could be, I think, more of us on the magnitude upper bound or perhaps lower. I'm not sure about the direction of the bound, actually. But one thing is, I think one result related to that is we do see that actually the kind of selection on students into these more exposed majors has becoming more positive over time. And so if anything, this could be, in that sense, it could be a— upper bound or lower bound in magnitude that these are really smart students that are coming to these institutions that are taking these majors, and even they are having a tough time finding jobs?

[01:50:46.13] - Tarek Hassan

Hey, thanks. These are great questions. Let me maybe start from the bottom. So Samuel, an important question, of course, is whether these parameters that we estimate are stable. for the period of history that we can see, they are remarkably stable across different technology waves and so forth. So I'm going to say a tentative— OK, stable for all the history we can see, stable for the future, I'm not going to make any claims. Marina, the supply side of college, of course, the college wage premium is a massive puzzle in economics. And what makes it worse is that so many people went to college. So let me kind of say sort of like the overall college wage puzzle is deepened by the fact that we have more people going to college and the price of college-educated work went up at the same time. The model fits perfectly if I hold the supply constant. If I move the supply, it explains a third of the variation. Which is good because lots of other people also wrote papers explaining why the college wage premium went up. So we need to leave some room for other things going on. Okay, so now to the bigger— so now I think the other questions were all about Tito, like different versions of Tito's, is this time different?

[01:52:14.37]

Okay, so let me kind of say I am less confident. I would say that we have the least— the part of the data, the paper that we have the least data about is the age gradient. So that's sort of like really only— I can go into details, but that doesn't come out of our data organically, but it's more something that comes from census data and so forth that was about one specific thing, which was the IT revolution. And in the IT revolution, you know, like it's pretty clear that young people used computers first. We don't know about any of the other technology vintages. And I think that

is probably the part that maybe is the most different this time. But on other dimensions, it really doesn't look all that different in the sense that the skill— I just looked this up as we were talking. So the skill bias of the— it's actually in my data, it's 3% of job postings mention generative AI. it's massively educated people. It's 72% require a college degree and it's all pretty highbrow engineering type things. Some are applications and marketing and so forth, but most of them are pretty highbrow jobs.

[01:53:32.04]

So in that sense, it's spreading much faster in a way, but the general dynamic of the new technology is complementary to education, that is 100% there. Now let me say another thing, and this goes to Antonio's question. We have some other parts of this research program that have been looking at how do new technologies contribute to the development of more new technologies. And if you start thinking about this way, maybe saying it's spreading incredibly fast is somewhat misleading. Because really what the IT revolution, for example, was, it wasn't just like we invented the computer, but we invented a gazillion things. And so far we've seen one thing, which is like the chatbot on your Claude, you know, that we all know and love. But so we have strong evidence that, you know, technologies contribute to the production of new technologies. So from that perspective, we're probably just at the beginning of AI generating many new technologies that all, according to our framework, are all going to be skill-biased. So we're just going to have another hump into the wage inequality. And I think that's really the main thing to focus on from a policy perspective, that accelerated technology creation will increase wage inequality even more than we've experienced over the last 40 years.

[01:55:06.04]

And, and that's— so that's a challenge. I hope I covered everything.

[01:55:14.44] - Tito Boeri

Well, to answer Florent's question, I think the evidence about the effectiveness of retraining schemes in Europe is mixed. There is not, you know, very strong evidence. In general, what is found is that target policies seem to work better than broadly based. This may be good news for AI because AI allows better targeting of policies. So you know way more about potential applicants and so on. At the same time, clearly, not so easy to devise what to train for because clearly you really would need to do training on skills that are complementary to AI. And some of those training programs that have been used so far were not moving in that direction. Well, the discussion about this time is different. It would be very nice, but we have no time.

[01:56:20.44] - Antonio Spilimbergo

I think we have to wrap up.

[01:56:23.07]

Let me

[02:02:35.23] - Åsa Johansson

Interesting conference. So I welcome everybody back to this next session.

[02:02:43.08]

I give you a couple more minutes to sit down. Okay. So, let's kick off this 4th session as we have a a relatively tight schedule so we can keep on going.

[02:03:10.48]

So my name is Åsa Johansson, and I'm from the Economics Department at the OECD. And I'm very happy to be here to share this session on AI and its implications for emerging market and developing economies. So we have had a really, really rich discussion over the past day and a half, and I've learned a lot from the presentations.

[02:03:30.34]

So what I took, one thing that I took from many of the presentations, including the ones this morning, is that Often the discussion on the papers start out from a task-based approach, looking at how different occupations consist of different tasks and how these different tasks can be exposed to AI.

[02:03:48.28]

And that's a very useful way of looking at the implication of AI for labor markets.

[02:03:52.58]

Many of these papers have focused on, I would say, mostly advanced economies.

[02:03:57.16]

And it's possible that the different occupations can exist in emerging and developing markets. And also that the combination of tasks in those occupations are different from those in advanced economies, meaning that AI can have a differential impact in these economies. And it's also possible that the exposure to AI and implications on labor market also varies with how integrated countries are in global supply chains and also their position in global supply chains. So that's another interesting aspect, and all of these things can differ between advanced and emerging market economies.

[02:04:31.17]

This is exactly what this next session is going to look into.

[02:04:34.41]

We have a really, really great panel with 3 presenters. We first have Herman Winkler from the Inter-American Development Bank. Then we have Gabriel Marin from the University of Pennsylvania. Then the third speaker is Filippo Pontadini from LUISS University in Italy.

[02:04:51.15]

We also have a fantastic discussion with David Kaplan from the Inter-American Development Bank. The format is the same as the previous sessions, 25 minutes for each presenter and then a lead discussion for 10 minutes. I'll please invite the first presenter, Hernán Winkler, to the stage. Thank you.

[02:05:19.31] - Hernán Winkler

Hi. Hi, thanks, Asa. Thanks everyone for being here. I'm going to present this paper. This was joint work with Pavel Gomirek from the ILO and Marina Vilas from the World Bank. I worked on this paper when I was at the World Bank, and this was a background paper for the World Development Report 2026, The Promise of AI. And the question we tried to tackle here is if AI brings both disruption through displacement and benefits through productivity growth, will those also reach developing countries? And if so, will it be at the same speed? So what do we know about AI? We know a lot already. AI targets non-routine abstract tasks. These are the type of tasks that mostly more educated workers, white-collar workers do, which is different from previous earlier waves of automation. We also know from experimental evidence that that AI has productivity benefits in very specific occupations. And those benefits have shown to be benefiting lower ability or lower experienced workers the most. So that could have some inequality reducing impacts. And also there is evidence that AI is already affecting the real labor

market. And the evidence comes mostly from online platforms. We saw some of these papers today.

[02:06:42.59]

This is just a very brief literature review, but we already know that vacancies in exposed occupations are declining and young workers are being affected the most. And we also know from this very interesting paper from Betaya and Chem, more related to Al Guido, focused on international outsourcing through platforms, where they show that vacancies in exposed occupations have declined dramatically and the value of the contracts have declined dramatically. So there is already an effect taking place in the developing world. So what is missing is that most of the literature is US-focused or developed country-focused. Most of them use data from O*NET. It's a dictionary of tasks that US workers do. And developing countries differ in many ways. The mix of occupations is different. The role of complementary technologies, so the digital divide, will likely mediate any impacts of AI on the labor markets. And also, it's different what workers actually do. Most of the work uses the same occupational classification across economies without considering that workers actually may be doing— the basket of tasks may change with economic development, even for the same occupation. And finally, not all tasks are equal, right? The effect will be different whether the expert task within the bundle is automated versus the more mundane task, the more routine task that a worker does is automated.

[02:08:10.22]

So this is— I'm going to talk today also about a brief extension. That's why I mentioned this. So what we do, this paper fits into the exposure literature. As it was mentioned before, this literature has the limitation that it is about potential impacts, not actual impacts. But it has the benefit that it can— Our paper has a global coverage. So when we look, most of the literature on impacts comes from data from online platforms, which are not the most representative job in developing economies. So what we do with this exposure method is try to see where the impacts, the distributional impact of this technology, which segments of the population, which segments of the labor market are more likely to be affected. So we cover 135 countries. That is 2/3 of the global employment. It's closer to 100 because we don't have data for China. And then we add the digital divide layer where we see among those who are exposed to AI, which workers have access to a digital technology and which ones do not. Because clearly this is something that you need to have access to, to be affected by AI directly. And finally, we look at how the task content of jobs for narrowly defined occupations changes across the development spectrum using data from PIAC and STEP for 46 countries.

[02:09:41.20]

We measure this and do an adjustment of ILO occupational scores and show that this matters and show the importance when we try to analyze the potential impacts of this technology on developing economies, importance of actually look at what workers do because that varies substantially through economic development. We have 4 main findings. The first one, not surprisingly, exposure to AI increases with economic development. As I mentioned before, most jobs exposed are white-collar, high-skilled jobs, so that explains the correlation. What is more interesting is that the exposure is different. This development gradient is different for jobs exposed to automation versus jobs exposed to augmentation. The gradient is much more— is stronger for automation. For augmentation, the gradient is a bit flatter, which could potentially be interpreted as good news in the sense that poorer countries and rich countries have the same odds of being benefited by AI. However, once you add other layers, the story changes. When we control for the role of the digital divide, we see that in our sample, around 442 million workers could potentially benefit from using AI at work, but 67 of them lack internet access. And that gap is more severe in developing economies, but also we find an asymmetry.

[02:11:11.56]

We see that among jobs exposed to automation, they're already highly connected to the internet, even in low-income countries. So the pace of automation, this channel, may not be that much slower in the developing world than in the rich world. In contrast, the jobs that could benefit from AI, here we see that the digital divide is much more severe. So the disruption may come faster than the benefits. Then we discuss the idea of a potential white-collar bypass. And what do we mean by this? We mean that when we compare on the aggregate, yes, rich economies have a much higher exposure to AI than developing economies. But when we compare college graduates across the development spectrum, that exposure is quite similar. So a college graduate in Mali has the same exposure to AI than a college graduate in the US. And what explains this is mostly in the developing world is mostly clerical jobs. This may not sound like the most exciting segment, but it is. In developing countries, good jobs are scarce and these clerical jobs are overrepresented among those good jobs. And these are jobs that are already connected. and historically are the jobs that in today's rich economies allow many people to reach middle-class status and without huge educational investments.

[02:12:37.50]

In the '80s, when this peaked in the US, 1 in 5 workers were in clerical jobs and have been declining since then. In the last quarter century, in developing countries, it has been flat. And to provide some context, when we talk about automation jobs, we're not thinking about mass layoffs. It's mostly more about the generational turnover that was mentioned in the previous

sessions, but it means that the rung of that job ladder is going to be missing for developing countries. It's this typical story that moving from agriculture to manufacturing to high-skilled services where critical jobs were very important, that rung will not exist. And that's a challenge for us thinking about how, what will be that rung in the job ladder for developing countries. And then when we adjust by what people do at work, we see that this matters because even the non-exposed jobs in the US have higher non-routine cognitive jobs than the most exposed jobs in low-income economies. So this also says that the impact will likely— is smaller in the developing world. So I'm going to talk today about how do we measure exposure, the role of the digital divide, how do we do the task adjustment, and then I'm going to briefly talk about an extension that we are working on.

[02:14:01.10]

So here we use the methodology of the ILO, which I don't want to spend a lot of time here because it's not the contribution of our paper. But what they do is that using a survey, they ask people, workers, to classify tasks that they're familiar with say what is the likelihood that that task can be delegated to AI. Then this goes through an expert validation. This is a sample of tasks, but there's a total number of 33,000 tasks that would be very difficult to do manually. And then they use AI to scale this classification to the whole dictionary of 33,000 tasks. And for each occupation, they use 2 moments. to say whether a job is exposed to AI and whether it's exposed through automation or augmentation. They use the average score of tasks and the standard deviation. Intuitively, if you see the red dots over there, those are occupations that have on average high exposure to AI and very small standard deviation. That means that most of the things you do can be done through AI, and they are very similar in terms of their exposure to AI. So likely that those occupations are at risk.

[02:15:15.04]

And then we have in the middle the light blue and dark blue areas where we have some moderate exposure but huge dispersion in how your task can be done through AI. So here the argument is people can focus on the task where humans have a comparative advantage, delegate the rest to AI, and there goes the productivity benefit. Of course, there is some degree— it's arbitrary to choose this threshold. But we wanted to— the idea is to have a typology of jobs. And these are just 2 examples because I think it's important when we talk about exposure to tell people what we mean by that. 2 extremes here, we have financial analysts, a sample of tasks, and education methods specialists. Again, a sample of their tasks. For financial analysts, all of the tasks that they do have a high likelihood of being done well by AI, which is interpreting data, preparing plans, monitoring trends, analyzing financial information. Does this mean that this occupation will disappear? Likely not. They will most likely be

transformed, right? We can think about a scenario where a financial analyst will devote more time to building relations with customers, more focused on the interpersonal part, delegate the rest to AI.

[02:16:31.23]

But If all of the tasks have a very high exposure to AI, most likely that occupation will be transformed. On the other end, we have educational methods specialists where we see a more mixed set of tasks. 2 of them, which is developing the structure of the curricula or documenting your courses, evaluating courses, those things can be done well with AI. The specialist can delegate this to AI. There is a big component of interpersonal tasks where they have to evaluate teachers or be in a specific physical space. Those cannot be done with AI so easily. So here's where you argue, okay, this person can delegate and focus on the ones where they have a comparative advantage. We do a lot of data integration here. I don't want to bore you with details, but this is a result that we show. Again, not surprisingly, the exposure to AI increases with economic development. What is more interesting, I think, is when we separate by augmentation and automation. You see that when you look at automation risk in rich countries, that the fraction of jobs exposed to this is 15% or more. In low-income countries, it's near zero. But when you look on the left panel on the augmentation risk— on the automation augmentation opportunity, you see that the gradient is flatter and there is a lot of dispersion on the bottom among low-income countries.

[02:18:00.06]

And again, you could interpret this in a little bit of an optimistic way saying the chances of benefiting from AI are more equally distributed, but things start to change when we introduce the digital divide. Again, here you see microdata for several countries and some imputation exercises, we separate exposure, which is the height of the bars here, which is the same numbers I showed you before, exposure to AI that goes from around 10% in Rwanda to more than 36% in Switzerland. We divide this among those who have access to the internet and those who don't. And you see that the divide between rich and poor countries increases dramatically. If you take the case of Estonia and El Salvador, which are kind of in the middle. They have the same share of jobs that are exposed to AI, but of course Estonia is a digital champion. Everyone is connected to the internet. El Salvador, about half of people who are exposed do not have access to the internet. So this will affect, this will mediate the impact of AI in the developing world. Now it's more interesting when we focus on what's happening within countries.

[02:19:15.59]

Here we show 2 groups, high-income economies and low-income economies, the aggregate. And we show in dark color the share of jobs that are exposed to automation. In dark gray, those are exposed to automation without internet access. In green, those who are exposed to augmentation who can benefit. And in the light green, those who can benefit but do not have internet access.

[02:19:43.30]

First, if we look at the top line over there, we see what we saw before. High-income economies have twice the share of jobs exposed to AI. But when we focus on the segment of college graduates, that gap reduces significantly. It goes now in high-income economies is about 38% of exposure, in low-income countries, 33%. Also, Interesting here is that among those who face the risk of automation, which is at 9% among low-income economies, those already live in areas with very high connectivity. And these are mostly clerical jobs in developing, in low-income economies. So that automation channel may not be as slow as you would expect in the developing world through this channel. Now, when we look at the augmentation potential effects, you see that low-income countries could have that benefit even among the lower educated segments, but most of those segments do not have internet connectivity. So these benefits may take longer to realize. Now, when we look by sectors, not surprisingly, high-income economies, the sectors most exposed are transport and communications, financial and business services, and the public sector. In low-income countries, those sectors matter too. They have high exposure, but the sector most exposed is commerce, which is one of the main employers in the developing world.

[02:21:15.41]

This is, of course, mostly informal jobs, low-productivity jobs, and most of them are not connected to the internet. as you can see there. And what explains this difference between rich countries and low-income countries is the fact that, to put it in simple terms, retail in rich countries is large supermarkets. Retail in low-income countries is mostly a one-person show, someone who has to get the inputs, put them on the shelves, do the pricing, do the bookkeeping, everything. And for that reason, because they have such various set of tasks in their job bundle they could benefit from AI. And then the last part is where we do this adjustment through what people do on the job. We use skill service, PIAC, the latest round. And for those countries that were not in the latest round, we use the previous round. And also use data from STEP from the World Bank that are a bit outdated, but there is no other information as far as I know for low-income or lower-middle-income countries. We do some exercises harmonizing these datasets and classify tasks into these 5 blocks: non-routine analytical, interpersonal,

routine, and routine manual. These are some results for high-income countries for which we have very granular occupational information.

[02:22:48.29]

First, on the horizontal axis, we measure exposure using the ILO index. And on the vertical axis, we measure what people report to be doing at work. And we see that the non-routine analytical content increases with exposure to AI. So the index is validated in the sense that People actually do the tasks that are targeted by AI as you move through the exposure categories. Also interesting is that the routine manual tasks are very low among the most exposed segments. Computer use is very high among segments— occupations exposed to AI. And finally, the non-routine interpersonal tasks, they peak in the occupations that have more that theoretically are more likely to benefit from AI. Confirming this idea, these are the occupations where you have substantial variation what people do. They have many tasks that can be delegated to AI. Some others where you need a lot of this interpersonal, physical component, interpersonal component, not routine, because those who are more likely to be automated have a lower, as you can see at the end of the line on that chart, a lower interpersonal component. And these are the averages by income group, country income group. And as you can see in the middle chart, non-routine analytical tasks, again, the target of AI, even in the segment that according to the ILO index, the segment that is not exposed to AI, those workers have a higher non-routine analytical content that workers, the most exposed workers in low-income economies, the workers, lower middle income economies, sorry.

[02:24:37.10]

And upper middle income countries are in the middle, but still their average content of non-routine analytical tasks is considerably lower than in high-income countries. When we do a simple statistical adjustment of ILO scores to take into account the different task content of jobs across economies, we see that on the left here we have the raw ILO index, which shows this correlation that we documented so far. Richer countries have more exposure. But when we adjust this by what people actually do on the job, the gradient, the slope is higher. We see that the dots for Kenya or Ghana go down because people do in practice much fewer of the tasks targeted by AI. In contrast, Japan goes to the top. And again, I think here the main message is when we need to study the spread of AI and the potential impacts or any exercise related to new technologies in the developing world, it's key to have access to this type of data. and what people actually do at work. And unfortunately, for the developing world, the data is already old. It's from several years ago. And just to finish, I want to give you a flavor of some work that— another background paper for WDR with a complementary angle to what we did here, which is extending the work of David Autor and Thompson on the role of expertise.

[02:26:13.11]

to the developing world because what does expertise mean is domain-specific know-how that a task requires and that has a market value. And again, it's not the same if AI replaces that, an expert task, than an inexpert task. And here you can see when we— It wasn't an image. OK, so in this chart, I wanted to show you the relationship between expert jobs and jobs exposed to automation. And here is for high-income countries, but when you look at this through the development gradient, the bubbles in the upper right quadrant, those are the ones that become fuller because you have more employment in those jobs that are more exposed and are more expert. But then we use this information to do a very simple counterfactual. What will happen if we— all the tasks— what will happen to expertise of a job if we remove the tasks exposed to AI? And these are the 3 main groups where expertise is targeted by AI. We see that among professionals and technicians, expertise increases, meaning that these occupations can delegate the more common tasks to AI, and the barriers to enter the occupation increase. For clerical support, which is a broad occupational group that is already suffering automation, elimination of jobs, for the jobs that survive, the jobs are becoming less expert.

[02:27:44.53]

So the barriers of entry decline, entry to these occupations is easier, likely earnings will decline. And this will accelerate something that was already taking place for decades, right? This is what I was mentioning in the beginning on clerical that in the US they peaked in the '80s. 1 in 5 workers had this type of jobs and they have been declining since then. And also in Switzerland and Brazil, in Switzerland and France, I'm sorry. And as you can see, the developing world, they have been flat over the past 25 years and likely they will not increase given all these trends on AI and trends that have been happening, technological trends even before this. So the challenge here is thinking thinking what is the rung of the stepladder that these countries need to build to have a skilled labor force. So to conclude, exposure rises with development. It is important to consider the digital divide. This will have an asymmetric impact where occupations already exposed to automation in the developing world are largely connected to the internet, so the automation channel may not be that slow. For augmentation, yes, those jobs happen to be less connected to the internet, so this benefit can take slower to realize.

[02:29:03.08]

And finally, it is important to consider the fact that even though on the aggregate exposure is significantly low in the developing world, it is highly concentrated among the few good jobs that

low-income countries have. And this creates a challenge when we think about how to build the job ladder in the era of AI in the developing world. Thank you very much.

[02:29:36.47] - Åsa Johansson

Thank you very much, Hernán. That was a very interesting paper with lots of insights on how important it is to actually take into account what the workers do in the job.

[02:29:45.37]

as this difference between advanced and emerging economies.

[02:29:48.08]

So thank you for that.

[02:29:49.01]

Our next presenter is Gabriel Marin from University of Pennsylvania, and he will present his paper on the Informal Buffer: How Labor Market Duality Mediates Automation and AI Exposure in Mexico. The floor is yours.

[02:30:12.59] - Gabriel Marin

Okay. Hello. Good morning, everyone. My name is Gabriel Marin. I'm going to be presenting some joint work with Guille Lagarda at the IDB and Boston University's GDPC. The usual disclaimers apply. The paper is called The Informal Buffer: How Labor Market Duality Mediates Automation and AI Exposure in Mexico. And basically, I would like to run you guys through a story of how automation affects in Mexico, how informal markets work in this case, and how as policymakers we need to be rethinking the way we tackle informality because there's some nuances over it. So I will begin with the typical story of the Autorndorm framework, and I know most of you have seen this graph a billion times. But on the left, you see the employment changes from 1980 to 2005 in the United States. On the right side, you see the real wage changes for those same occupations. They're ranked by wage— by the wage distribution. So on the further right, you'll see people that earn a lot of money. On the left, people that earn less money. And you see a typical hollowing out in the middle. And they hypothesized that this was because occupations who carried a high routine component were being displaced by automation technologies.

[02:31:42.59]

Whereas the other ones in the polar opposites, on the right side, high-skilled workers had a more abstract component to their tasks, and on the left side, more manual for those who are low skill. And that drove into wage polarization. Now, I'm going to show you this same exact graph, but for Mexico. And for Mexico, we don't see something as clear as the Autor and Dorn framework. On the left side, you will see the blue and the red. Blue is the Mexican wages. On the red is the 1980s same distributional. from Autor and Dorn. And we do see the kind of U-shape in the employment. So there's a hollowing out in the middle. But if you go and check the wages, you don't see that typical polarization story. So there is a generalized compression in the formal markets. So for this paper, we ask the following question. Does labor market duality shift the incidence of routine automation from wages to jobs? And we find 4 results. The first result is that formal clerical jobs fall 0.44%age points per 10 points of exposure, that formal wages do not change in the reduced form. We do a 2-stage least squares.

[02:33:11.35]

I will get into details further. And on our main result, we find that estimating a buffer where the informal sector absorbs through self-employment around 2/3 of the jobs that are displaced from this automation shock. And now we go to the present, which was the future, but now the present, like AI workers using the Olundu classification, which is the typical thing that we we use. And we find so far that there are some effects in the clerical and retail occupations. There's some declines, but no wages. So informality acts as a buffer protecting formal wages, but it has the cost of social protection. To discipline ideas, I will present this baby graph of a routine labor market. And as you can see, you have an elastic labor supply. demand for routine jobs, and then an automation shock hits the economy, which makes automation technologies price cheaper for formal firms to adopt. So this compresses the demand for routine jobs and it acts through wages and less through labor when you have an inelastic labor supply. But then if you have an elastic option, so an elastic labor supply and big informal market like Mexico, 50% of the labor force, then you see that the effect is more through labor market adjustment, like the people, instead of the actual wages.

[02:34:52.01]

And we then rationalize this simply through an occupational choice. So your formal wages minus your deductions plus a taste, which is this epsilon, must be higher than your self-employment. And with that, we obtain that the elasticity of supply is an increasing It's increasing in the share of clerical informal share. And solving for the model, we find that the buffer is the ratio of the elasticities. Simple result.

[02:35:21.01]

So the model predicts that the formal jobs fall, that wages hold, and some— and the workers that are displaced either go informal, either self-employed or unpaid. And unpaid, I'm talking about people that work at a family firm. Okay.

[02:35:42.27]

So I'm going to spend maybe like 2 minutes here. I don't want to bore you with all the details, but I do have to explain it. Here, our unit of analysis is the local labor markets. And this is work that was done by Professor Hansen, who is not here, but I wanted to thank So I'm on the record thanking him. He did some really good stuff with the Bank of Mexico to create commuting zones. And this is what we're leveraging now. So you can see on the left, this is the variation that we're identifying. This is the exposure of all the labor markets in Mexico. The darker are those who are higher in routine shares, and the clearer the opposite. So we define exposure as Autor and Dorn did with routine task intensiveness. And those occupations that exceed the 66th percentile, we classify as exposed. And we instrument through predetermined industry shares of 2005. So basically, identification by shares. And we use this to instrument the routine share of the start of period employment. This is the typical specification. We do this for all occupational groups, which are going to be 6. And I'm going to be presenting them throughout.

[02:37:07.22]

There is something that I do want to say is that you see the little boxes of first stage, pretrans, robustness. I don't have those here with me. I'm happy to share them with you if you're interested. But for the spirit of just sharing a small PowerPoint presentation, I removed them. But if there's something that's bugging you, I'm happy to send them to your email after this. So—

[02:37:34.38]

oh, sorry. And I am going to be presenting graphically beta G. OK? This is going to be the occupational groups change. And here, I want us to focus on the clerical and retail occupations. This is our preferred specification. And we find that there is a decline from 2005 to 2019 using the ANOE, which is Mexico's local labor market survey. And we find that 10 points more formal employment in routine jobs, the clerical share grows 0.44%age points less than otherwise. Similar market. I will abstain from talking about the managers and professional ones because we do find that there is some stuff in the Rodenberg weights. This result is mainly driven by the apparel industry, just one industry. If you remove it, it's no longer statistically significant. The Cleric on Retail result is robust to a humongous amount of tests that you can imagine. On the reduced form, on the other hand, we don't find any kind of detectable response. And this is

basically just the instrument with the log changes on— from 2005 to 2019. And now our main specification is we cataloged those workers who are in the payroll those who are self-employed or classified as unpaid, out of the labor force and unemployed.

[02:39:11.44]

And we run a 2-stage least squares again for the routine share. And we find that those payroll loss positions— and this is a change in the coefficient. I just put the 64% for readability. 64% of those who are displaced reappear as self-employed or unpaid. And we find that the unemployed, 'cause one might think, oh well, these people are displaced, maybe they're going to unemployment directly. We find that this zero is actually statistically significant. So there doesn't seem to be this story that people are just becoming unemployed, but rather they're changing the way they carry out economic activity.

[02:39:59.16]

So why does this matter? Well, this matters because, and I am Latin American, so obviously I am prone for drama and soap opera, and I will give you guys the story of this person now, 'cause we're working with a policy audience, so we need to see who is actually being affected by this, right? So meet Maria. This is the most common name in Mexico for women. She's a secretary, formal, has access to Social Security, so she has IMSS. She's 38 years old, has high school, earns 3,600 pesos per month, and her task composition is routine intensive. She's top third in the formal jobs. Of those clerical and retail, almost half are secretaries, and of secretaries, 64% of them are women. Okay?

[02:40:53.47]

So before, she's formal, she's a secretary, she's earning her IMSS. But now she faces the shock and reappears as an informal self-employed worker earning a wage cut or unpaid. So she goes to work with her uncle, with her aunt, with her brother, whatever. But you're not able to see how much they're earning, but you know that they're not having access to Social Security. So she's informal either way. Now, what are the policy implications for this? Well, I'm here, I'm going to spend a little bit of time. We have a lot of work that is talking about informal markets as being very bad for the economy. a necessary evil. Now, they're perverse, they're persistent in Latin America. Okay, that is true. But here we are finding this kind of not so bad result. Like, the existence of an informal market is actually allowing her to not just be unemployed. So when one wants to think about policies that tackle formality, we need to be more nuanced in the way that we carry them out. I will elaborate. If one government would want to promote policies that

increase formality through payroll jobs, or maybe the story is, oh, well, these people— it's very common in Mexico.

[02:42:37.14]

Like they just take you out of Social Security, but you're still working for the same person in a payroll. We do see a generalized decline in those payroll jobs. If you're trying to tackle those through formalization policies, you're actually not hitting the tax base that you intended, but rather those who are self-employed are the ones that are actually being absorbed in this case. So one needs to think of a more refined way of doing so. And I'm going to be quoting work from Santiago Levy, who has done a lot of this. And maybe we can just remove the payroll tax, generalize it with a universal health coverage, and couple it with a universal VAT. That would alleviate the distortion that creates the occupational choice of the worker trying to either go to formal or informal and would allow both of them to have this access to Social Security. And finally, one has to be balancing out informality acts as a buffer. absorbing these people that are being displaced. But you also know that the informal firms are the least productive now, and we see them through the whole conference, no? So just some food for thought in the way that we can think about this.

[02:44:12.06]

So now I'm going to go to the future, now present day artificial intelligence, and we use the Ilundu classification. We find no results whatsoever for informal markets yet. These are not identified, and hopefully we'll be able to find some relevant results once we update the NOI. But don't quote me, don't wait for them yet. We do find some decreases in clerical and retail occupations and managerial and professional for formal markets. If you ask me, I do think find it plausible that there is a difference between formality and informality in AI adoption, one might think this technology hypothesis could be something that is driving these results. Nothing identified yet, so take me with some limitations, of course. And this one— this specification was done through an identification by shifts rather than by shares. And this is mostly because we have a small change in years, 2022, 2025. So we didn't seem to use— like using 2022 shares was not something I was feeling comfortable with because they could be— them being possibly exogenous doesn't seem to be the case. And now the reduced form for this one, we don't find anything in particular, no result whatsoever.

[02:45:57.22]

But we do want to improve upon the paper and start designing something with respect to hiring margins because that seems to be where the catch is. And I'd love to pick your brains after this

with respect to that. So finally, what happens to Maria, this person? Well, her job's still exposed. AI still classifies her as potentially exposed. So formal jobs like her shrink, her pay does not move, but where she goes is still open, at least in the age of AI. With automation that we know that the informal market has absorbed, people that are similar to her. So us as policymakers need to be thinking about this nuance. So formalization policies can protect the vulnerable, but then these are the least productive firms. So how can we balance this out? Okay.

[02:46:56.47]

And finally, the informal buffer is protecting wages. at the cost of social protection. Formal jobs shrink. This is consistent with some work by Brambilla et al. It's Argentina, Mexico, and Brazil. They use robots though. And the self-employment result is consistent with work from Dix-Carneiro, but they use trade liberalizations in Brazil. So maybe you can extrapolate in a way. But we do find differences in the wages. They do find wage polarization in their paper. We do not. So that is one of the differences from our paper and our work. So it's the same shock, but we have a different margin. Okay.

[02:47:42.31]

And that is it. Thank you.

[02:47:55.31] - Åsa Johansson

Thank you very much, Gabriel. That was really an interesting paper to hear about, more about emerging markets and the specificities and how informal jobs can be a buffer. The next speaker, the third speaker, is Filippo Bontadini from LUISS University. So he's going to talk about his paper called Exposure to Digital Automation and Jobs in Global Value Chains. So Filippo, the floor is yours.

[02:48:29.02] - Filippo Bontadini

Okay, thank you. Let me see. Yes, great. Yes, so this is joint work. You can hear me? Yes.

[02:48:34.53]

Yeah. Okay, great. So this is joint work with Tommaso Cerli, Santiago Picasso, Maria Savona. It's about— it takes more of a global outlook. It's more aggregate, and we try to study exactly Essentially what it says on the tin. So how different exposures to digital technology— so we take a broad understanding of what digital technologies are. It's not just about AI. It's actually, it's just a little bit about AI, but it's more broader technologies. And we look in particular at jobs and in

particular global value chains. So we try to bring also a trade perspective and a production linkage perspective to what's been said so far. So our key starting point is that we can think of technological change in a broad way as affecting employment through some productivity gains. There'll be obviously some displacement and some new tasks, but also crucially, and I think this is especially important for emerging economies, it will also affect the way in which countries and industries plug into international production networks, so global value chains. To the extent that obviously you can have a country sector that might be exposed, and I'll discuss a little bit later what we mean by exposure here, to a certain technology without necessarily adopting it, but that technological change that might happen at the frontier will reshape the way in which production chains are configured, and that will have an impact on the kind of employ— the amount of employment that is activated within country industries.

[02:50:09.22]

So our key contribution here, we rely on a highly granular measure of exposure to digital technologies. The measure has 40 digital technologies grouped into 9 families. The results I have are for the 9 families. And then we try and decompose— well, we try, we decompose employment changes into 3 main components. One is labor productivity. One is what I call here GVC linkages, although you can think of it more in terms of production linkages, and then final demand, of course. And this is based on the work by Pal and Cuartos. It's a pretty standard structural decomposition analysis. I'll dwell a little bit on that in the next couple of slides. What I think is also important is that we distinguish between GVC jobs, so jobs that are activated by foreign final demand. So we are able to say which portion of total employment within a country industry comes from final demand that is total domestic and foreign, and the portion that is just activated by foreign demand. This is especially important because naturally participation to international production linkages often follows different technological paths. And so it's important to distinguish, and we do find different patterns. So our overall research question is, how does the emergence of digital technologies relate— emphasis on relate, we don't have any causal evidence, it's more about understanding what different patterns are afoot— to employment changes through labor requirements in particular and GVC linkages or production linkages more broadly?

[02:51:52.00]

So to try and help think about this, the way we've thought about this is that you can So, with our decomposition framework, so at your far right, you have employment growth, which you can decompose into these 3 main components. One would be obviously final demand growth, and then you have labor requirements, which is the inverse of labor productivity growth, and then production linkages, which says how much of total final demand you're actually supplying to.

What is your— you can think of it in a way as a sort of market share concept. And then you have technological development at the frontier, which obviously is directly related to final demand growth. And so it will both be a demand pull and a technological push story. And what we do is that we interact that technological development at the frontier with a measure of digital readiness. We try different measures. I'll discuss in a second which one we end up picking, but it's sort of, they move in a similar way. And so through that channel, that will affect both on the one hand, how productive, what is the dynamic of productivity in a country industry, but at the same time, how much of the final demand that country industry will be able to supply to satisfy.

[02:53:00.41]

And so these will have different effects. We don't have a straight— and this is something I'll be happy to get suggestions about— we don't have an obvious relationship between technological change within a country industry and final demand at the aggregate level. the extent to which that what happens at the, say, in the Taiwanese electronics industry, the extent to which that could affect final demand for automotive industries. It's not obvious that we are, for the moment, we use it as a control, the final demand channel, but it's an open issue and we're still thinking about that. What is it we do? We rely on the wonderful data that the OECD has compiled about international country input-output tables. They've finally, I would say, finally for us released previous year prices input-output tables, which is obviously great because we got a lot of comments about the fact that obviously nominal productivity isn't labor productivity. We now, as of this summer, have rerun the entire thing and we got— so we can decompose employment into E prime, that is your employment per value added at constant prices. V is value added share in gross output. B that you have up there is the Leontief inverse.

[02:54:19.32]

The Leontief inverse, it's a core component of input-output analysis. It essentially allocates final demand to the production factors. So we'll be able to say of the final demand for, which is F, which is say, a car bought in Germany, where is value added and through labor productivity employment generated, not just in Germany, but say in the Vietnamese sector that provides the electronic components that go into a car that is sold in Germany? And this would apply obviously for aircraft in Toulouse or fridges in Poland and whatever you wanna decompose it that way. And so we have that you can decompose the log change in employment into a contribution from labor requirements. So how much employment is needed to produce a unit of value added, which is at constant prices, it's the inverse of labor productivity. You will have a change in VB, which is the value added adjusted which sums up to 1 along the columns. So it's saying for each unit of final demand, what is the share that each country industry provides? And then you have final demand, which is a sort of scale effect of the more final demand grows for

automotive, for say cars in Germany, that will trigger more production throughout value chains trickling down all the way to the Vietnamese, say, electronic components.

[02:55:57.11]

And so we focus here on the first 2 terms that I was just mentioning, labor requirements and GVC linkages. If you have a negative contribution of labor requirements, that means that labor productivity has gone up, so you need fewer workers for the same unit of output. And that would be, for example, a process innovation kind of story where you get better at producing the same unit of output. You can also have a positive, or what we find in reality is a less negative contribution, that would be linked to the idea that innovation generates new tasks, which might be newer technologies that need more workers and less standardized. So that is also familiar with some of the presentations we've seen this morning. GVC linkages, if you have an increase in the GVC linkage contribution, that usually means that for an increase in final demand, you are getting a larger share of that final demand is coming from that country industry. The country industry is becoming a bigger supplier, a more central supplier into a value chain. You can also have a negative contribution, of course, that doesn't happen too often, but it does happen in some country industries.

[02:57:05.13]

The way we interpret that is an outcome of an outsourcing or a production reallocation. So as technology develops, you might not need as much of the same input as before, and that might therefore lead to a restructuring. structuring of production activities across the chain. I've already mentioned we use constant prices. We also use obviously the TIM, so trade and employment data also kindly supplied by the OECD, which we then complement with ILO stats to cover certain countries and industries that are missing from the trade and employment data. And obviously we do this for the total employment, so that's what you see, and then we also have a version where we set to zero all of the cells in the Leontief inverse that correspond to own domestic linkages. So, that nets out the jobs basically that are activated by final demand coming from within the country. So, those are GVC— the way we define GVC jobs are all of the jobs that are activated all of the way down at the end of the value chain by foreign demand. So, So to take again the example of the Vietnamese electronics industry, if those workers are contributing production that goes to Vietnamese demand, those would not be GVC jobs.

[02:58:24.53]

If they are contributing to inputs that are satisfying final demand, say, from the German automotive sector, those would be GVC jobs. So a little bit on the descriptives that we have. So

these are the 3 components that I just discussed by high-income and middle-income countries, you find a pattern that is well, I think, known in the literature. You have final demand that generally supports employment growth. This is a scale effect. And on the other hand, the blue bars are labor requirements contribution. So that's overall good news. Employment, labor productivity has gone up over— this is for the period 2012-2019, by the way. has gone up, which usually means that, you know, you need fewer workers for the same unit of output. Linkages contributions vary quite interestingly across countries. I also wanna emphasize that the labor requirements is particularly strong in middle income. You probably can't see this, but it's, you have Côte d'Ivoire, India, and China are the biggest, are the ones with the biggest negative labor requirements contribution. If you're wondering, that big positive blue bar for labor contribution for AGO, that's Angola. It's something that we're still investigating why we find in the data that something has happened in that.

[02:59:51.45]

And if we portray the same patterns for GVC jobs, the main difference is that GVC linkages become much more important. So it seems that for international linkages, the role of the position of country industries is particularly important to— as a contributor to employment changes. We then do this also by sectors. And so what we find is that overall, there's a clear shift. So the diamonds that you observe are the employment log changes. And so there's a clear shift towards services. Services have overwhelmingly grown, seen a growth in employment. While in manufacturing, this is only concentrated in a few mostly high-tech sectors, motor vehicles, pharmaceuticals, fabricated metals, and electronics and optics, as well as other manufacturing repair. And this seems to be largely driven by a negative contribution of labor requirements, which again speaks to the difference in labor productivity dynamics across the 2 macro sectors. Again, if we have the data as well for the GVC jobs, similar patterns, the GVC linkages contribution is much bigger. Okay, so we have these patterns as a decomposition of employment growth, and then what we try to say is we try to relate this to this measure of digital technologies.

[03:01:21.34]

What we do is that we rely on this measure called tech— of tech exposure, which has been developed by Pritikova and co-authors, where they use a technique similar to what's already been discussed here. But essentially, it's a text analysis of patents from PATSET and the descriptions of 4-digit ISIC industries, which they aggregate up to 2 digits as well. And that gives us a measure at the industry level of the exposure of each industry to the different technologies. Here you have the heat map, so you can see that digital technologies overall affect largely services, although there's variation in manufacturing industries. The first column that you see

more geared towards manufacturing, that's 3D printing, which obviously is especially important for manufacturing. And then we say, okay, well, we have that over the period 2012, 2019, there's been a lot of development in these technologies at the frontier, which vary in terms of their exposure across industries. And then we say, naturally, it won't— these— the way in which these will affect our variable, our observations are not exactly the same across countries. So we have an index of digital readiness, which only varies across countries. And we use here fixed broadband subscription per 100 people.

[03:02:39.55]

We tried a variety of different measures. The results are— the patterns are pretty consistent so that that doesn't particularly change. So what we have here essentially is this idea of an exposure to frontier development, that our measure of technological exposure isn't country-specific in and of itself. It says something about how much each industry is exposed to technological change happening across the globe, but though especially in the US, Europe, and to some extent China since this This is based on patenting data. And then we interact this with an idea, with a measure for technological, for digital readiness. So what we then try to estimate is the log change of employment and its 2 contributions coming from labor requirements and GVC linkages. We have final demand as an additional control, obviously to control for the fact that labor demand for final demand will have changed differently across countries and industries. And we include regressors for our technological exposure and digital readiness interaction across the 9 families jointly. So what we have are net relationships. We focus on non-primary market industries, and we run a weighted regression using employment as a weight. These are semi-elasticities, and the coefficients of it add up to the total observed coefficient for employment changes.

[03:04:14.40]

So what we find overall is that if you look— if you take employment, total employment as a whole, we don't find— we find fairly muted results. You observe obviously that labor requirements and GVC linkages often work in opposite directions so that we seem to find a parallel of this pattern of on the one hand you get more productive, but on the other hand you also get more competitive. And so these 2 trends seem to net each other out. What is interesting is that, for example, for computer vision, which includes as well machine learning, here we find labor requirements actually have been— the contribution has been overall negative, but less negative in those country industries exposed to labor requirements. And on the other hand, we find GVC linkages that seems to have shrunk. So that suggests that, yes, there's been some technological change which has made new activity available in those country industries. But at the same time, this is counteracted by the fact that production has also been

restructured away from these very same industries. For digital services, we find the opposite pattern. This is more of a typical— there's more productivity, but also more bigger involvement in production linkages.

[03:05:27.11]

When we then look only at manufacturing, we find a very different pattern compared to what we have. So in particular here, what jumps out is e-commerce. E-commerce includes digital advertising as well as more kind of standard e-commerce, so coupons or digital auctions. And I think it's interesting here in particular if you think that E-commerce obviously does not— is not necessarily adopted within manufacturing industries, but it does affect the kind of activities that are carried out within those industries because e-commerce will largely be adopted by retailers and the way in which products are shipped rather than necessarily manufactured. And then if we try and focus on GVC jobs, so the portion of employment that is activated by foreign final demand, What really jumps out if you look at the entire sample, labor requirements pattern are very similar. However, GVC linkages now become quantitatively much, much bigger. And that is especially true for digital services, but also computer vision and health tech. And it's true across the board, I would say. So here again, you find that for those jobs that are activated by foreign final demand, the role, the way in which these country industries are integrated into production networks is especially important.

[03:06:47.10]

If we then look at manufacturing GVC jobs and we compare them with manufacturing overall jobs, what jumps out again is the GVC linkage component, which flips. So you find that digital services that used to improve or increase the amount of value added that you you supply to final demand. If you focus on manufacturing jobs in GVC, that becomes a negative component. So there seems to be an outsourcing pattern for those. For computer vision, it's the opposite. In jobs in manufacturing industries activated by foreign final demand, GVC linkages have had a positive contribution for those country industries that have used more intensively, that are more exposed to computer vision among which machine learning. So you find again these differences really in patterns between manufacturing, GVC jobs, services, and also crucially across technologies. So let me just conclude with a couple of points that we've established. That's a significant— and I think it's important for the policy discussion in particular— there's a large and significant heterogeneity across different technologies. So whenever we discuss the emergence of digital technologies, it's important to bear in mind that not every technology will act, will behave in the same way. You have within aggregate employment changes opposing contributions coming from productivity patterns, but also insertion into production networks, and especially in GVCs, so international production networks.

[03:08:22.51]

Manufacturing and services exhibit very different patterns. And so it's also true whenever we think about emerging technologies, it's not just an automation story, it is also the creation of new tasks, of new activities, and the way in which different value chains, different sectors and countries can insert into production networks. And this is exactly the last point. So GVC employment seems to also follow different patterns compared to overall employment. You'll notice I haven't spoken obviously here about salaries, but there's obviously a significant amount of evidence that shows that those GVC jobs tend also to have higher salaries, higher productivity. And so you can— I think there's evidence to suggest that it's important to consider these different production networks based on whether those are domestic or global. And that is, I think, especially important for emerging economies. Naturally, future research that we are undertaking at the moment tries to distinguish between this idea of exposure versus adoption. We do our best here using these measures that we have about digital readiness. I'm not particularly enthusiastic about what we do, but it's what we have for preserving the coverage of countries and industries. And I think it becomes very important to also distinguish obviously the underlying mechanism that at the moment we only have in terms of aggregate patterns.

[03:09:40.41]

But I think this is some work that I hope contributes some nuance and this idea of having a more careful approach when we discuss digital technologies as a whole. And especially if you want to then apply to artificial intelligence, which will have all sorts of different applications. I think it's important to bear in mind how different technologies interact with different sectors, different countries, and different production systems. I think I'll leave it at that. And any questions are welcome. Thanks.

[03:10:22.58] - Åsa Johansson

Thank you very much, Filippo. Another fascinating paper that shed light on some other related issues on technology and GVCs. We have now a discussion with David Kaplan from the Inter-American Development Bank. So David, please take the floor. What do I do? Ah, there we go.

[03:10:51.50] - David Kaplan

Well, first of all, thank you so much for the invitation.

[03:10:55.16]

Let me start by saying, so I sent these slides on Monday, and then yesterday I was told, please mention that the IDB just published its its flagship report this year on AI, From Digitalization to Artificial Intelligence: Turning Promise into Prosperity. Please download it and most importantly, if anyone asks, please say that I mentioned it.

[03:11:20.43]

Anyway, I'm not going to exactly comment in detail about these 3 papers.

[03:11:24.55]

I'm going to try to put them together in a certain way and try to focus on some of the policy issues. I want to come up with this as a little bit of an AI skeptic. Let me just say that's my position, not necessarily IDV's position. In a lot of ways, I'm not 100% convinced that we're facing something so revolutionarily new that we won't recognize the labor market in a few years. So first, talking about the paper on Mexico and talking about the informal buffer, I mean, it goes to some, first of all, some very basic issues. So I really should have put, is automation displacing workers in emerging markets rather than AI? That's most of the paper. But I think this paper does a good job in not only sort of trying to address that question of, is automation displacing workers? And I think the answer is a convincing yes. But also what happens? And I think this dimension of unemployment versus informality is a crucial one. But I would issue I would offer 2 comments. First of all, they've been around a while. There's always some shock, and that shock is always doing something to someone.

[03:12:59.07]

I think to the extent that we could not only document, are there impacts? Of course, how could new technologies not impact the labor market in some way? But is the magnitude of this shock bigger than earlier ones or not. I mean, I think that's one of the issues. As far as policy implications, I would at least put on the table the following issue. Perhaps it's the same distortions that generate the informality that you argue, and I think convincingly, is somewhat of a buffer, but that also could be part of what's causing the displacement of the workers. And in particular, and this if it were true would be particularly worrying, you could be observing something that Asimovglu and Restrepo would call automation, socio-automation or socio-technologies, in the sense of It could be, again, the very same reasons why there is such high informality, that hiring formal salaried labor is so costly, so difficult, that it could be that when an automation technology comes in, even if it's not more productive, if it gets me out of the whole mess of hiring formal salaried workers, a firm might choose it even though it's not a productivity-enhancing technology.

[03:14:34.29]

And that, of course, would be a real tragedy. And maybe I would just point out on the policy side, maybe you could convince me to be tolerant of informality, although sometimes informality equates to illegality and it's hard to be tolerant of that. But the low productivity that goes along with informality, for me, is very hard to tolerate. And maybe this is like a snarky comment, but if all employment were in subsistence agriculture, probably automation technologies just wouldn't have any impact. But I don't know if that's the world we want. So now moving to the disruption without dividend paper. Again, I think this is also very good, and I think the main contribution is not telling us how AI is going to affect the labor market. I just pulled a quote, the aggregate implications for labor demand and a bunch of other things continues to be highly uncertain. That was true before they wrote the paper, after they wrote the paper, after seeing all the papers in the conference today. But I think the real issue is it's sort of putting new policy questions on the table about things we should worry about, which, not to criticize, but I think it's a huge improvement over X percent of jobs are exposed.

[03:16:07.44]

I would always think, first of all, I don't know how you get that number, but even so, I don't know what to do with that number. Whereas, I mean, I'll try to show a little bit later, I think some of these more nuanced approach, not just what percentage of jobs are exposed, but what are the characteristics. I think they do lead us to some policy implications. But before I get to the policy implications from this paper, I want to talk about the more general term. How can we have policy implications if we don't know what's going to happen? Interestingly, this is something that came up yesterday. It's probably dangerous to aggressively pursue policy changes that are only going to be useful if our models, if our guesses about the impacts of new technologies, AI, are right. But I was interested. I saw— I forget who said it yesterday, but some of the policy recommendations are stuff that we've been saying for a while now. And actually, even if it turned out that AI wasn't quite a revolutionary technology, we'd probably do well advancing on these agendas as well. Anyway, going back to the so-so technology example I mentioned, a lot of us for a long time have been saying we should probably change labor regulations, Social Security regulations, tax regulation, tax codes to avoid punishing the hiring of formal salaried labor.

[03:17:41.44]

By the way, a part of me thinks no one's listened to us for like 30 years, so if AI is the excuse, maybe I'm all for it. Maybe I should say, yes, because of AI we should do that. But the truth is, a

lot of us have thought we should do it for a long time now, continue to think we should be doing it, and it is true that in the case that AI really has the effects. We think it might, you know, it might be even more important in the near future than it was when it was already important in the not-too-distant past. And I just wanted to mention, and this I think shows how a paper like this can really, you know, put new issues on the agenda and things I think we would have confidence recommending for policy recommendations, you know, it's unlocked at least one new fear for me, that despite being a skeptic, I don't know what's going to happen with AI. What if in emerging markets we bear the full brunt of the labor substitution, the labor replacement, but then it turns out because where digital technologies have penetrated or not, what if we don't get in the emerging economies the benefit of the labor augmentation?

[03:18:55.25]

That's definitely something we don't want. Leading to policy recommendations I think we could feel confident about, improving the digital infrastructure of emerging economies is something—it's got to be a good idea anyway. To the extent that in an AI world it might be even more important, this feels like, despite all of the uncertainty up the kind of policy recommendation we could get behind. Because again, if I said 10 years ago we should help digital infrastructure in emerging markets, I don't think anyone would have terribly criticized the policy recommendation. Now there's an additional reason why such things might be even more important. Then there's this bit that AI might attack these office-based occupations, which are disproportionately a share of good jobs in emerging markets. I was struck by that, but to be honest, I didn't know what to do with it. I'm now worried about it, but I don't know if other people have thoughts. I couldn't think of an obvious way to mitigate that problem. Now to the last paper on global value chains. I think this is important. It's obviously, of the 3 papers, the one that sort of most seriously takes into account equilibrium effects and puts, again, new things on our radar.

[03:20:28.07]

Well, at least for me. I can automate, I can replace workers with AI, or I can offshore, or I can reshore, but instead of reshoring workers, I can reshore or capital. There's all sorts of ways this can be worked through. But I also want to mention, as sort of a— when I say an AI skeptic, I don't mean AI won't affect anything, but the idea that we probably won't see radical, radical, radical changes. We saw this in the presentation. A lot of times these equilibrium effects cancel each other out. Maybe There's a labor-replacing technology and that increases productivity and maybe that on its own would lead to lower employment, but then when you work through the equilibrium effects, some other things compensate. Again, it would be a terrible injustice to the paper to say, and it says in the end, nothing happens. course, plenty of things happen. There's differences across countries, differences across industries. I mean, legitimate concerns maybe

for inequality. I don't want to say that the paper is saying nothing happens, but on the other hand, I do— in fact, I would argue it's very important that these sort of papers get worked out because they frame the way we think about other sorts of studies that say, you know, almost all jobs are gonna be destroyed, and it's something like what jobs are exposed.

[03:22:04.11]

And I think it's important when you work through the equilibrium implications, sometimes things aren't so terribly wrong, or the change, the net changes may not be so terrible or so radical. So I just wanted to clear, so I just wanna leave this kind of idea about policy as you know, maybe we don't need radically new policies, especially because, by the way, I mean, I've been listening for a while. I haven't heard any like radical, radical like policy ideas. I never, you know, I hear things like training and, you know, the universal Social Security. You know, I hear a lot of policy recommendations that, you know, I feel like are not so new. But actually, I think that's okay. I mean, maybe that's the best way to think about it. We have new technologies coming in that we think might cause labor displacement. Maybe just going back to some of the things we've been saying for a while now, correcting labor distortions, labor market distortions that could lead to the adoption of social technologies. Again, I think we should have been doing that for a while, but there might be more reason to do it now.

[03:23:22.16]

Ensure countrywide access and equal access to digital technologies. That seems like a perfectly reasonable policy recommendation. I'm not an expert in digital technologies, but I'm sure that's a quite reasonable thing to do with or without AI. Strengthening skills for policies, ensuring access to vulnerable populations. I mean, I think, again, we've been saying this for a long time, and if anything, it's more important than ever, these sorts of new policies. And just to make one mention to the IDB flagship that was published on Monday that you should download, again, that's from digitalization. I forgot the title, but I know it's from Promise to Prosperity. AI digitalization.

[03:24:17.33]

I forgot the title.

[03:24:18.38]

Anyway, the flagship on AI. Despite my skepticism, it is likely that AI will cause a reallocation of labor. That's almost certainly the case. And in fact, I would say to take advantage of AI, we need a reallocation of labor. And that's, I think, one of the main messages from our recent flagship

report. So strengthening the mechanisms to protect people as they will undoubtedly need to transition from jobs because of the adoption of AI, because of aging, because of climate change, because of a whole whole bunch of shocks that will continue to hit the labor market. So anyway, thank you so much again, and I apologize to the authors if I misunderstood anything, but thank you so much.

[03:25:25.13] - Åsa Johansson

We'll have a Q&A now from the audience. I invite the presenters and the lead discussants on the table. So, thank you very much for this very insightful presentation.

[03:25:51.58]

I see that we have already people lining up from the microphone for questions.

[03:25:55.52]

So I'll hand over the mic to the person at the— but please introduce yourself as well when you ask your question. Thank you.

[03:26:05.43] - Audience

Norman MacRae Foundation. I'd like to play devil's advocate because I thought the India AI Summit actually started with the hypothesis that maybe AI in developing countries could create more livelihoods. I come from a bias. I'm a statistician, so I always like to see what data we're doing AI on, because I see AI really as machines with mathematical capabilities beyond our individual minds. And if we apply, you know, there's a lot of health data, there's a lot of agricultural data, there's a lot of energy data at the solar level, etc. But the other thing was I I visited Bangladesh 15 times in 2010s because just as the developed world did not really understand how mobile phones would be used in poorest countries, I would like to question if, you know, AI can actually be a complete advantage if you're scaling what communities actually need to franchise because they've got desperate needs. And if we match the technology with the data, we should be empowering solutions. Thank you.

[03:27:28.06] - Åsa Johansson

Thank you. Do we have other questions? Yeah, Florence.

[03:27:35.22] - Audience

Thank you. A couple of questions. First for Gabriel. So you showed that some of the people who were in formal employment, you know, go to the informal sector.

[03:27:47.53]

I was wondering if you are able to distinguish, you know, self-employment or entrepreneurship that's out of opportunity or necessity, because one storyline around AI is that it will make it easier to start, you know, your own business and become an entrepreneur.

[03:28:06.08]

And I was just wondering if some of what you observed could be this, you know, or not, and whether your data allow to look at this. The other question is, you know, what room is there to upskill the people who were in clerical jobs maybe before, rather than have them, you know, maybe being demoted to low-skilled service? service jobs? And that relates a little bit to what Herman discussed, that, you know, you mentioned the lack of access to the digital infrastructure that prevents benefiting from augmentation from AI. But are there other, you know, kind of obstacles that prevent this upskilling? And then one question for Filippo. So again, relating to what Gabriel showed, that there is some disappearance of formal jobs, can you clarify if in your analysis you saw some reshoring of these, you know, kind of good jobs in emerging markets, you know, back to advanced economies where maybe they can more easily be automated or augmented? Sorry, augmented. And if so, in which sector? Did you see this? Thank you.

[03:29:26.21] - Åsa Johansson

Thank you. We take one more question and then we let the speakers answer. So I think it's Jakob.

[03:29:32.19] - Audience

Jakob Kirkegaard from the Peterson Institute. I'd like to just take a step back and just sort of focus on the title of the panel, which is AI and the Implication for Emerging Markets and Developing Countries, and ask a fairly sweeping question, actually. I guess it's sort of, if I sort of superficially look at AI and growth models in emerging markets and developing countries. And I sort of take a look at, say, the traditional flying geese manufacturing sector. And now these countries are having to deal with China. They're having to deal with AI-assisted automated manufacturing. So labor costs probably doesn't matter very much going forward. And then when I look at the services sectors, the growth models that we have seen, you know, call sector-driven, Philippines, other countries, obviously in India, the Infosys and onsite IT services model is probably also going to be somewhat negatively affected by AI. So I guess my question to

anybody on the podium is, what is the post-AI dominant growth model that these countries actually have available? Thank you.

[03:30:50.40] - Åsa Johansson

Over to the speaker. I don't know who wants to start. Maybe we go in the order these friends—Hernán first.

[03:31:00.14] - Hernán Winkler

Hi, thanks. Regarding the other obstacles besides the internet access, I think foundational skills is always a topic that comes up, right? To use the technology is not so much that everyone has to become an AI master. the next day. It's having the foundational skills to make the transition from one type of job to another. And that's a big challenge to help to build those educational systems and also lifelong learning. But clearly, I mean, something where developing countries are lagging behind. And going back to David's question, in some points, it's something that we have to do no matter what, whether there is AI or not. So it's one of those kinds of no-regrets type of policies. Regarding the post-AI growth model, that's a very difficult question. And I don't think we have the answer yet. I mean, for developing countries, there is a lot of uncertainty, like you said and like David mentioned, and especially because this is the type of technology that is exogenous to developing countries. It emerged endogenously in the rich world, but it's not solving a bottleneck necessarily in the developing world right now. So the challenge is how to think how this technology may help a bottleneck, which may be access to skills that countries, developing countries do not have.

[03:32:23.46]

But clearly, it's something that we need to pay a lot of attention to.

[03:32:28.12] - Åsa Johansson

Thank you. Maybe Gabriel, did you have some thoughts? A view?

[03:32:37.46] - Gabriel Marin

Sorry, first time on a P. Okay, so a couple of questions that I got. I can definitely refine the analysis to understand whether self-employment out of necessity is a thing or not. I'd have to check it and be annoying, but you can observe questions that relate to labor market, like, did I leave this job because of this reason or not? So maybe you can proxy that and maybe I can answer some stuff and I'd love to send that if I find something interesting. And then I will answer

the first question with the combination of the displays dominant growth model AI. I— okay, so people are displaced. So what? Like, what can we do? If this is a problem that this person working in the manufacturing is being displaced by this machine that now is AI-powered, can we think of ways to have that person have ownership of that machine? Like, maybe the answer should be thinking of some kind of AI fund. If I'm thinking of ownership of AI and AI is displacing, maybe those who are displaced should be having a share of that, no? So, and this is some work that I have been doing with Guillermo as well, AI ownership and growth.

[03:34:13.14]

So maybe the answer is more how the government can leverage its power, no, with protecting those who are most vulnerable through creating financial markets, like financial instruments that can help them hedge. Because this is not— this doesn't end up being just a short-term issue, but rather OK, these people are displaced today, but then their children might not even work. Like, in a very far future, we can think of AI just replacing everything. OK, then maybe we should be thinking about how we can leverage those rents for everyone. So that is my food for thought. But it's a lot of work for theory, for sure. No data yet. Thanks.

[03:35:03.04] - Åsa Johansson

Thank you. Interesting thoughts. So Filippo, what's your take on this?

[03:35:07.38] - Filippo Bontadini

Yes, thanks. Let me first try to answer this question about reshoring and the reconfiguration of production processes. I think, so to give you an example, what we find, and I think it's quite interesting, is that employment has grown the most in digital services. So ICT and computing services. And this is despite a massive increase in labor productivity, and an increase in final demand, which is what you'd expect. But the input-output linkage component is as big as the final demand contribution. So that tells us, for example, that those services is not just that they're being consumed as final demand, but they've become an integral part of production processes that might happen in all sorts of other industries. So in order to produce a car in Germany, you will need much more final contribution from digital services. That would be, for example, a way in which workers that, for example, were— or activities that were located inside the automotive industry, say, have been moved to industries of ICT services so that it's not strictly speaking a reshoring pattern. I think to do that, we'd have to tweak quite significantly what we do because reshoring means that you decompose final demand backwards.

[03:36:33.22]

What we do is downstream facing. So we look at where is the demand for jobs in a country coming from. And so in that sense, it's not exactly the same, but we do pick up on a lot of these kind of restructuring of production processes across sectors that change, therefore where employment is being generated. On the— I thought the flying geese, I think that that's quite interesting. And I think it speaks a lot with what we've seen yesterday about structural change as well. I do wonder, in a way, I'm torn between a very optimistic kind of view of, well, AI should democratize knowledge, make geese fly faster in that metaphor, or even if you want to use another animal, leapfrog faster. frog leap across entirely new production processes. And so therefore the sky's the limit. I do think that, however, these kind of technological changes don't happen in a vacuum. So even though you might have, say, at the very micro individual level, new ideas for a new entrepreneurial project, and in order for these innovations to scale up and have aggregate impacts, I think you do need a lot of infrastructure, not just actual infrastructures the way we, you know, data centers and whatnot, but also intangible assets and knowledge capital that exists, which is at the moment very unevenly distributed across the world and largely still located in high-income economies.

[03:38:10.02]

So I wonder whether the extent to which these AI developments can actually alter these development kind of parts unless they go hand in hand with a redistribution of other kinds of assets, which aren't necessarily AI, but more knowledge-based assets, I think.

[03:38:26.41] - Åsa Johansson

Thank you. David, what's your take?

[03:38:32.26] - David Kaplan

I think I'm only going to try to tackle the AI dominant growth model question. Although let me just say, I've been a little bit lost. I didn't understand the geese and the frogs, but Unfortunately, I have a very Latin American focus just because of my work, but I want to talk about the growth model we've had so far in our region. First of all, it's been largely demographics. It's been the working age population as a percentage of the total population has been growing. Also, I want to—

[03:39:07.04]

by the way, that's about to end. So we better find a different one. And then I wanted to talk about sort of the economic structure from a different point of view. In a different recent book we published specifically on labor markets, Making Labor Markets Work, which you should also

download, one of the things we make a big deal about is the very high percentage of workers in very small firms. So some things like unpaid work, which I've heard you mention it, is mainly sort of small family businesses, own-account self-employment, so self-employment but without any employees. So that's probably not entrepreneurship and just employment in firms that are 5 people or less. That's on average in Latin America and the Caribbean about 46% of employment, way over 50% in some countries. If we want to talk about that, well, the good news is that shields us a bit from some of these shocks, I guess, but it's a very low productivity way to do this. Again, once demographics stop being a key engine of growth, we're going to need to turn to productivity. I would argue— first, I want to talk about foundational skills. Which, of course, are obvious for the obvious productivity reasons, but we may underestimate the extent to which they could really transform the types of firms that there are.

[03:40:38.18]

Now we have people who can be managers of more modern firms. It's not just individual people are more productive. It is that, but it's also that may spawn a more structural change than realize. And then also the redesign of social protection, which was also mentioned, which first of all will protect us against these shocks, but also could potentially promote more productive forms of employment, among other things, by lowering the costs of formal salaried labor. In that sort of world, we would at least have, I think, the enabling conditions and a less distorted economy to have more confidence that we could protect people who are affected by AI, and AI could come in and be used for productive means, which I would argue right now is Far from certain. As someone who got caught up in the flight disruptions on Friday and the IDB, maybe it's not related, but we also had some kind of cybersecurity issue, which I know they're not cyber attacks, but maybe they are. And was affected by the email outage a couple days ago. with all the uncertainty regarding what AI is going to do, I think we could go back to the basics, try to eliminate the underlying distortions, which would almost certainly imply that AI is deployed in a more productive manner in our economies.

[03:42:28.45]

So I think I'll leave it there. Thank you very much.

[03:42:32.15] - Åsa Johansson

Thank you so much.

[03:42:32.52]

I thought we had some further questions, but I see that I have Yeah, I think we're about time. So it's lunch. So let's give a hand of applause to the presenters and the discussants.

[03:42:47.19]

It was a great session. Thank you.

[04:34:48.22] - Adam S. Posen

Hello everyone. If we can ask you to take your seats. I'm Adam Posen, President of the Peterson Institute, and it's my privilege to be the moderator of today's high-level panel as part of our 2-day conference joint between the IMF, the OECD, the World Bank, and the Peterson Institute on structural change and adaptation in labor markets in the face of AI and demographics. This is the aptly, if unprettily named, high-level panel It is the high-level panel because of the distinguished people on it, but more importantly, it's the high-level panel for 2 other reasons. First, it's an indication of the seriousness with which the 3 previously mentioned international institutions take this subject that they have sent some of their most senior officials to engage today. And we at the Peterson, but also everyone at the working level in those institutions on these topics, are grateful for that. And second, it's the high-level panel because we're going to be talking about things at a slightly higher level of altitude than the very serious research paper results we've had presented over the last couple days. To rapidly introduce the people in the order they'll be speaking, running from my left going down.

[04:36:18.50]

First is Mr. Bo Li, who is Deputy Managing Director of the IMF. He oversees the fund's work on about 90 countries. Prior to that, Mr. Li had a distinguished career at the People's Bank of China, including as Deputy Governor. He worked in monetary policy, and he played an important role in the drafting of anti-money laundering law. internationalization of the renminbi. In particular, I want to acknowledge DMD Lee has just coming to the end of his time, his term here as one of the deputy managing directors of the Fund. He's been here since August 2021, and he has contributed greatly to the work of the Fund and to international governance and economic affairs. So on this occasion, we just want to say we're grateful to have you in your final weeks here. Thank you for participating. Next to him is Stefano Scarpetta, who is the OECD Chief Economist, head of the OECD Economics Department, and has been G20/G7 Finance Deputy along with that since April 2026. Many of us know him because he was, of course, previously Director for Employment, Labor, and Social Affairs at the OECD for 13 years and ending up as Deputy Chief Economist at one point.

[04:37:49.36]

Prior to that, after a long career at the bank— I mean, excuse me, at the OECD, he worked several years at the bank again in labor market issues. Stefano is here both as a representative of top-level OECD, but also as a working, researching labor economics. So thank you, Stefan. Finally, as one distinguished person gets ready to depart, another distinguished person arrives in Washington. It's only us think tank people who are here forever. That got less of a laugh than I hoped. The new Chief Economist of the World Bank is, of course, Professor Michael Kramer. He is the joint winner of the 2019 Sveriges Riksbank Prize in Economic Studies in memory of Alfred Nobel. He is University Professor in Economics at the Harris School of Public Policy at the University of Chicago, where he directs his ongoing work of the Development Innovation Lab. Prior to that, he was a professor at Harvard University and MIT. And I had the good fortune and the humbling experience of being in first-year grad school school with Michael. So, but I'm delighted to have you here to talk today. So with that, um, let's try to set the stage. We, we have the technological boom in AI with all the potential very major positive effects.

[04:39:23.43]

We have very well-known risks of job displacement, issues of inequality, What do we think we can do that isn't passive, but it would be effective to deal with AI changing labor markets? And what do you think is the risk, upside or downside? We're economists, so risk goes in both directions. What do you think is the risk from AI that policymakers are probably underestimating today? Let me start with Bo.

[04:40:07.22] - Bo Li

Thank you, Adam. The, you know, we see several potential channels where AI could impact labor market, but also impact productivity and growth. One, of course, is automation. Through automation, workers may be displaced but productivity may get increased, especially for those routine cognitive or administrative tasks. Automation— the possibility of automation is quite high. Another channel is augmentation and innovation. With AI, we could have faster R&D. We could have better matching between skills and jobs. We could have faster diffusion of knowledge. Those can also impact productivity and growth, but also labor market. And finally, in terms of public sector, AI could also support better public service delivery by better targeting, better, faster transfer, et cetera. And especially for those developing countries with limited capacity, like limited technical capacity, limited administrative capacity. AI could be quite useful to boost the productivity of the public sector. So all those channels could potentially have a significant impact on labor market, but also productivity and growth. So in terms of risk, job displacement, of course, is a risk that could be very challenging for some countries. I would

actually highlight the risk of rising inequality, the potential widening divide across companies, across countries, and across workers.

[04:42:34.33]

You know, in terms of companies, those Those companies that are close to the frontier of the technology can probably benefit a lot, but other companies with, you know, smaller companies, less productive companies may struggle to make the necessary investment to capture the benefit. In terms of countries, countries with good digital infrastructure, with good labor force, and with good institutions could probably benefit more than those countries who are still trying to establish or build those foundations for growth. For workers, depending on the skill, some workers may benefit because their skills may be more complementary to AI. But there may be other workers that may face the The risk of automation and displacement. So those risks could potentially be very challenging for policymakers. So for policy, the key question is how can policy support AI's role as a leveler? instead of being a divider. So AI could also be a leveler, right? If you think about developing countries, there is a shortage of doctors. There is a shortage of good teachers. AI could potentially help them to help developing countries to leapfrog. So in those sense, AI could be a leveler. But AI could also be a divider at the justice Because different companies, different countries, different workers could have very different— face very different prospect in the face of AI.

[04:44:39.20]

So, I guess, for policymakers, the trick, the challenge is to figure out how to diffuse the technology as widely as possible, how to equip workers with the ability to adapt, How to manage the transitional risk? Those challenges, we as the International Monetary Fund are trying to— we, of course, we don't know the answer yet. We try to figure out, we try to work with the country, our country, our member countries to figure out, and this will be a challenge in the years to come. I stop here for now.

[04:45:20.14] - Adam S. Posen

Thank you.

[04:45:23.32] - Stefano Scarpetta

Thank you. Thank you very much, Adam. Good afternoon to all of you. And let me start by thanking you and all the Peterson Institute for participating in the organization of this event together with our colleagues from the World Bank and the IMF, but also for hosting us here in these beautiful premises. I think, taking your question, the first thing I'd like to say is that we

have all to be quite humble because on the one hand, I think our institutions are required to provide advice to policymakers. At the same time, I think the context of the evidence base is evolving very rapidly. And so we are navigating to some extent in quite a high degree of uncertainty. I think, and here I underline the term potential, but the AI in particular, GenAI, has indeed the potential to fundamentally change the relationship between human skills, work, and productivity. So to some extent, one of the risks we might be running, and we discussed that in the previous session, is to the previous evidence, the one for which we have the data about skill-biased technological progress, and maybe not capturing the fact that this might be somewhat different type of technology, not just a new vintage of the digital technology itself.

[04:46:30.44]

The reason why I'm saying that, and maybe aware of the macro hat before getting to the labor market, is that the evidence is very limited at this stage. So if you look at the impact of AI and GenAI, At the macro level, this is largely confined to basically the huge investment on AI by AI modelers, the few firms that actually are leading both in the US and to some extent in China. We see huge investment in the infrastructure, I mean data centers, computers, and so on. And this, of course, has spillover effects throughout the supply chain, not just in the US and China, but in many countries. We do not, apart from some specific occupational sectors, see the benefit of AI and AI on productivity. And this, to some extent, in my view— and probably I will get back into that in the second round of questions— the key point. So here at OECD, we have used exactly the same approach that has been pioneered by Daron Acemoglu and Philippe Aghion and others. We look at 3 dimensions, starting from the micro, about, to some extent, the capabilities of AI model. We look at exposure, and then we look at diffusion or adoption.

[04:47:36.10]

I think we know something about the capabilities. Because AI models have been improving dramatically in their capacity to mimic and to basically perform at the same, if not higher, level than humans. Just want to open and close quickly a bracket. We've asked GenAI models to perform our PIAC survey. This is basically the Adult Skills Survey. We did it in 2023, and already sort of the large language models were performing at the 90th percentile on literacy and between 50th and 60th percentile on basically mathematics. So this was in 2023. We know that since then there has been significant progress. So I think we know something about the potentially strong, very high capabilities of AI. We know the exposure, and I think we have discussed that because we can map tasks in which AI can be used already and to be very productive and who perform this task in different occupations. So we can get a degree of exposure. We can bring it for that to what would be the potential sectoral effect, and then we can aggregate it up depending on the sectoral composition at the aggregate level. When you do

that, and again, we are in the range between different estimates across the different literature, for the US, we are talking about half a percentage point to more than 1%age point annual growth in labor productivity.

[04:48:51.15]

This is sizable, taking into account in the US, labor productivity is around 1.5%. So potentially, this is a significant booster to productivity going forward, but it's something that is not yet realized. And I think the important point, in my view, the essential point, is the missing part of this modeling, which is the adoption, the degree of adoption. Where do we stand on the adoption of this technology? What can we expect going forward? And what are the type of policies that we need in order to promote a widespread adoption of the technology? But before I get into that, I just would like to highlight a few points because This is a technology that's evolving very rapidly. There is a huge concentration among a very few sort of companies developing the frontier models. 90% of the models are basically in the US and China. So one of the policies, in my view, is actually to promote and to maintain a strong competition. And there is competition at the frontier among the models who are at the frontier, but this is very much restricted to a few companies. So one of the challenges going forward is how do we maintain competition?

[04:49:49.58]

We make sure that there is a sort of the ability of everybody to enter this market and to compete. Second, I think there is something interesting because there are 2 approaches, very different. On the one hand, we have the proprietary approach, which is led by the major companies in the US. On the other hand, we have the open, basically, weight models, which are basically, again, at the frontier but not at the top of the frontier, basically cost less and perform relatively Which one of the two will dominate is an open question because given the immense investment by the proprietary companies, there could be some disappointment on the market about the return to this huge investment if they do not materialize in the short term. I think in previous technological revolutions, we have seen that there was a bubble before actually we saw the benefit on productivity. We might not be completely exempted from actually going through the same process in which Maybe there is too much exuberance, there is a correction, and then we see sort of the diffusion. But I think the important point I wanted to underline here is that we know also what might be some of the barriers for further adoption.

[04:50:52.40]

We see now the use of AI more in advanced knowledge sectors and much less so in less advanced knowledge sectors. We see more use in some of the advanced economies compared to developing and emerging economies. So we see the potential of a growing divide associated also with GenAI. But I think, as was said in the previous session, GenAI has potential. This is sort of the evidence from David Autor, the potential to actually boost the capabilities of the middle-skilled worker, not necessarily those at the higher end of the skill distribution. We see that on the difficulty that some college graduate young people have in entering the labor market. So there could be a difference completely. compared to the previous vintage of technological changes. So diffusion, in my view, and adoption is a key word here. And there is an area in which policy, in my view, can play a major role. The only thing we know, and I think the evidence is pretty clear, that AI, like all previous digital technology, does not mean simply plugging into the system and then suddenly you see the productivity benefits. It requires basically a change of the way in which companies are organized.

[04:51:57.29]

requires different type of skills or new skills, but I think require basically investing in skill for the entire workforce, not just for a few workers in some specific function. So this is a massive endeavor that I think we have to, we have to consider. And maybe I'll stop here and maybe we'll go back to the adoption in the second round. Thank you.

[04:52:19.01] - Michael Kremer

Thank you. You know, in terms of risks, there are obviously concerns about risks of serious AI catastrophe, including through engineered pandemics or through making conflict more costly or more likely. I'll leave those risks to others who are more specialized in them, and I'll focus on the risk that we fail to realize the full potential benefits of AI for low and middle income countries. And then we get the scenarios of widening inequality that Beau Le has mentioned. So I think that in terms of the potential, AI and digital tools have tremendous potential for low and middle income countries, both to increase productivity, creating more and better jobs and to improve public services in services such as health and education. Without going through trying to be exhaustive, AI tools can be very useful in digital agriculture. Farmers are some of the poorest people in the world, some of the most exposed to changing weather patterns. They could get advice, scientific advice on weather, on appropriate seed choice, on soil chemistry, a whole range of different decisions. Teachers who sometimes have quite low human capital levels could be empowered with taking advantage of the latest— or detailed lesson plans that take advantage of the latest knowledge in education and that could even be tailored to the student needs.

[04:54:21.53]

And they could even get AI teacher coaching to help improve their skills over time. Community health workers. There's shortages of highly skilled health workers in a lot of low and middle income countries. Community health workers could handle many more types of conditions than they can currently handle if they had some appropriate to— access to appropriate AI tools. And they could refer other conditions. Government services could be made much, much more productive. For example, targeting was mentioned. Bodil mentioned targeting of of transfers, but one could also think about taxation systems or procurement systems being made much more efficient. And judiciary, there's remarkable results on speeding up of clogged judiciary systems are a big toll, create a big impediment to business. Those could be speeded up not by having AI judges, but simply by handling the caseload and the assignment and reducing delays. Potholes could be identified by satellite data and then automated processes could be put in place to address them. There are all these opportunities, but I think there are some very important market failures that mean that we can't count on these things, these tools being developed. There will be some cases where commercial incentives will be enough under existing frameworks for the private sector to both create and adapt and deploy the software tools that are needed.

[04:56:09.01]

But I think there are a wide set of cases, including many of those I've mentioned of services for very small businesses and services for governments, where incentives for really high-quality products And developing those high-quality products requires not just— you can code very quickly right now with AI, but you really need to test these, you need to refine these, you need to do a lot of work with users, you need to adapt the social systems, including government systems in the case of software for government. All of that takes a lot of resources. And I don't think incentives for that are typically large enough to do— to get really high-quality systems. We do have a lot of proliferation, in fact, of low-quality, not very well-tested AI products, which, you know, some of which look amazing on a slide deck, maybe look amazing with a little bit more examination. And, you know, there's great marketing material, looks wonderful. But These often are not actually tested to see if they work in real-world contexts. And there's not much effort to adapt them to user needs. Things like device availability, things like the— something can look great if you can type in, give me a lesson plan that does all these things.

[04:57:34.19]

But is the teacher actually going to do that in a low-income context where they— And is that going to work with the Ministry of Education? All of those types of development of high-quality tools and integration with government systems, those things often don't exist. So I think in part

because of that, while there are lots of great examples you can point to that look very impressive, if you actually ask how many of these have scaled, it's a much, much smaller number. Now, I think the fact that there are weak commercial incentives to make— to do the full investment means there's at least a prima facie rationale for public or philanthropic investment to create these tools, particularly open access tools, to test them, to refine them in a number of different areas. But there's another important— how to organize that is a difficult question. And there's— it's important to note there's an asymmetric information problem as well. It's often difficult— each developer is going to be touting the benefits of their system. It's often very difficult for funders or governments that are doing procurement to judge product quality. And so that has multiple bad consequences. First, there can be a race to the bottom.

[04:59:07.19]

Producers don't invest in producing a high-quality product. Rather, they invest in producing something cheap and then invest in marketing or establishing connections. And even in some ways, I think the equilibrium that we're in is not so much that bad quality products get widely adopted, although there's some of that. The equilibrium is one in which governments, they hear stories, things have been tried. They hear lots of bad stories, things have been tried and they didn't work out. And therefore, it's a classic market for lemons situation. Governments are hesitant to adopt because they know there are many bad quality products out there. And then that in turn deters the investment because even if, say, a philanthropy wanted to develop a high-quality product, are they gonna get demand for that product? Well, they don't really know. So I think there's a lot of market failures around this. I do think, you know, why did I pick this example of all the risks out there to highlight? I think this is something that with appropriate investments, fairly modest investments, and with some concrete policy steps. Those of us in this room may not be able to deal with the problem of engineered pandemics.

[05:00:28.35]

So I think that's a serious problem. But I think many of us are involved in economic policy with fairly modest commitments of resources, fairly modest policy changes, I think we could get a lot of creation of tools to both increase productivity and improve government services.

[05:00:53.58] - Adam S. Posen

Thank you very much, Michael, Stefano, Bo. For a high-level panel in altitude, you're doing a good job of getting deep, which is not a surprise given all of you. So let me take up some of what Michael was saying and turn it to the others initially. What do we think is the keys to unlocking the potential of AI for the developing world, for basically the world that isn't China and

the US, but the developing world in particular? And essentially anyone who's not at the technology frontier. And let me start, Stefano, because you, while in your time at the OECD and some of your predecessors in your position and colleagues in the room worked on important research done in the past where the OECD made the case that a big chunk of differences in productivity trends across members was the extent of diffusion from leading companies to the rest of the companies within countries and so on. So I mean diffusion not just along the international dimension, but the inter— inside a country in the economy dimension. So maybe you could say a few things about what can get SMEs or other parts of the economy going in the AI space.

[05:02:27.53] - Stefano Scarpetta

Thank you, Adam. That's, I think, is an essential point. I think the latest evidence still suggests that there is very little spillover effect from frontier firms to busy follower firms in most of the advanced economy for which we have data. The question mark I have is that maybe GenAI, generative AI, can make a difference in terms of spreading out the benefit of productivity growth across a wide range of firms. And I think the crucial question is, what will make adoption actually be widespread? Frankly speaking, we do not really have the answer to that because we do not have the data. What we know is that in the US, the Census Bureau, there is a module in which they look at employment and the adoption of AI. Even in the United States, if you like, adoption is around 25% of companies using it. When you look at intensive adoption, so the kind of adoption that will make a potential positive impact on productivity, you're down to basically a single-digit number. So what we have done at OECD to try to put on the table a bit of evidence, we have run a survey to 62 countries.

[05:03:27.56]

I will not fully spoil the presentation that Felice will give later, but I think you won't be surprised if across these 62 countries, actually many people are using GenAI in their own private life. When you look at how much this is used in the workplace, you go down to basically 20% in the country that use it the most. And when you go into what you can call intensive use, so the kind of use that potentially would have an impact on productivity, again, you are down to a single-digit number, 6, 7, 8%. So we are still in an early phase of the adoption, which would be actually what will make a difference in terms of AI really boosting economic growth and productivity across the board. In several countries, I think Michael made a very good point that technology has all the potential to benefit many countries and actually benefit private sector firms as well as improving the efficiency of the public services. There is a huge potential out there, but I think we are still in an early stage of this process. And the question is that what is preventing a more— given the capability that I mentioned before, what is preventing a more widespread adoption?

[05:04:29.59]

So here, let me just refer to a few of other surveys we have run. We did one actually in '22, Leader Did We Know that at the end of the year, basically, there was the introduction of the first GenAI model. Back then, cost was one of the barriers for adoption. And the second one was actually the lack of workers with right skills to work side by side with AI. Now, we run a similar survey, this time focusing on those firms that potentially have more difficulty in using technology. So the small and medium-sized enterprise, those with less than 150 employees, we did that in 2023. in 7 countries, the costs were no longer the major issue. Actually, the number one issue was that I really do not understand how AI can use my own business. So there was a lack of understanding of the potential of the technology to really make a difference in terms of boosting productivity in the company. The second one, again, was the lack of workers with the right skills. So I do not feel that my workers will be able to use this technology. But the third one was actually, I'm very concerned about privacy concern

[05:05:30.43]

biases, and managing of risks. So if something goes wrong in the use of this GenAI model, which has fantastic potential out there, who's going to be managing the risk? I'm responsible as a company using it. Is there those who designed the model or actually the data that we are using? So in my sense, and I think throughout the conference we have been talking a lot about training, retraining program, my sense is that we have to train the workers. There is no question about it. The second is actually providing provide guidance to managers, especially in small, medium-sized companies, on how to use AI in an effective way. And that's easy to say, much more difficult to actually do in practice. Here, different countries have different approaches. The European countries have actually regulated AI. The US is using sort of voluntary practices and so on and so forth. But there seems to be a bottleneck, which is related to the fact that companies do not fully yet appreciate the potential benefit of the technology. but also are very concerned about possible risk and how they're going to manage this risk. Now, I switch hat and I put the labor market hat on the top.

[05:06:30.36]

Whenever I hear that we have to invest massively on training and retraining, I get a little bit concerned because I think all of the labor economists in the room know that the evidence about the effectiveness of training programs for adult people are very low. So the returns on this investment, I mean, there is not a very high score. Now, we have made progress by targeting. There has been significant progress in many countries, but the simple notion of basically scaling

training and retraining everyone or most of the workers is really something we have never experimented even close to that. I mean, even in the countries that spend more on training and retraining, you're talking about a timeframe, I mean, 60% of the workers, but from time to time. Now, and just putting on the table as an as a potential— I don't have the evidence on that— is that AI can actually be a tool to improve the way in which we deliver training for the adults. You can think about much more tailored training courses, much shorter training courses, training courses that can be delivered basically remotely instead of the large, long training course that basically many workers do not even dare to participate in, or maybe their employer is not willing to let them do it.

[05:07:41.50]

So the AI can actually be a tool to improve the targeting and the efficiency and maybe the outcome from training programs. But if we want to avoid further divide among high-skilled and low-skilled workers across countries and so on, I think human capital becomes one of the essential factors here. And we have to think about how we are designing training courses that can benefit many workers, not just a fraction of them. And the other big concern which which has to do with the labor market is what we discussed earlier on, which is if we are moving to a system in which entry jobs are diminishing because basically many of the tasks performed in entry jobs, even sort of college graduate type of jobs, are diminishing, then how do we allow the young people to get the training which is needed in order to move on in their positions? And here again, I heard some discussion in some OECD countries about working with the lawyer to develop apprenticeship programs in which young people basically learn and then at the same time also work in the company to some extent to allow them to gain the kind of on-the-job experience that will eventually lead them to move on into higher positions.

[05:08:47.32]

So I think we have some evidence about some of the bottlenecks, some of the barriers for adoption, but I think adoption, in my view, is really the key word here because at Otherwise, we might see sort of the bursting possibly of the AI bubble and not much of the productivity potential that is certainly out there, but it's not yet materialized.

[05:09:08.42] - Adam S. Posen

Great. Now, Bo, the IMF has essentially every country in its membership, has a slightly different remit than the OECD. Can you talk a little bit, but it also has a similar narrower remit in the macro financial space. Can you talk a little bit about where you see, or the Fund sees, the

binding constraints on this kind of diffusion and adoption, particularly coming out of financial infrastructure capacities, but whatever you think best?

[05:09:46.04] - Bo Li

Thank you. I very much agree with Stefano. You know, adoption is the key And in terms of adoption, we think 3 factors are important. One is called preparedness. Second is exposure. And the third is access. In all 3 dimensions, developing countries are facing challenges. Let me start with preparedness. We have prepared a— our research department has prepared an AI Preparedness Index. And uniformly, developing countries score, you know, lower, actually significantly lower than advanced economies. There are 4 dimensions of this preparedness. One is digital infrastructure. The second is education, training, human capital. The third is innovation and economic integration. And the last one is regulation and governance. So on these 4 dimensions, as you, you know, you could guess, many developing countries are facing challenges. In terms of exposure, developing countries are also doing worse compared to advanced economies. You know, advanced economies have a larger share of service and also more workers in managerial and professional occupations. So advanced economies are better positioned to capture the benefits of AI compared to developing countries, many of which have a large sector in agriculture or traditional manufacturing. So in terms of exposure, that's the second factor. the developing countries, the average exposure, we call it 26%.

[05:11:50.04]

That means 26% of the occupation in developing countries are exposed to AI as compared to 60% in advanced economies. And finally, in terms of access, access to semiconductors, access to large-scale computing, access to good data platforms. Again, developing countries are facing challenges. So what can developing countries do? Our advice is to focus on these 3 factors: preparedness, exposure, and access. But that does not mean developing countries have to buy expensive semiconductors, have to establish big computing centers, et cetera. That's why economic integration can be useful and important. If you are integrated in the global economy, you could be able to buy, purchase, buy computing from— From the US. from large international providers. You could establish a good data platform with technical assistance from the international community, et cetera, et cetera. So economic integration, we think, is important for developing country to benefit from AI. Of course, developing countries still need to invest. With their limited fiscal space, they should invest their precious fiscal resource to the most important part of this matrix, which is digital infrastructure, training, education for workers. These are probably most important areas for investment. Of course, in addition to public investment, developing countries also need to think about how to attract private investment into this AI preparedness.

[05:13:57.06]

So, as I mentioned before, we do think that developing countries have the potential to leapfrog in some of the bottleneck areas, including government service, including agriculture, including health. healthcare. But whether a particular developing country can realize that potential depends on the policy environment, the investment, and the kind of economic integration that this country could achieve. And we are, again, we are working with countries. We are hoping that with international partners developing countries could overcome some of these bottlenecks. I'll stop here.

[05:14:46.11] - Adam S. Posen

Thank you. So coming back around to Michael, and don't worry, next time around I'll start with you, but just picking up on the very tantalizing thing you said initially about some affordable, achievable policy missteps could have a big impact. both Bo and Stefano have sort of gone their direction. Can you give a more specific couple examples maybe of what you have in mind?

[05:15:13.30] - Michael Kremer

Great.

[05:15:13.43]

Yeah, no, I'd be happy to give an example. And let me start with one example, which is an area that I've been involved in, but I think it also illustrates some general principles. So my first answer was sort of very abstract market failures of various types. So I'll try to make up with that with something more specific here, but then I'll go back to abstraction a bit. Okay. So here's the setting on AI weather forecasting and for agriculture in particular. So billions of people worldwide depend on rain-fed agriculture. Weather in the tropics typically involves rainy and dry seasons rather than summer, fall, winter, et cetera. So farmers have to time their planting to the rainy season. Ideally, you want to plant as early as possible in the rainy season to maximize your growing period and your yields. On the other hand, if you see some rain and you think, okay, the rainy season started, I'm going to plant, but the rainy season, it's just some scattered rain, not really the rainy season. That's a financial disaster for the farmer because the seeds don't germinate. If you put any fertilizer in there, the value of that's lost.

[05:16:33.31]

The farmer, the family, farmer family can go hungry. They can have to sink into debt. They have to pull kids out of school. So trying to anticipate when the rainy season is, critical. There's often a lot of traditional knowledge about when the rainy season's likely to hit, but that traditional knowledge may not work anymore because weather patterns are changing. Now, there's a half dozen randomized trials from around the world suggesting that when farmers have access to high-quality scientific weather forecasts, they change their agricultural decisions. They respond to it. This being one example, but obviously there's ways farmers can respond to weather throughout the year. And that increases incomes, it reduces the risk of catastrophic crop losses. The financial gains vary by context and by season, but a reasonable— the cost of developing weather forecasts and transmitting them by SMS messages to farmers, that's very, very low. So, the benefits are likely to outweigh the costs by more than 100 to 1. So, there's— that's the setting pre-AI. But AI makes this much, much— really creates much more potential because there's been a revolution in weather forecasting. A few years ago, AI models weren't very good.

[05:18:08.53]

And people thought it'd be a long time till they would surpass regular models. Well, it turns out they moved extremely quickly and you can have better accuracy, you can have longer lead times that you can predict weather. And these run— developing the models takes time. Running the models, it's 80,000 times faster than a a physics-based model that you run on a supercomputer. You can run these on desktops. So you've got the potential for really democratizing access, but also you have the potential to develop forecasts for specific applications. If you're a farm— I ride a bicycle, I want to know when is there a 20-minute break in the rain that I can ride home? If you're a commodity trader, you want the same information, a millisecond before all the other commodity traders. If you're a farmer, you want to know, are we going to have— is the rainy season really started, or are we going to likely to have a dry spell in the next month? That was— now, that's not something that most of the global AI weather models have been designed for. So there's multiple open access AI weather models, probably the— you could argue about what's the best one in the world.

[05:19:36.48]

Is it European Center for Medium Range Weather Forecasting? Is it one of the Google models? Is it another one? But it all depends on what you're trying to predict. Well, none of them were designed to predict the start of the rainy season. And what But it turns out that if you take these open access models that already exist and you combine them, and you combine them with some historical data, some statistical model, that together they can do a very good job. So this year, the Indian government created a model that combined their own sort of older, older but not worse, just different physics-based model with open access AI models. and a statistical model,

getting the best elements of each one. And the combined model beat any individual model in accuracy and lead time. There was a precursor version of this that was just using the open access models that are publicly available that predicted an abnormal pause in the progression of the Indian monsoon in 2025. And that— so the monsoon starts in the south of India, it moves up, It just paused one year and didn't move up further. That's a very dangerous situation for farmers.

[05:20:55.19]

The model predicted that successfully and no other model did that. And that was sort of the 2025 model. The 2026 model is better than that. So that's one element, predicting the weather. But again, to pick up on the point about you need organizational changes throughout. The Indian government also piloted and conducted A/B tests of messages. How do you transmit this information to farmers in a way that farmers will understand it? And that proved critical to boosting message comprehension. And how do you roll it out so that you can send messages in the right language for different parts of India and do that all efficiently? And they put that in place as well. So this past year, they sent messages to 50 million farming families. We have preliminary evidence on the impact of this. Looks like 7% to 8.5% increase in farmer profits over the past couple of years. That's early work, so I wouldn't swear to it. Okay. So Ethiopia is now doing this as well at a pilot level. Let me note a couple of lessons of this. First, these were open access systems, and the open access nature created several advantages. I'm not saying everything in the world should be open access, but I think there are advantages to having some open access benchmarks.

[05:22:29.16]

First, that meant India could keep improving the system each year. Second, it meant— and they did it to make changes. They didn't need to renegotiate a contract with the provider. They could just do it. They also wanted their own model. So there were initial open access models, but then they created their own Indian national model, including their own data, their own physics-based model. And that sovereignty is important to many countries. Third, as long as there's an open access model, that's initially created, then this can really be a global public good because Ethiopia was able to take that open access model, make other changes to it to reflect Ethiopian data and the Ethiopian circumstances, and get something that would be optimized to predict the rainy season in Ethiopia. So I think that having open access is important. That's the first general lesson. Second, to get impact, you need to combine the model development with all of these other elements— human-centered design, A/B testing, and intensive work to integrate with government systems. Then a third element, which might not be as obvious, is you have to be able to select the right model. And this goes back to the issue of the proliferation of models.

[05:23:52.30]

There's all sorts of claims out there about models. And actually, you know, these major, you know, European Center, Google, all these things are all— they all publish papers on what they do. They're all pretty high quality. There's also a bunch of startups that make very strong claims without much evidence based on them. And I think in some cases, you know, deliver pretty dangerous advice. Having, you know, together with colleagues, we just for the case of weather, we issued a short paper that listed some of the ways of judging the quality of models based on how well they predict and some very basic things like if you're going to make claims, say what years you tested the data on or what exact data you used to test it. Just as AI models have benchmarks, We need systematic independent standards to help judge AI tools for public services. And then that needs to be complemented with some testing in real-world conditions. So I think that with these, the total cost of this was a few million dollars. And if you think about reaching 50 million farmers so far, this is something that I think can be expanded greatly.

[05:25:14.43]

But I think that once farmers are checking their phones for the weather information, it's then you can start to add on a whole range of digital agriculture services. That's one example. But I think, as I guess one thing which I won't discuss right now, but developing this required a partnership. developing and scaling required NGOs, required firms, it required philanthropy. And, you know, in light of my new job, I'll say I also think there's an important role for multilateral development banks as well as, of course, national governments.

[05:25:55.06] - Adam S. Posen

Thank you, Michael. Just an information point for both our set of conference participants in person and those you watching online, this, because it's such a high-level panel and high quality, I hope you'll agree, we are— the session is meant to go till quarter after the hour. And so we will be coming to the in-person audience for on-the-record questions in about 10 minutes. Let me just do one more round through the panel. the group. And as promised, I'll start with Michael. So far in this panel, we have been emphasizing the making use of AI, the productivity impact, the ways to make it be more widely diffused and adopted, all of which is great. But obviously, a lot of concerns raised in the course of this conference and frankly throughout the world are about the impact of AI on currently existing labor, the impact of AI on unemployment, on churn, on youth versus older workers. I mean, again, some of you have made reference to this. I'm not suggesting you've slighted it, but I'd like to sort of focus us one round on the labor market. market impacts of AI. And, excuse me, so starting with Michael, you know, something

else which you've worked on in the past, and obviously a lot of colleagues at the Bank have worked on through the years, is the question of quality education for workers to be able to adapt.

[05:27:46.29]

That's partly the retraining issues that Stefano already mentioned, but it's also much broader, longer-term issues of education, both in nature and quantity. So do you see the same kinds of positive possibilities you just mentioned in the agriculture space for AI to help us with education in a way that helps workers?

[05:28:11.13] - Michael Kremer

Yeah, I absolutely do, actually. I think As background for those who are not familiar with the context, there's been tremendous expansion of access to education in low and middle income countries. Most kids are at least getting primary school and increasingly secondary school. And that's new. But unfortunately, learning levels are stagnant and at best, in some cases, there seems to be declines in learning levels. Now, there are a number of factors, but it's hard to maintain quality when you expand very rapidly. There are, even without technology, there's some very promising evidence-based approaches. There are a number of them, but let me give an example of something called structured pedagogy. Structured pedagogy teachers get detailed lesson plans. So even if the teacher doesn't have that much subject matter knowledge or that much— isn't amazing on the pedagogy themselves, they can— with a structured— with a detailed lesson plan, you get tremendous gains. Actually, let me say, it's a detailed lesson plan plus teacher coaching plus some other elements.

[05:29:35.16]

OK.

[05:29:36.35]

That you get— there's versions of that without technology. But you can imagine if you combine that with technology so that teachers are getting their lesson plans on a tablet, well, then that could be adjusted to respond to students' learning levels. And that's been shown to be very effective. You could— an expensive element of structured pedagogy is the teacher coaching. You need skilled teacher coaches who have to travel around to provide coaching. Well, you can't eliminate the human element in that, but you could do a combination of a human element plus— I think this technology is on the cusp of being ready— of having the teacher record on the tablet the lesson that they're giving and then get automated feedback as well together with

the human feedback. I think once those systems with AI, once those systems are developed for one curriculum, one language, you could— and it's fairly inexpensive to do at least a first draft translation of all those materials to another country, another curriculum, another language. And then that'll obviously need to be adopted. further by humans. But I think this is something that could dramatically improve learning levels. And that empowers teachers to do a better job.

[05:31:05.44]

But it also means that their students, when they hit the labor force, are going to have much more of the preparation that Bolu and Stefano were talking about as things that are important. to take full advantage of AI in the workplace.

[05:31:23.56] - Adam S. Posen

Thank you very much. So shifting to Stefano, just broadening this a bit out with your labor expertise hat, we have had a number of events in this room in recent months with Peterson Fellows Martin Chorzempa, Anton Korinek, Jed Kolko, who's of course one of the co-organizers today, Anna Stansbury, but also Luis Garicano from LSE, talking about how much you can unpack jobs into tasks versus— Luis and his co-authors have the concept of messy jobs. Obviously, they're not alone in thinking about these concepts. But just more broadly, How does that complexity and ability to decompose jobs into component tasks affect design of education, design of training? People are now starting to talk about things like, well, we're not so much worried that jobs will disappear, but the jobs may get much more commodified in a very sort of Marxist sense. It's broken down into much more oppressive tasks. So anyway, there's a whole set of issues there which I know you've thought deeply about, so maybe you could say something in that regard.

[05:32:51.41] - Stefano Scarpetta

Thank you, thank you very much, Adam. I think in much of the work we do at OECD, we move away from jobs and occupation. We look at tasks that are performed in the jobs because, I mean, if you want to have a good understanding about how the labor demand is evolving, you really have to look at what tasks are in high demand, which tasks are in less demand, and how they're bundled in different types of jobs, and how this bundle can actually be changed. I think I want to build on what Michael will say about all the potential of actually using GenAI also on the education front. But I think given the speed of evolution of the technology itself, we have to think about how to adjust the education system, but also how to, to some extent, improve the adult learning system itself. And here, both on the education and on adult training, there is a sort of a possible syndrome that we might run. Since we don't know exactly how the technology will

evolve, let's wait and see before we actually invest in changing the education system or changing the training and retraining program.

[05:33:48.40]

And I think by doing that, I think the risk is that we're always lagging behind, that we are not really providing the kind of services that actually are needed. So again, I think the evidence is important here. If you look at the the online job vacancies, which cover a wide range of occupations, not all sectors, but I think for advanced economies they have a pretty good coverage, you see a tremendous change in which kind of tasks are required by the workers in these jobs. And I think in a matter of a few years, you see a significant decline in basically coding, and you see a significant increase in what are more properly, if you like, human skills. So the ability to work in team, creativity, problem-solving, critical thinking. It's quite impressive to see how the demand for labor at the task level has changed dramatically. And this, I think, has to be basically brought into both education, but also, of course, into training and retraining program. Now, the evidence from the survey I was mentioning before is that the sentiment varies a lot among workers, whether or not in their own company AI is used.

[05:34:50.44]

So those who are not tend to be extremely scared about the use of AI. Those who are actually using AI in the workplace seem to be much more confident about the benefit of that. If you go one step further from another survey, you see workers that were told that AI was used in the workplace, then their performance, but also quality of job, perceived quality of job, improved significantly. If then they receive some training to work with the AI tools, then their performance at the same time improved even more. tend to improve further. So there's a lot to be done, again, not only in designing the right training program, but actually in involving the worker into the process. Now, if I look at the advanced economies, again, the OECD countries, and they look at how many— what is the share of training courses for AI users, you're talking about in between 3% and 5% of basically the different range of offer in terms of training for adults. Very little. There is quite a bit on basically to become AI developers, there is much less on training for AI users. That's where we have to invest, in my view.

[05:35:51.53]

This is a combination of small, basically a minimum of digital skills, plus probably some well-targeted kind of training depending on the type of job the individual works. What is an additional source of concern is that in many countries, and this was part of this conference, we're also aging. And with aging, there are 2 things. One is that we know that older workers tend to be less

innovative, training and retraining. And we also know that older workers tend to be less mobile. And I think with AI and the diffusion of AI, we can only expect more mobility. So not only sort of changing the bundle of tasks, but also the job will evolve and workers might have to move from one company to the other. So I'm trying to bring the AI and the demographic changes together because they're happening at the same time. And I think there is a number of opportunities. opportunities, but also a number of challenges that we have to take into account jointly, I would say.

[05:36:43.13] - Adam S. Posen

Thank you. And just again, a small programming note for those of you who have been here or are watching, or those of you who will be coming back, because all of our videos and papers are posted permanently on the PIE website and the partner institutions. The demographic theme and, and challenges have been a day of this conference, we simply chose to focus this panel more on AI, but the demographic challenges are a key point of research and policy as well. So just before we go to questions from the audience, Beau, you have already mentioned a couple times your own and the Fund's concerns about the impact of AI on inequality of incomes, inequality in labor markets, and other kinds of outcomes. Could you say a bit more about which workers, whether it's by country, by age, by job type, that you fear are going to be most affected, and more broadly, other steps and concerns you have in this regard?

[05:37:48.33] - Bo Li

Thank you, Adam. Let me start by saying that at the aggregate level, we have not seen large-scale job loss. We have not seen significant displacement at the aggregate level, partly because there are a lot of new jobs being created, partly also because we are still at the early stage of AI adoption. We haven't seen significant productivity gain yet. And there is also something called the Jevons paradox, where as productivity increases, products get cheaper, the demand may go up, and there may be more demand for labor. So regardless of the reason, the aggregate level number, we haven't seen significant job loss yet. However, at the sectoral level, we see Not all sectors are equal, and not all workers are equal. So that comes to your question, Adam. There is a distributional impact. So let me maybe give several examples of potential losers versus potential winners. One potential— one sector of potential loser is those sectors that are highly exposed to AI, like IT, in particular software, marketing, finance, those sectors, we already start to see weaker labor demand, fewer vacancies, and slower wage growth. Our research has shown that in the US, level in AI-vulnerable occupations are lower in regions with higher AI skill demand.

[05:39:55.47]

We have actually have number, it's like 3.6% lower compared to other regions. So we already see that when you have sectors that are highly exposed to AI, you could have, you could face job losses. That's number one example. Another example would be young and entry-level workers. Young entry-level workers typically undertake tasks that are easily automated. And so this has implications not only for job displacement but also for for training and growth, because if younger entry-level workers, if they're not getting the job, they will not get trained and they cannot grow. So that has medium-term productivity implications and also implications for inequality. Let me give a final example of a potential loser, which is middle-level, mid-skill, middle-skill workers. The middle, the mid part may get squeezed. Again, our research shows that for high-skilled workers, they could utilize AI to be their assistant and to support their work. So they don't have to rely on— traditionally, they rely on assistance, rely on support from the supporting department. Now they rely on AI agents. And so the middle-skilled workers may get squeezed. Actually, the lower-skilled workers may not be too bad, actually, because if you have high-skilled workers earning high wage, they're going to demand services.

[05:41:42.49]

They're going to demand— they're going to go to restaurants more often. They're going to go to vacation more often. So for those low-skilled workers in the manual sectors, Right now, AI cannot replace them. And they actually benefit from the high-skilled workers' higher wage. So what gets squeezed is in the middle, the middle-skilled workers. So that's another potential group that may face challenges in the age of AI. But that's a potential losing group, but there are also potential winners in this AI age. If you look at the people with AI-related skills. Our number shows that, for example, in UK, workers with AI skills command a wage premium of on average 7.5% to 8% wage premium. In the US, AI developers command a wage premium of 8%. Even those non-AI new skilled workers, they could get a wage premium of 2.5%. So if you have new skills, you may benefit from AI. You may benefit from AI. So there is an uneven distribution of impact that has policy implications for government in terms of which group to support, what kind of training can be provided, and what kind of adjustment educational system needs to make. can adopt to try to support young people, support workers.

[05:43:21.39]

I stop here.

[05:43:23.24] - Adam S. Posen

Thank you. This has been a very rich panel, but we have an audience of scholars, practitioners, researchers, and even multilateral development bank staff. If you would like to pose a question

or a very brief comment, Questions can go a little longer than comments. To the panelists, please identify yourself. To the front tables, Jessica holds a roving mic. There's also a standing mic at the back. One other thing, even if you spoke earlier at the conference or you're world famous or both, please identify yourself for those who tune in online. Anybody?

[05:44:09.15]

At the back mic, please.

[05:44:12.56] - Audience

Hi, I'm Jed Kolko from Peterson.

[05:44:15.26]

So much of the research on AI in the labor market and the discussion today focuses on the labor demand side.

[05:44:22.47]

And a lot of the policy response and policy concern is around managing possible shocks to labor demand from AI. But I'm also curious to hear your thoughts on how you imagine possible impacts of AI on labor supply to be?

[05:44:39.25]

Thank you.

[05:44:40.04] - Adam S. Posen

And maybe we'll get a second question and then let our panelists respond.

[05:44:45.26] - Audience

Hi, I'm Gabriel Dimambine. I'm at the World Bank. My question is for Stefano. Stefano, you talked about the critical role of management skills and knowledge in adoption, and you said you were optimistic. I think that AI could be used to enhance worker skills around AI, you have similar optimism around the possibility of using AI to enhance management skills towards AI.

[05:45:11.40] - Adam S. Posen

My thanks to both for their concision. As I see one more person at the mic, we'll gather that question and then have responses, please.

[05:45:20.10] - Audience

Rebecca Benson, Auckland University of Technology. My question was around Stefano's comments on managers who are maybe unsure of how they can benefit from AI, are there any policy levers to kind of help them see what's possible and to actually make that possible?

[05:45:39.01] - Adam S. Posen

Okay, great. So thank— again, thanks to all 3 who were very concise and clear. Maybe Bo and Michael can respond to the first question, and then Stefano can take on the 2 referred directly to him.

[05:46:04.09] - Michael Kremer

Thank you. On labor supply, I earlier discussed the impact that AI could have in strengthening education. Let me note a few other ways that I think it could boost labor supply. I think you can boost labor supply by improving health as well. and thereby improving human capital. I think it could help with upskilling and on-the-job training. If you think about, for example, I saw a paper recently on the problems that are caused by difficulty of getting repair services for machinery in rural areas. Well, and argued that had a significant macro impact. Well, repair— and it also caused people to all want to use the same model of equipment and thereby reduced competition in supplying capital equipment. Well, if everybody has access to training videos and can video them and can take pictures of themselves trying to do a repair and get step-by-step advice, that's going to increase the amount of effective labor. Finally, I think there's potential gains in matching processes to match labor supply and labor demand. And we've seen some at least preliminary evidence that AI can help improve match quality.

[05:47:32.41] - Adam S. Posen

Thank you.

[05:47:33.22]

Bo?

[05:47:37.05] - Bo Li

I think the answer Depends on different possible scenarios, depending on whether AI will be purely labor augmenting or AI will actually be more labor replacing. I think there are different scenarios. We probably don't know right now which scenario is gonna prevail. Of course, you can imagine in a labor-augmenting scenario, you will not worry about labor supply because labor supply will adjust to labor demand. So if we can create a lot of new kinds of jobs that people can continue to have, we can continue to enjoy full employment as we know today. Labor supply, as Michael just described, there are ways actually to to boost labor supply in terms of health benefits and other benefits and education. So that scenario to me is actually straightforward. But the more difficult scenario would be if AI is so revolutionary that it's replacing more labor than producing more jobs. And that one, we don't know what's the answer, right? Because in the future, if we have UBI, UBC, Maybe many people do not have to work, at least not work for income. If that's our future, then we don't have to worry about labor supply either because we don't have to worry about to find a job for everyone because there is a UBI and UBC.

[05:49:14.29]

So again, it's a matter of which scenario will be our future, and we don't know. We will have to, I guess, have to continue to observe and experiment.

[05:49:31.29] - Stefano Scarpetta

Maybe let me go to the 3 points building on what has been said. On the labor supply, I think there is some interesting emerging evidence about the fact that AI can make our labor market workplace more inclusive. Think about the potential benefit for allowing people with disability, physical disability and mental disability, to actually participate in the labor market. Again, We were talking about demographics. In a number of countries, advanced and emerging economies, we are aging rapidly. And of course, I think, again, we can think about workplaces which are more flexible, agile, and therefore will allow all the workers to adjust their working time, their working conditions in a way that will allow them to stay longer onto the labor market. So I think there are potentials out there, if exploited properly, to actually have a significant impact on the labor supply, especially when there is or more structural labor shortage, which is apparent in quite a few countries. I think the point about helping individuals to manage their own skills, that's in my view is one of the fundamental issues, because when I mentioned the little sort of the fact that many workers do not participate in training, it's not just because the employer do not give them access to training, it's actually because many of these workers do not want to participate or do not see the benefit of participating in training.

[05:50:42.22]

So to some extent, AI can be a tool to profiling what actually individuals know, what are their competence and skills, and what they might need to know in order to move into the new job or actually to keep their job that is evolving because again the task bundle is evolving. So I see a significant potential also for workers themselves to learn better what kind of training they might need in order to keep their job or to move into a new job. The last point I want to make, because this is the most difficult one in my view, is about how do we overcome the resistance that many managers, especially in small and medium-sized enterprise, have to use this technology that they don't fully master because this is a black box. You know what comes in, maybe, maybe not. You know what you get out. It makes sense, but you don't know how to explain it. And I think, I mean, what we have done at OECD, not because I want to publicize the work of the OECD, since 2019 we set up this OECD recommendation on trustworthy use of artificial intelligence, and we are working on what that means also for using AI in the workplace.

[05:51:43.08]

I think there is a lot that can be done without necessarily setting up very strict regulation on basically identifying the kind of measure, the kind of guidance, the kind of support that actually middle-level manager or companies, small and medium-sized companies, may have to be aware of in order to navigate this technology, which again, they see the potential maybe, but they're also concerned about the possible risk associated with that. Basically everyone all companies should be able to run a risk assessment. These are the potential benefits. These are the potential costs of using AI for specific tasks. And this is how we are going to minimize these risks. And some workers would still face the risks themselves. But I think there is a lot to be done there in what I call sort of creating some guardrail in which managers feel a little bit more comfortable about using this technology.

[05:52:31.02] - Adam S. Posen

Fantastic. This has been, I think, an outstanding panel in substantive and policy terms, reflecting the work of the teams of these 3 individuals, but also the people behind this conference. Just, we are going to take a short break, and in approximately 8 minutes, around 12:25 local—excuse me, 2:25 local time, we will have a further but much more research paper-based panel on AI diffusion moderated by Jacob Kirkegaard and with an opening setting up by Filiz Unsal of the OECD. With that, I want to thank Michael Kramer, Chief Economist of the World Bank, Stefano Scarpetta, Chief Economist of the OECD, and Bo Li, Deputy Managing Director of the IMF, with a special thank you to Bo both for making time in your last couple weeks in Washington and this role to engage with this conference and these topics, and for your past 5 years of service to the international economic community. Thank you very much.

[06:00:25.32] - Jacob Funk Kirkegaard

Indeed, honored to welcome you to the final panel of our 2-day conference, which I, somewhat biased, will say is probably the most exciting of all the panels. So, I like to think of it as we've had, you know, a 2-day conference, which is sort of the intellectual equivalent of a full 3-star Michelin full meal. We just had the high-level panel, which I suppose you could say was a little bit of a sugar-high dessert. And I am here to give you the double espresso shot that's going to make you sit up and pay attention even though it's at the end of the day. Because we're going to talk about what really matters. which is AI and diffusion. Everybody has an important AI agenda these days. Obviously, if you care about geopolitics, you read in the paper that 50% of all NVIDIA chips are sold to Singapore for use, really used for training Chinese open weight models. Obviously, if you care about overinvestment and bubbles like that, the Anthropic IPO prices in a few hours. It will start trading tomorrow. And obviously, if you care about whether your plane is on time or not, there is a cascading version of a cascading number of cyber or network reliance issues.

[06:01:56.07]

But as economists and economic policymakers, we obviously care about AI because it's a general-purpose technology with the capacity to potentially fundamentally reform, restructure, change workflows in all parts of the economy. And we care about who, how, how many, and how AI is diffused from frontier firms throughout the entire national population of firms and from sort of the core AI developing frontier economies, US and China, to the rest of the world. And that's exactly what we're going to talk about in this final panel. And we're going to start up with probably what excites certainly economists like myself about data geeks, with what economists like more than anything. Namely new data. Stefano sort of gave a preview of it. We're going to have the introduction or the publish, I believe the first time public, of a new OECD dataset for adoption and diffusion covering 60 countries, etc. We're going to have Filiz Unsal from the OECD do that just now as a stage setter. Then we're going to have 3 very exciting papers that I will— whose authors I will introduce in due time. And then we're going to have the discussant and the ability to all of you to answer questions at the end. But without further ado, Filiz, the floor is yours.

[06:03:47.32] - Filiz Unsal

Thank you. Thank you, Jacob. Good afternoon, everyone. It's a pleasure to be here. It has been really intense discussions and we learned a lot in this conference. One thing, I'm maybe the only one in this room that have worked on 3 out of 4 institutions, and I noticed that I am also the only one setting the stage. I guess that's what you get to do when you work 3 out of 4, what I thought. But joking aside, the reason is that when we were putting this data together, this

conference program together, we weren't entirely sure whether we are actually going to have a paper to present, and setting the stage seemed relatively cautious. So that's what we went for, but we were also very keen to share the data with you. But we managed it. So today I will be— I am very happy to present to you our work on Global Economic Survey on AI adoption, which we'll be publishing very soon and releasing the data. Can we put the slides, please? Okay, yeah.

[06:05:03.06]

Thank you. So as it has been mentioned earlier, we all expect AI to deliver several product— quite a bit of productivity benefits. And a lot of people have looked at this, and at the OECD also, we take this approach and looked at also across sectors and across countries. In our work, one of the key ingredients of what is the main driving of the benefits of productivity benefits of AI is actually AI-intensive AI adoption. So this is a little bit different than whether people are using AI or not. This is not a binary yes or no question, but it's more about integrating AI in the business processes and for people to work AI intensively. in their core tasks. Currently, this type of information is lacking across the globe. There are some surveys for a limited number of countries, but the systematic evidence on AI adoption is one area we thought we should fill the gap. And for this reason, we partnered with Google and Ipsos to run a survey across 62 countries. We focus on individuals, and this will be published, as I mentioned before, in a couple of weeks, and also will be published in our Economic Outlook chapter in mid-December.

[06:06:32.04]

And we are hoping for regular updates going forward. Okay, let me show you what we have found. As Stefano mentioned earlier, AI use is actually quite widespread, but what we find is that the intensive use is actually still very limited. And on the one hand, you see that the overall AI use for personal and work purposes is around 60-75% across countries. However, this is true across the board. However, this is, if you wish, a bit of a quick and a bit of a superficial number. When we focus on the intensive use for work, adoption appears to be rather shallow. In fact, the share of people who are using AI for work ranges about 20 to 40% across OECD economies, and the share for intensive daily work users is even lower at the average of 6%. So differences across countries, tend to reflect general level of economic development, as you can see in the chart. And OECD economies and high-income non-OECD countries lead in AI adoption. However, this also suggests potential AI divide between higher and lower income countries that we also discussed this morning. So as I mentioned before, at OECD, we looked at the productivity impact of AI, so one may ask the question, okay, so what your productivity numbers would imply with this type of shallow but widespread adoption numbers?

[06:08:18.11]

So here is the— this is from the literature. You see at the end with the blue bar, the OECD, what we estimated before across countries between 0.2 to 0.7%age increase in annual productivity growth. As Stefano mentioned earlier, this is actually quite a big number considering that, you know, in the last 2 decades, the productivity has been at most around 1%, let's say 1.5%, which is the United States. So we estimated earlier with some estimated adoption numbers about 0.2, 0.7. Now, when we incorporate our global AI adoption survey, survey, the recent evidence that I just presented to you, actually, that's pretty much very close to our model scenario, still around 0.2, 0.6. So the point I'm making here is that even though with those widespread but shallow numbers that I showed you, they imply quite substantial productivity gains. So now, let's look at digging a bit further and looking at to the different occupations, we see that adoption is not only shallow but also uneven across sectors. We compare the task— the share of tasks that could be performed with the help of AI as it has been identified in the literature and called exposure.

[06:09:43.53]

That's the blue line in this chart. With the share of workers that have actually adopted AI in the survey is shown orange line. The point here is that the orange and blue line actually quite follow each other, as you can see. However, we do observe that the adoption rates at work, that's the orange line, still remains below the exposure. So this gives you the idea that there is still some scope, further scope for adoption. And as you can see, the daily use at work, that's the green dots, are even much smaller than what the orange, orange that's use at work shows you. So at some point in the future, we may expect that AI use at work, that orange line, to be quite big, probably 100%. But what's important for us for the economic impact is more about the AI use at work. That's the green bar. So we have quite a bit of scope there to proceed. So still looking at the workers, we also see that almost half of the workers use AI, but only 20% of the workers are using it for core tasks. And what's interesting is that actually this 20% number is universal across education levels, or it doesn't change with the exposure level of the tasks.

[06:11:10.48]

So currently in all of the jobs, higher exposed, lower exposed, regardless of that, we see that only 20% of the workers are using AI in their core tasks. Of course, this goes in the direction of the literature that suggests that in many of these jobs, the different tasks are combined, right? And some of these jobs involve human interaction, that needs human interaction. And many of these things we cannot still do with AI where human judgment is required. So we then looked at what type of barriers we can talk about if people are not using AI at work as much as perhaps we think they do. Interestingly, One of the main reasons in high-income countries you see here, the main reasons are more about the data privacy concerns, bias and manipulation risks, or the

transparency concerns. However, when we look at the reasons why low-income countries are not that much incorporating using AI, then we see a different set of factors. problems, which is more about the lack of training, limited access, or low-quality AI output. But this lack of training and limited IT access also indicates that there may be some further infrastructure needs in developing countries that we may need to focus on, which was also considered today in the session.

[06:12:49.26]

And in fact, when we do some scatter plots, of course, this doesn't mean any causation, but it does seem that AI use for work is positively related to digital infrastructure, skills innovation, and general economic development. That's probably not very surprising, but still it confirms that for developing countries, maybe certain infrastructure needs may be critical. Now, also going now a bit more on the sentiment of AI, what type of barriers that we can identify in the case of AI use, what we see is that the share of people who report being mostly excited or excited about the possibilities as opposed to people that are concerned or mostly concerned about AI risks. You see that sentiment varies significantly across countries with a high share of concerned respondents in Anglo-Saxon G7 economies and a lower one in Asian and non-OECD economies. In the right panel, you will see that there's a stark concentration of concerns among never users. So people that actually never used or experienced AI, they tend to be much more concerned about the risks of AI. However, just next Next to it, you see the people that are using AI daily, and they are concerned— they are much more excited rather than concerned about AI.

[06:14:21.05]

So this shows that as people perhaps, you know, more experiment with AI, these user numbers can also increase going forward. Now, looking more closely at the negative sentiments among non-users, we find that actually across income level, across countries, the non-adopters in high-income countries tend to be much about the trust-related issues. That also goes back to the previous discussion that in high-income countries, people are more worried about the AI risks and transparency and bias. This is especially the case for never users. And so, you know, this suggests that addressing some of these concerns on AI could be an important policy direction to increase adoption, especially in high-income countries. So this is just fresh off the press. I will say we just get this data almost a month ago, and our wonderful team put together this for this event. So we are happy to— very happy to share with you, and there will be much more work from us coming on this. So please stay in touch. Thank you.

[06:15:45.48] - Jacob Funk Kirkegaard

Very good.

[06:15:46.15]

We now move to the presentation of the 3 related academic papers. And the first presenter is Professor John Graham from the Fuqua School of Business at Duke University, who will talk to us about AI productivity and the workforce and what corporate executives think about that issue. And as John comes up, the eagle-eyed among you will have probably noticed that John's name is not on the table here among the presenters, those who will be able to ask questions, because John has to leave a little early. But that just means you have to square away all potential questions while you're here during your 25 minutes.

[06:16:25.24]

Thank you.

[06:16:33.07] - John Graham

Okay, well, thank you for that introduction, and again, my apologies to the organizers and our discussant for having a schedule change, and I had to take off after pretty much the second paper. So I'm going to talk about a CFO survey. We've been doing it, we've been surveying chief financial officers at Duke for over 30 years now. quarterly, so over 120. In the last 6 years, we've been doing the survey jointly with the Federal Reserve Banks of Atlanta and Richmond. So you can see all the list of authors there. Those are my colleagues who work on the CFO Survey in Atlanta and Richmond Feds. I should say that what I'll say today are my own comments and don't in any way reflect Federal Reserve policy. On this particular project, we also had some CFOs from NASDAQ and a membership organization, FEI, join. join us. Okay, so this is going to be a survey on your typical firm's use of AI. All right, so about a year ago, year-end 2025, we asked about the usage for 2025 realizations as well as their expectations for 2026. We had 734 executives respond. I should say this will be a disappointment for this audience, but it's just for the United States, 2025 realizations, 2026 expectations.

[06:17:51.20]

And think of this as the usage of AI at your typical company. So we're not talking about data centers here. We're talking about just the vast majority of companies, how they're using AI. And of course, this is just in the early stage of adoption of a new technology here. We very purposely kept the survey very short because who can really predict very much what's going to happen 2 and 3 years from now anyhow? Having said that, we obviously need to update this. At the very bottom, if you go out and find the paper, there's a data dashboard. You can play around with the

de-identified data if you want to slice it and dice it in some way other than what I present. All right, so here's kind of a rough outline. I'm going to walk through these things and interject some slides and such, and we'll start by talking about AI adoption and spending. And as our last speaker just said, we're going to find widespread widespread but shallow usage of AI in these US firms. And we'll talk a little bit about the type of spending also. Okay, so if you can't read that, the blue is 2025 usage, the brown is 2026 expectations, and we have 4 industry groupings I'll refer to: manufacturing on the left, low-skilled services, high-skilled services, and then on the far right we broke the finance sector out of high-skilled services.

[06:19:09.14]

And basically in a nutshell, increased use of AI from 2025 into 2026, close to 90% or so companies say they're using AI or expect to use AI in 2026. All right, this one you definitely can't read, but we broke out the dollar spending of AI by these various companies. I should have mentioned we have a wide range of companies from small, medium, large, public, private, and across various sectors here. So on the left, I'm not going to talk about is the total dollar spending. I'll highlight the right instead. This is per capita spending of AI on AI. And we can see I'm breaking it out by large companies versus small. Large company is 500 or more employees. Small is less than 500. And just looking at the numbers on the top there over on the right, small companies about $350 per employee, large companies smaller than that, right? So it's just not a large dollar amount. It's a lot of subscription-based spending at your typical company. All right. And as other people have hinted at, I think this is one reason why we're going to see the AI diffuse pretty quickly because you don't have this huge upfront upfront expense.

[06:20:24.29]

You can just jump right in, like, just like you did when you subscribed to Netflix in terms of using AI. All right, now what are they spending on? All right, on the far left on this graph is hardware, not so much at your typical company. On the far right is operational spending, that is subscriptions. The small companies are in brown, the large companies are in blue. For small companies, up to 2/3 of their spending is on these subscriptions, very little on hardware. The middle is kind of extra spending you do on internal development with this external product perhaps. But most of this is subscription spending. All right, not much hardware, and that theme will come back up again as we go. I'm going to skip this one.

[06:21:08.28]

This is just—

[06:21:09.07]

go and find the paper if you want to see why some companies say they haven't fully embraced the use of AI yet. All right, so here we had kind of an ex ante question. Why did your company adopt AI? And at the top of the graph there, the things that companies were after were increased production efficiency, increasing productivity, and also introducing new products or innovating. That's kind of why they said they wanted to get into using AI. On the bottom, Not so much about reducing costs or reducing employment, at least at this stage, this initial stage of where we are in AI usage. All right. Now, let's talk more about what are the effects of actually using AI. That previous graph was ex ante. And it'll overlap significantly. It's about efficiency and productivity, not about cost reduction. And I'll talk about what types of tasks AI seems to be helpful for at companies today.

[06:22:12.53]

All right, so what we've got here, I guess there's a lot on there again and another graph with a lot on there for you, but this is the ranking of things that what is AI affecting at your companies. And number 2 there with that bluish circle there and the arrow, output per worker, labor productivity. This is where companies are saying it's one of the chief things they're observing. Now here, probably have to take the exact number with a grain of salt because it's bigger than what We've been hearing others quote here today, but that circle says about 3% increase in productivity because of the use of AI. And it's broken out by— and those are for the 2026 expectations for small firms in brown, large firms in blue. So that's, again, pretty large number. There's some other things on there I'm not going to get into, but down on the bottom here on the left, also circled, is what about an increase in revenue? Here we've got— so the increase in revenue or sales due to the use of AI is more in the 1.8% range. So I'll come back to this discrepancy between supposed productivity increase versus revenue increase.

[06:23:20.29]

I'll come back to that a little bit later. Over on the right, what are we seeing? Just like the ex ante expectation, a bunch of zeros for costs, whether it's labor costs or non-labor costs. cost. Nothing much happening there on cost savings. And on the very bottom is employment. And it's basically a bunch of zeros too. The very bottom has a little blue bar pointing to negative territory. Large firms, there's a hint that they may be reducing employment in 2026 because of the use of AI. Now, I'll come back to that also. You're actually going to see this graphic again in just a minute. All right, then we listed out a whole bunch of tasks that companies might use AI to perform. And what does it actually help you with and what does it not? I'm not going to read all the things that are on there, but over here in the bold, basically what AI is currently helping companies with is routine pattern recognition, text generation, fairly simple tasks. Continuing on

this same graphic but on the next page, What is AI not helping with at this stage? High-stakes decisions, case-specific judgment. And if you can't, I know you probably can't read most of this in the back, so that's why I'm reading a lot of it out loud.

[06:24:34.59]

But at the very bottom, AI is not helping making merger and acquisition decisions or tax strategy or investor relations or, or management type, partnership management type issues at this stage. All right, on to productivity. So here we're going to see, yes, AI is said to be increasing productivity. We'll dig in a little bit on that, and we'll introduce what we're calling a productivity paradox. We didn't invent the term, but a new example of a productivity paradox. So there's the same graph from before. When I talk about reported increase in productivity, that'll be that output per worker, about 3%, that number I quoted earlier. But we're also going to measure productivity by some kind of measure of output over labor. Okay, the Y there, the output, we don't have a number of units variable on our survey, so we're going to use the revenue number as our measure of output. And then we have a measure of labor directly. So I'm going to use that. We're going to call that the calculated or the implied change in productivity. All right, so let's talk about the reported productivity and compare it to the implied calculated productivity. So what I've got here, I'm back.

[06:25:53.30]

I'm going to have these 4 industry groupings again. What you're seeing right now is just the manufacturing and construction. So in the, in the solid blue number there, that is the 2025 reported increase in productivity in manufacturing, and then the solid brown is the expected in 2026. Then these hash numbers here are the calculated increase in productivity in 2025, and then again in 2026. The first takeaway from this is just productivity effects are expected to increase. All right, not surprising. If I put on the other categories here, low-skill service, high-skill service, and finance, we see that same statement is true across the board. Productivity increases from the use of AI. Second thing we see is the increases are expected more in the high-skill services and the finance in particular industries, all right? Not in low-skill industries so much or manufacturing in particular. The third takeaway from this slide is, again, the productivity paradox phrase I used before. If you compare the solid, that's the reported increase in productivity, to the hashed, that's the calculated change in productivity that we calculate, the reported is always greater than what we calculate. In other words, going back to Solow's quote from the 1990s with the computer revolution, Productivity is everywhere, but in the data, we're seeing that here too, where the reported effects are greater than kind of the measured effects, if you will.

[06:27:34.22]

The 4th point that you can come to here is this— the red dashed lines tie into this. And I don't want to overinterpret this, but it is kind of remarkable that the number we calculate for 2026 lines up with the number they're reporting for 2025 almost exactly in all 4 of these sectors. So I wouldn't take it literally that there's going to be only a 1-year delay, but it suggests this thing about delay in when it's going to show up. Normally, we'd say in the bottom line numbers, but in this case, it's revenue. It's actually the top line number even. There's going to be a delay. So are productivity effects coming? We think so, but we're not exactly sure when. It should be at least a year. All right.

[06:28:18.23]

Yeah, I know you can't read this one, so I'll give you the summary. So we took a Cobb-Douglas production function, and we backed out how much of the increase in productivity is coming from capital deepening, from the expenditures on physical capital. And that's shown in the blue. So I know you can't read much from where you're sitting, but it's very small. So for this AI revolution, at least for your typical firm, there's not a big effect of capital deepening, if you will, of the increase. And the rest is just TFP, the residual that we would like to try to explain. So first point here is it's not physical capital. The second point, and it's really— you can't really make it out from up there, but we could tie the increases in productivity to the things that companies say they're doing with with AI. And to make a long story short, the increases in innovation and satisfying customer demand, including with new product or better product, that's where the connection is between productivity increases and where it's coming from. So innovation and meeting consumer demand. And of course, and I'll come back to this again a little bit later, but This is different than other products, than other tech revolutions, right?

[06:29:39.28]

For the computer revolution, you had to buy the hardware to take advantage of it. It doesn't seem like you need to, at least at your typical company. Why? Because somebody else is doing a lot of the hardware investing. It's happening somewhere in the economy, but not at your typical company. All right, now we'll get into workforce, and this ties into what we've heard people say here. Talk about the magnitude of the effect and what types of jobs might be affected now. So that little top bullet up there, we had an open-ended question on the survey. Which jobs at your— or roles and responsibilities at your company will be replaced by AI? Blank line. Now, if you're doing a survey and you have a blank line, rarely do enough people write the same thing that it adds up to much. So open-ended questions usually get pretty low responses, but fully half of the people wrote in no effect on employment from AI on this, which is very large for an open-ended question. Now, on the other hand, the other half didn't write that. So for the

other half, we took our sample and we weighted by sector and by size of firm, and we tried to get up to represent the US economy.

[06:30:52.15]

And if you add up the types of job loss they're talking about, about 500,000 jobs overall. Now, that's 0.4% of the US workforce. I used the adjective modest, maybe I should have written moderate, because that's about 40,000 jobs a month. That would be pretty nice to add to payrolls right now. But anyway, that's the ballpark of what we're seeing in terms of job loss. What types of firms are laying people off? This graph, the brown lines coming down, Those are for the large firms, 500 or more employees, mostly in the high-skilled services and finance sector. You can't see it because it's pretty small, but there's a little bit of positive blue. Small companies might actually like to hire more employees because of AI. And there are a lot of small companies, right? So that, that might add up to something. We'll dig in a little deeper on this. So here we've got what types of jobs are being affected. All right, on the left we've got routine clerical work, and the x-axis is firm size. Okay, so large firms over there on the right of this left graph, they're the ones that want to reduce routine clerical work.

[06:32:03.19]

And over here on the left, small firms kind of hanging in there on routine clerical. On the right graph, we see small firms would actually like to hire technical workers maybe to take a little bit of advantage of the AI, better advantage of AI, and large firms holding steady. If you put it all together, job loss is coming from large firms in the high-skilled services industries, particularly in routine clerical work. I think this lines up with what we've heard other people saying over the last couple of days, and maybe a little bit of a bright spot for at least small companies liking to hire technical workers. Now, if you went out there to the paper and played around with the data that's online, we've made some word clouds and things like that. This is, again, from the blank line. What did people write in on what types of jobs and roles and responsibilities would be affected? There's several words here. I'm going to highlight marketing and accounting because we'll see them again on the next slide. These are jobs that could be replaced. These are the job— a separate question, which jobs will be enhanced by AI, blank line.

[06:33:16.25]

And here you see marketing, accounting, finance on the low-skill side are the jobs that are enhanced by AI to the extent, if we combine it with the previous graph, some AI is so good at doing some tasks, we can get rid of some of those employees for these particular types of sectors. All right, so we took what people wrote in in that blank line again, and we're going to

take all those words and match them to the BLS jobs, the 343, I think, job categories that BLS lists. Got a little AI humor for you. We thought this is a great task for AI. We're going to take these words and match it. AI was terrible at this task. So that was a joke. Maybe not. So anyway, we had to hire people to actually do that matching. All right. And I know you can't read it from where you're sitting, but the small print is these exact BLS jobs. And oh, we created a ratio of the replacement divided by enhancement. So anything greater than 1 means there's more replacement being mentioned than enhancement. And so up at the top here, office and administrative support, bookkeeping, accounting, clerk jobs, This is literally what people are writing in.

[06:34:31.03]

And what's different about this, just to emphasize something in case it got lost in the shuffle, I mean, this is from inside the business, what they're saying they're actually doing or planning to do, rather than much of the work we have is about what might be exposed to AI. This is what the job's actually being affected. Down here on the bottom, engineers are at the other end of the spectrum, and management actually. where the enhancement is much better than the replacement.

[06:34:59.56]

All right. Here's another breakout you can find in the paper if you want to see it by sector rather than by job title. Okay.

[06:35:08.20]

So let me just summarize a little bit. With the AI wave, we try to compare it to other technology waves. A couple of things that are similar in aggregate, Because large companies have a lot of employees, they spend a lot. They are sort of where the frontier is right now on the rolling out that's similar from other tech waves. And the productivity paradox, we've seen that before in other tech waves also. Here are some differences. AI spending is broad but shallow, and the dollar spending is actually surprisingly low, at least at this stage. AI adoption does not require the capital deepening, as I mentioned before. That's different than other ways where you had to buy a computer or something. The third one down here on the bottom, it's tied to innovation and demand, not to cost cutting. That's different. And finally, employment. The service industries are being affected most now, not like manufacturing with automation. There were a lot of results here, a little bit of a fire hose here probably. So if you go out and find the paper, Table 1, Splits out a lot of the results by small firms, large firms, and where you can find the evidence should you be interested.

[06:36:23.27]

Okay, so I'm not going to read the first 3. Those are just kind of summary results, probably summarized them enough. Future research.

[06:36:29.16]

Okay, look, we did a survey of 2025 and 2026. So much has changed in the last year. Obviously, we need to repeat the survey as well as other people doing a lot of updated research. And we are going to repeat it in the first quarter of 2027. Second little bullet on the bottom is, well, just AI is changing so fast. That's another reason to repeat it. The third one, we're looking at a typical company use of AI. Okay. Obviously there's this other stuff with data centers and all going on. That other research, you know, you really have to complement what we're doing is only part of the equation. You need to add the other stuff in to get a good macro sense of the effects on the economy. And the last bullet point, and this will be my last comment, is that there is this hint that innovation and new products are where the enhancement in productivity is coming from. Well, that means there are new companies being created that are actually hiring people. We couldn't survey them because they didn't exist a year ago type of thing, right? Or they aren't in our panel yet.

[06:37:31.16]

And so there is this other piece that hopefully others are doing good work. on that innovation piece that we're unable to do. With that, I'll leave you. Thank you.

[06:37:48.25] - Jacob Funk Kirkegaard

Thank you very much, John. The next presenter is Hannah Wei from the IMF, who's going to present her paper on emerging trends in global technology diffusion and whether AI differs from other new technologies. So is this time different?

[06:38:06.59] - Hannah Wei

Good afternoon, everyone. Thank you for having me. My name is Hannah Wei, and I'm an economist from the IMF. The paper I'm presenting today is joint work with 3 economists from the World Bank— Paolo Bastos, Catherine Stapleton, and Daria Taglioni. So the question we ask in the title is simple. Artificial intelligence is spreading fast. But is it spreading the same way earlier technologies did? And our short answer is no. I'll use the session today to show you why. So let me start with why this matters. AI is widely seen as the defining technology of our time, yet we

know surprisingly little about how it actually spreads. For earlier frontier technologies, we actually know quite a lot. 2 well-documented channels, for example, are proximity and firm networks. For example, when a firm is located near where a technology was first created, was usually the first to adopt that technology. And likewise, for firms connected to the innovators through the value chain, were often the next in line. And policy so far has also built on these exact channels for good reasons. Place-based policies such as the innovation clusters or regional technology hubs tend to bring firms together closer close to where technology was created.

[06:39:22.56]

And then joining the global value chain also has been a proven route for countries to catch up technologically. And meanwhile, export control policies in some countries were also designed to limit the frontier technology out of their competitors' reach. But what if AI spreads differently? What if its open-source models make proximity or trade linkages matter less for AI's diffusion, then some of these earlier policy instincts may need rethinking. So in this paper, we ask, how does AI diffuse, and is it any different from other technologies? So here's what we do. We compare AI with 28 other frontier technologies, the full set identified by Blue Metal. And we measure diffusion through job postings data, which was also used by some of the earlier sessions yesterday. And our dataset is very granular. We have nearly 600 million job postings across 17 economies from 2014 to 2022, with more than 100,000 firms also matched to the supply chain linkages. And we have 4 main findings. Relative to other technologies, AI-related job postings is first of all more dispersed across countries and regions from the very emergence of AI. And second, AI is also more tied to income. So richer countries tend to shift their technology-related postings more towards AI than towards other technologies.

[06:40:53.10]

And third, AI is also shaped more by distance within countries, but not between countries. And last, we also see that AI is less transmitted through firm networks. especially in the early stage comparing to other frontier technologies. So AI seems to actually spread more through the local labor market than through global firm networks comparing to other frontier technologies. If you remember one line from this talk, this is the one. And then let me turn quickly to data. Sorry, I think I clicked something wrong. So for data, job postings data has been an increasingly popular barometer to measure how technologies spread over time. And several studies previously find that the job postings data from Lightcast represent the underlying economies reasonably well. For the US and the UK, they cover around 85 to 90% of all vacancies in their economy. economy. And for the 17 advanced economies that we cover in the paper, they together account for around 60% of world total GDP. And we focus on the technologies identified by

Bloom et al., who set out the 29 most disruptive frontier technologies of the 21st century using text analysis from patent data and earning calls. So a job posting will signal adoption if it mentions one of these technology keywords from Bloom et al.

[06:42:27.05]

And so what counts as AI in our dataset? Bloom et al. define it more broadly. So any keywords related to machine learning, natural language processing, or neural networks would be considered AI. And most importantly, the AI that we cover in our dataset is more pre-generative AI because our data ends in 2022 before ChatGPT. So what we cover, more precisely speaking, is the early prediction and adoption models instead of the later, more generative large language models. And then to explore the role of firm linkages in AI diffusion, we also use data from FactSet, which allows us to link the job postings data from firms with their buyers, suppliers, partners, and affiliates. And then let's take a look at what the data shows. So here on this chart, I show the share of all job postings mentioning AI-related keywords against the average share of job postings mentioning other frontier technologies, with both set to 100 in 2014. And we see that from 2014 to 2022, the share of AI job postings went up by more than fivefold compared to other technologies, which only grow by 1.6 times. And the growth rate for AI-related job postings was around 20% a year, making it the second fastest growing technology in our sample.

[06:43:53.24]

And then we also compare the growth of AI-related job postings across our countries. So the hollow dots here on the graph show the share of AI-related postings in 2016 against the solid dots, which is the share in 2022. And here we see that the rise in AI is seen in every single country, but their levels differ. So the top leader on the chart, which is Singapore, had more than 5% of their AI-related job postings by 2022, followed by the United States. Whereas on the other end, we have countries like Italy, New Zealand, and Belgium with less than half a percent of their AI-related posting in 2022. So there's also the distinction between all these advanced economies over time in their AI adoption, and this also motivates us to explore what really matters for their diffusion. And this leads us to our next 4 stylized facts about diffusion. So starting from the first fact, we, we find that AI demand measured through job postings data is more evenly spread than the demand for other technologies. To explore how widely spread AI is, we use the standard measure, which is the coefficient of variation here. And we show it for AI in particular against the average for other technologies.

[06:45:15.07]

So the lower the value of coefficient variation, the more widespread the technology is. And we see here for both panels across the 17 countries against— at a more granular level across the 156 regions, AI adoption shown in the blue line has been sitting well below the orange line, which is the coefficient variation for other technologies. So in other words, AI has been more widely spread from the very beginning comparing to other technologies. And another observation from this chart is that at the country level, AI diffusion has been more stable with little variation over time, whereas the diffusion across the regions has been— the coefficient of variation has been decreasing stably over time, suggesting that much of the earlier AI adoption that we see from the job postings might have been more happening within countries instead of between them. And then we— This is not my slide. What did I do? Okay, almost there. Right, so coming back to our second slide, stylized facts. So for here, we explore the role of income in technology diffusion. And on the vertical axis, we show the share of AI-related job postings relative to other technology-related job postings. And on the horizontal axis, we have GDP per capita.

[06:47:09.50]

And so the question we want to explore here is not whether richer countries tend to hire for more AI, they do, but rather whether richer countries tend to hire more for AI comparing to other technologies. And so the chart here shows a positive correlation of 0.44 suggesting that income indeed matters for AI than for other technologies. But one potential concern for this correlation is that there are quite a few countries like Singapore, Norway, or Luxembourg on the chart chart that have a very similar economic structure in terms of the fact that their share of finance or professional jobs are very common, which just tend to hire more for AI. So we also, as a robustness check, control for the share of finance and professional services jobs in these economies. And our correlation is still positive at around 0.4. And then we also explore what's the role of distance in AI diffusion. So before going to the details, there are 2 relevant definitions. The first one is the pioneers, which we take from Bloom et al., who define them as the regions that together hold half of a technology's citation-weighted patents in the 10 years before it emerged.

[06:48:24.30]

So in other words, these are the places where the most influential invention behind a technology took place before the technology started spreading out to the rest of the economy. And for AI, this is mainly the San Francisco Bay Area. And then the distance measure is just pretty straightforward. It is simply the kilometers from a country, state, or firm to the nearest pioneer region. And then going to our specification here, we have the outcome variable, which is the normalized share of job postings mentioning each technology. And we regress it on the distance

measure and also the income measure. Uh, we also interact them with an AI dummy to particularly see how AI differs from the rest of the technologies with, uh, technology and year fixed effects. So our coefficients of interest here from the equation are beta 2 and gamma 2 because they are the ones that show us how AI differs from other technologies. And then here, uh, we have the results for the AI interaction terms. Um, so all the coefficients here tells, us how much more or less distance and income matter for AI diffusion comparing with other frontier technologies. And let's start from the country panel.

[06:49:35.58]

So across the 17 countries, the coefficient for distance is positive at 0.12, suggesting that across all the countries, distance seems to hold AI back less than it does for other technologies. And income is positive and significant, which reinforces our our earlier slides on the positive role of income. And then moving to the US states, we see that the distance picture now flips with a negative coefficient of -0.09, whereas the income stays positive as before. And then going to the firm-level regression, we're having again a negative coefficient for distance, and income here at a firm level seems to matter slightly less. So taken together at the country level, it seems that distance does not seem to be the main barrier for technology— for AI diffusion. But within countries it does, and also within firms. This also fits the idea that AI tends to be more closely linked to the local labor market. So within a country being far away from the technological hub does put you at a disadvantage. But across countries, what matters more is the income and the associated skills and talent with that income level, which seems to be a bigger priority than distance itself.

[06:50:57.32]

So let's now look at our final diffusion channel, which is through the multinational supply chains. There are 2 relevant set of firms which matter for our analysis later on. So the first set is called the pioneer firms, which we take from Bastos et al., who identified about 5,000 pioneer firms in a technology's pioneer regions and also their lead industries before the technology emerged. And our other set of firms are the network-linked firms, which are simply the preexisting buyers, suppliers, innovation partners, and affiliates of a pioneer firm, which we gathered from FactSet. So here on the chart, we simply show the share of job postings mentioning these technologies for the pioneer firms against the other non-pioneer firms. And we see that for both AI and the average other technologies, the pioneer firms have a very big advantage in technology uptake. And in particular for AI, the share of the job postings for the pioneer firms went from just below 2% all the way to nearly 8% over time. And then here we show a similar chart but for the network-linked firms against the unlinked firms. So we see that from both panels, again, the

network-linked firms seems to have a strong edge comparing to the unlinked firms, especially for the other technologies.

[06:52:24.15]

And then on the left panel for AI, we see that early on in 2014, the links to pioneers does not seem to matter as much. But over time, we see that the gap between the network-linked firms and the unlinked firms started to widen. So in short, links to the pioneers seem to carry AI only very weakly in the beginning, but they did more later on. So to test these ideas more formally, we have the following specification. The outcome variable is, again, the share of the normalized share of job postings over time for each technology firm and year. And then we include a pioneer dummy reflecting whether a firm is a pioneer firm or not. And then we also have linkage dummies reflecting whether the firm is a pioneer firm's buyer, supplier, or innovation partner. And we interact all of these variables with an AI dummy, which could tell us how AI differs from the rest of the technologies. And we have firm technology and year fixed effects. So to also see whether the effects that we have for AI changes over time, we also interact each term with the number of years since the technology had emerged.

[06:53:42.13]

So in short, what we ask here is for a firm to hire more technology, does it benefit for the firm to be a pioneer or does it benefit for the firm to be linked to a pioneer through the networks and whether that effect is stronger or weaker for AI. So here we highlight the 2 findings in the table from the 2 red box. And starting from the first one, we see that AI pioneers actually start far ahead of other technologies pioneers. This is seen by the interaction term between pioneer and AI as we're seeing very significant and positive coefficients across all columns. And then moving to the networks channel, we see that the links to AI pioneers actually adds much less than for other technologies. So in short, once again, the networks effects on AI seems to only carry AI very weakly. And next, we also see whether the effects remain stable over time. So starting from the 3rd red box in the table, we see that AI pioneers lead shrinks each year over time as we now interact the pioneer with the AI over the years since the technology's emergence. And then additionally, being linked to AI pioneers for firm networks channel, especially for buyer and innovation partner here, seem to grow more valuable over time.

[06:55:10.58]

as seen by their positive and significant coefficients across all specifications. And so to put together, the story for AI seems to be more of a convergence story. Their pioneers lose their exceptional lead over time, and then we have the network partners catching up more gradually

over time. So having seen these results so far, we have a central question that's still remains. While we see that firm networks do matter for AI, they seem to transmit less AI than for other technologies. And while we don't answer these directly in the paper yet, there are earlier studies that provide 4 explanations that fit our evidence so far. And so first explanation is that AI is more embodied in workers. It cannot be is physically shipped along a value chain. And second of all, the talent for AI is usually pulled more locally from the labor market because AI workers tend to cluster in just a few labor markets and multinational firms usually also keep their AI teams in core R&D hubs instead of spreading them more widely across all of their affiliates. And third, AI is more of skill instead of a buyer or supplier along the chain.

[06:56:32.17]

Much of AI's toolkit today is still open source, and so that means the technology is free to travel. But what is scarce is people who can use this. And lastly, much of our findings so far still surrounds AI development in the early days, as we only cover 2014 to 2022. So as AI matures and AI skills become more standard, we may actually also begin to see network effects being stronger over time. So let me conclude by going back to— this is not going back, is it? But I can also just quickly talk about my conclusions, although this will matter for the next presenter.

[06:57:26.17]

All right. So to just quickly conclude— all right, I'll just leave it here for the quantitative model for the next person. So to just quickly conclude, going back to our earlier question whether— all right, maybe I'll just leave it there. whether AI diffusion is different from the rest of the other technologies. And yes, it does diffuse differently. It is more widely spread across countries and regions. And income also matters more for AI. And in terms of firm networks, we are seeing much stronger effects for the firm networks channel for other technologies than for AI. So going back to what this means for policy? Well, we have an earlier discussion about how joining the value chains matter for technology diffusion. But for AI, joining the value chain itself is not enough because we have seen that for the pioneers, they only have a very weak effect for AI diffusion early on. So there are also other studies show that policies that target skills and talent and the labor market would also complement AI diffusion. And then for future research, it would also be interesting to see how the patterns still hold or not for the years after 2022, as our data only covers the pre-generative AI. And that's all. Thank you.

[06:59:05.17] - Jacob Funk Kirkegaard

Thank you.

[06:59:05.40]

And despite the, I guess, involuntary sneak peek of the last paper, it is nonetheless my pleasure to welcome to the podium Daniel Lashkarif, who joins us from the London School of Economics to give his paper. I think we should see if we can find the— you probably are best versed to find the first one. No, Daniel is going to give his paper on AI and Scale: Quantitative Task-Based Theory of Automation. Thank you.

[06:59:35.13] - Danial Lashkari

Well, thank you so much to the organizers for including our paper, to all of you brave souls staying until the very last presentation. I have to say that this is sort of an agenda, a research agenda we have around thinking about AI automation from a quantitative perspective. We have this very recent set of results. Thankfully, the organizers basically allowed us to— so what I'm going to show you today is a really recent set of results that is not exactly what you have on the program. Our discussant was kindly open to discuss this, even though this was a last-minute change, to be able to allow us to share where we stand in terms of our thinking. The paper is joint with Omid Mahzian, Matthias Mertens, and Neil Thompson. All 3 are based at the Future Tech Lab at MIT, a lab to which I'm also affiliated, as well as Christina Chu, who is at the University of Miami. Now, what I want to do is to start by contrasting 2 different views about the future path of AI automation. On the one hand, what I would refer to as the capabilities views, exemplified by this quote from someone that some of you might recognized now as being an infamous character in this world.

[07:00:53.01]

It's typically held by AI enthusiasts. The idea is that we know capabilities of AI are evolving and expanding over time, so it has to be that automation follows where capabilities go. Meters released this really by now well-known report last year that showed that the length of tasks that AI can successfully accomplish has been doubling every 7 months. So that's where we're headed, no questions anymore. But then there is what I call the cost view, exemplified by this other gentleman who's out there in the world trying to deploy these tools. And this idea here is that it's not necessarily the case that as soon as AI can do something, we would want to to actually do it.

[07:01:40.10]

Why?

[07:01:40.39]

Because it might still be too inexpensive and inefficient. What we want to do in this paper is to think about how these 2 forces evolve together. Over time, capabilities evolve. Frontier Labs are training these frontier models that are released, as Sonny mentioned yesterday. There was one of them that was released yesterday. Every one of that is released, it's larger in scale, it brings about new capabilities. But that comes with a price tag. Some of you might have already been grappling with the choice between Fable and Opus. And you know that if you use Fable for too many tasks, you're going to run out of your allotment. And then that's a problem. But then on the side of cost, there are also similar dynamics.

[07:02:29.23]

Why?

[07:02:29.38]

Because Innovators both on the side of these frontier labs designing the models and their architecture, but also on the side of hardware, those who are designing chips that then feed into compute, they're basically helping reduce the costs of using these machines, either it is through the cost of compute itself or through the productivity of these tools. On top of that, firms are presumably facing some frictions, regulatory, organizational, the likes of which we saw earlier in the earlier presentation. So, over time, as we are adjusting to this new wave of this new technology, these frictions are also likely to fall. So, what we want to do in this paper is to think about the evolution, the dynamics of these 2 sets of forces together in the context of the leading framework we have theoretically in economics for thinking about automation, which which is a task-based framework. What we more specifically do is that we're going to build a quantitative model. Its main novelty is that we assume that firms are faced by a menu of choices. These are menus of automation technologies as opposed to just one. And of course, in our setting, each vintage, each item in the menu is going to be an LLM model.

[07:03:49.23]

Now, In choosing which of these, if any, firms deploy for automation, they're trading off capabilities and costs. Do I go for the Fables of the world that are very capable but too expensive, or expensive, or do I choose a smaller model that has lower capabilities but it's cheaper for me? And of course, they also do this subject to additional frictions. The way we capture these additional frictions is in terms of some overhead labor costs that still they have to use in the deployment process. Then what we do is that we take this core model, we embed it in a multi-occupation general equilibrium environment, essentially a model of the macroeconomy, which then allows us to look at what happens when we feed a model like this with the trends

that we know has been at least historically happening in AI progress. in a decline in the price of compute, and then in potentially decline in frictions. And I'll tell you how we go about disciplining the last part. Now, OK, so this sounds like a very rich model. Of course, you're going to ask me, how do I ever have any hope of disciplining this model?

[07:04:59.29]

Well, that's where this new dataset that this lab, MIT Future Tech, has been collecting comes into the picture. So, what is this? What we do is that we take thousands and thousands of own-it tasks. And we create samples of those tasks, production of those tasks by lots of different models of different size and complexity. And then we give those productions to workers in the same occupation to rate them. What this allows us to do is to essentially learn a mapping between the inputs to the model in terms of, let's say, the number of model parameters, the amount of data used for its training, on the one hand, and its capability, its performance across a wide range of tasks in the US economy. And so this becomes the core of our model. Using this, we can actually fully characterize the whole capability versus cost side of this trade-off. And then we calibrate our full general equilibrium model. We need some additional moments from the data. importantly, of course, we do need here data on adoption. So the ones that we're going to focus on, the data by Alexander Bick that we talked about, but now that I'm inspired by seeing the earlier presentation, perhaps we have even further data.

[07:06:14.24]

It's important for us that this data has to be on the intensive margin, of course, right? We want to actually measure the productivity gains and the size of automation. And this is what allows us to back out the amount of frictions that is in the data because We— the model tells us what the capabilities and costs look like. And so, the gap between that and adoption, we explain through the frictions. Then, we're going to use this model to do some quantitative exercises. The core of it, as I mentioned, is exactly this idea that we feed this model with these trends. There are a whole host of results we have. I'm going to just highlight a few of them. So one of the key findings, first findings, is that the choice of the model is key in explaining the variations in automation, both in the aggregate and also across tasks. And another fairly stark result we have is that between the drivers that I mentioned, the one that wins the day is the algorithmic progress. It's not the scaling of the model. It is the algorithmic progress that if it continues at the rate that it has, we are bound to have to project really, really fast eras of growth ahead of us.

[07:07:27.26]

So, I'm going to hopefully come back to these points as I talk about the results later, but let's jump right ahead into the model. So, what is the framework? A lot of you might already be familiar with this core of the model. What I'm going to lay out now is going to be the production function for for one own-it task. Essentially, that's what you want to think about, right? And then we're going to embed it in the broader— in the larger model. So what is this like? I'm going to assume that that own-it task that you typically think about has a bunch of different varieties of it. Let's say the task is, I'm a professor of economics. I have to write some problem set, right? There are many different problem sets I can write. Each of them is new. And we're going to follow this idea that there are weak links or there are strong complementarities between these different tasks for my production. So lambda is less than 1. Now, if I were to use labor, I'm going to assume that labor has this task length— faces that task length tau, meaning that the hours of labor, direct labor LL, has to be divided by this to give me the units of production of the output.

[07:08:33.24]

But I also have the option to adopt this machine M, and I'm going to Be explicit about the machine M because, as I've already mentioned, you're going to be faced with a menu of these different machines. So but for now, let's start with one machine. Now, you're given this machine. For that variety— task variety new, this machine has a quality of production Q. So this is the capability side. For you to be able to accept its production, it has to meet some standard Q*. If it doesn't, then AI is not capable. At least, that machine is incapable in that variety.

[07:09:05.26]

If it is capable, then what you're going to do is that you're going to think about now actually how much production you're going to do with it. So, the way it works is that it's a perfect substitute with the direct labor usage, but it requires potentially not only just the compute input, which is K divided by gamma— gamma is the per unit of output requirement, compute requirement— and that's the cost side, compute cost side. In addition, it also requires you to potentially use some labor time. Your workers have to work on checking the results, figuring out how it fits with the rest of your production process. There's some essentially tax omega on labor that you still have to pay. So, the labor savings for you is going to be 1 minus omega times tau as opposed to being the overall tau. That's the standard. In the model, we're going to add this part to be able to capture what we see in the data. OK. So, this is the core of the model. Now— oops, this— Does it look good? OK. So now, let's parameterize this, or let's characterize this a bit more to be able to actually make it a quantitative framework.

[07:10:08.46]

OK, on the capability side, I'm going to assume that there is a distribution h of the qualities that you're going to get from that machine across all the varieties. So to the left— oh, I see. That's what's happening. So to the left, for now, ignore different m's. Just think about that curve. What that curve tells you is that you're going to only, for that machine, accept the results that are above Q*. So that gives you basically the capability share. 1 minus h is basically all the share of tasks that you can potentially, with this machine, automate. To the right, focus on the solid line. That's the distribution of costs. Now, when I talk about different models, let's think about different models for now just to simplify the setting in terms of the size of the model, in terms of the number of parameters. As I increase the number of parameters, the capability shifts to the right, meaning that the share of the task that you can actually— you're capable increases. But that comes at a cost because the cost curve shifts to the right. The distribution of cost shifts to the right. Now— OK, so the situation is getting dire as we move further into the— yeah, so this is— I made the mistake of pressing back, and now— we are in trouble.

[07:11:21.51]

I would need some help. I don't know what to do with this. Thanks. Some of the figures are also distorted, seem distorted. Well, I do know, but yeah, it's further forward. This is— we're in the appendix, and there are many slides.

[07:12:02.54]

We have to go all the way.

[07:12:19.57]

You're getting several previews of the— the slides, so it's probably a good thing overall. Okay, good, good, good. So, yeah, I apologize. Some of the slides is clearly— you're going to see them a little bit distorted, so hopefully— so for instance, this one is a little bit shifted. And I'm going to be brave and go forward hoping that this is not going to break us down. OK. So I guess it doesn't like clicks. So what's the automation decision like? For each task variety new, you're going to look at the left, the total cost of using automation, using AI. You pay the compute cost plus the labor cost. On the right side, that's the direct labor cost.

[07:13:02.00]

And that gives you the usual cutoff for automation. So, the cost— compute cost has to be below gamma star, which depends on the length for the task for humans, omega— 1 minus omega, the amount of time saving, and then the ratio of wage to the price of compute. Now, importantly, in a model like this, we get this neat result that the total share of tasks that are automated, phi to

the left, is the product of the share of capability— that's the expression we've already seen— times the conditional comparative advantage. What is the share of tasks that have compute costs below gamma star, right? And so, now we have basically the result that we want. We can go to the next slide and look at the machine choice. So, now that was characterizing the decisions for one machine. That gives me the marginal cost, cost, which inside the square bracket tells you what's— how much saving you're going to get. So, the saving, of course, grows with the amount of capabilities. The more capability, the more chances of savings you have. But then conditional on capability, there's this per-variety saving in terms of cost. That's this zeta function, which is again a function of the cutoff gamma star.

[07:14:17.11]

And what do you do? You basically the figure to the left that shows you how much you're going to get in terms of capability. To the right, larger models, that they give you more capability. But then, the saving declines as you go to the larger models. You look at the product of those 2, and you pick the model size that minimizes your marginal cost, right? So, that's the way the model works. OK. So, now, let's take that model, put it into this general equilibrium environment. On the side of the menu that the firms face. There is—

[07:14:50.17]

we're going to basically index everything in terms of one dimension. Again, think of it as the parameter size, the number of parameters. And there's the frontier of that, and bar is growing over time. To do that creation of the frontier model, you have to train it by the amount of data, some function D of M. We're going to come back to that in the next slide. So, then, final output Y, is an aggregate, CS aggregate of occupational output, Y0. Occupational output Y0 is a Cobb-Douglas aggregate of non-AI capital. And I apologize that now we have to rely on X to show you capital because we've already used capital in the standard way that in this literature we use it to show AI itself. And then, the non-capital part, and this is non-AI capital, by the way, this is the equipment capital for each sector structure of things. Then we have yj, which is the own task we were talking about. Each occupation o has a different mix of those identified by that vector mu of occupation o across task types j. So now, for each task type, we're going to basically rely on the same production function you saw.

[07:15:56.04]

I have to characterize h and g, capability and cost. H, I'm going to assume the quality is some average quality for a model of size m. for a task of duration tau j. j is basically this particular own-it task we're looking at. And then there's some dispersion around that that is u. Those of

you who are familiar with the idea of the jagged frontier, the jagged frontier is going inside that u, the stuff that is not predictable based on the size of the model. Then we have the cost distribution. Once again, there's an average log cost that is the gamma function as a function of the size of the model and the duration. And then there is this additional dispersion in costs. And then finally, we're going to assume that there is this adoption friction omega that varies by occupation and also over time. The investment— non-AI investment is pretty standard. This is not an exciting part of the model. Where it is more interesting is the computing sector. So computing sector, we're assuming, is also accumulating capital. What is this sector? This is basically NVIDIA. This is the data centers. This is This is the GPUs.

[07:16:56.57]

This is the stuff that you use to produce compute that then is used by AI to translate compute into the input for automation. Importantly, look at the price of compute, RT, that has 1 over AA is the investment productivity. This is going to be our first driver. This is where the enhancement in the productivity of the compute sector is going to lead to the decline in the price of compute. That's one of our drivers here. We're going to assume there's some potential wedge on top of that, what firms are paying. Perhaps there's some market power that those firms have. And then, AI foundation model sector, here we're going to basically assume things are exogenous. There is some leader of that industry that decides the rate at which they're training the frontier models in bars. They have to pay this cost, compute cost, that is a product of model size and data. Their tax finance. We don't get into the details of where their incentives come from for now. We clear the market. And now we're ready to go. OK, so now let's talk about the data. As I mentioned, what we are doing here is that we're basically looking at all the set of tasks that, back in the day when we started this project, GPT-4 decided that they're LLM-addressable.

[07:18:12.16]

There are 11,000 out of 18,000 tasks in O*NET data. 55%, 60% of all tasks. So then, as I mentioned, what we do is that we have 40 models. The data is getting completed. So what you see right now is what we have as of today. 800 occupations, over 10,000 task samples, tens of thousands of evaluators, and 41 models. So then now what we're doing is that we're going to these workers in those occupations, showing them an example, and we tell them, how would a manager view this response to the task? And importantly for us, 8 is the holy grail. It's the average worker quality that corresponds to Q* in our model. We also ask them about the task length. That's how we're going to characterize tau j, the length of time for human workers. And we also, of course, measure token costs and price per token. So we basically get to see the 2 sides. How do we use that information? We first estimate the average. Remember, capital Q is

the average distribution— average, the mean of the distribution of capability h. So now, instead of model size, here we're going to rely on these index— epoch capability ECI index.

[07:19:26.20]

I can tell those of you who are interested offline why. We also can do this with model size as well. But there are reasons to rely on this because when we get to the algorithmic progress, It gives us a more robust way of measuring things. Okay. So, what do we see? The first part— there are a bunch of other controls, but the important thing for you is to notice that high ECI, high benchmark models, they do perform better on everything. Longer tasks, more complex tasks, AI performs worse. High ECI models are particularly good in more complex tasks. And these coefficients, these slopes also vary. We allow them to vary by family task family. So they're different, for instance, for writing tasks versus coding tasks and other things. But this is just basically the broad outline of the kind of forces you're going to see in there. Okay. So we now have a model of our capabilities. We can do this for every task in the US economy. Then we move on to the compute requirements. Here, things become a little bit more nuanced. Again, I'm happy to chat about the details of why we do things the way we do them.

[07:20:28.35]

we do, but we break the cost into 2 parts— the number of tokens that a given model creates for an answer, and then how much it costs to create that token for that model. And so, for the first one, we use, again, the same specification as before. As you can see, larger model, they actually spit out more— they use more tokens to answer to the same question. Sorry, high ECI models, more advanced models. And they do so more for longer tasks, but not by as much. And then, as a function of the model size, producing every given token is becoming more expensive for larger and larger models. That's the second part. So this— and then we also have the dispersion of all these numbers. That dispersion is directly read off of the data. So we have the full distribution of G or costs.

[07:21:16.23]

OK.

[07:21:17.39]

So now we're ready— I'm sorry, how much time do I have?

[07:21:23.08]

So, we're now ready to move on to the calibration results.

[07:21:28.25]

Okay, good. Thank you. I'm going to spare you a lot of details. The bottom line is that we have a few additional moments that we read off the data, the data that we actually have to painstakingly put together. So, these numbers are not bulletproof, and once in a while when some company wants to go to IPO and some new information comes out, we have to go back back and change these numbers, as it happened a few weeks ago. But essentially, there are 3— the key moments are how far away from the frontier on average we see in terms of cost firms are. And those are the numbers we fit here. The share of spending on AI inference relative to the labor payment. And then importantly, to the right, these are the numbers at the occupation level for adoption. And the model can nearly fit all of them perfectly.

[07:22:17.31]

Why?

[07:22:17.42]

Because we have this overhead labor cost, this friction that actually acts as that wedge that ensures that the model— the gap between the model in terms of capability and cost and the observed data is closed. Except for farming, because farmers apparently use so much AI in the real world that our model is just throwing up its hands. It's like, based on the facts I see on my side, I don't know how they can adopt AI so much as we see in the St. Louis Fed data. OK, so now, first result. To the left, I'm showing you in 2026— by the way, we do this for 2025 and 2026. I forgot to mention that. OK, so now, in 2026, I have roughly 6% of tasks that are automated. I can break this down to the contribution. Remember, this is a product of 2 terms, right? Capability and cost. The cost side is pretty uninteresting. Basically, if you use a model, it looks like you basically go all the way. If I decide to use Opus, I use— I automate all the tasks that I can with that model. I max out the capability. But what's interesting is that it's not the frontier model that determines how much automation we're getting.

[07:23:26.38]

If I was using the automated frontier model, I would've automated 17% of the task. In fact, I decide to optimally use these smaller models. And that is basically what reduces the 17% where the frontier is to 6% what we observe in the data. This is also true of the cross-section of tasks. If I look at the variations in the cross-section across the own task, 75% of the variation comes from the choice of the model for that task. So there's a lot of variation across tasks in which

models we are using. So now, let's project forward. So, I have— I mentioned there are 4 different drivers. What is the first driver? I talked about compute price. That's where the productivity of the investment in AI capital comes into the picture. Historically, this number has been very dramatic, really high. So, our estimates are about 20%. 17% would be a benchmark number to use. But what's happened is that between '25, '26, this rate has declined substantially in magnitude. What's happened? Well, in the background, there are some capacity constraints. The demand has been surging so much that it seems like the sector, the compute sector right now is at this capacity constraint that raises the prices.

[07:24:43.31]

So, here, we have to take some scenarios. We have the historical trend, and then we have the more conservative, which is the recent trend, and then something in between, that's the moderate one. Adoption friction. So, in the adoption frictions that we've— the wedges, omegas that we've estimated, there's actually decline from '25 to '26. So, our approach is to take the change and then project the same number, assuming that's our headline number. But then, we could, just to explore the different possibilities, we could be more conservative, take a number that is on the 80th percentile confidence interval of our estimate. Then it comes to algorithmic progress. So I didn't mention this, but in the specification I showed you, I could basically now take the ECI and project it to 2 components. How much of it is explained by the model size evolution over time? How much of it is directly tied to the release date? It turns out that this contribution is really, really high. So the model of the same size, same exact size, in every way similar, if they're released 6 months later, their ECI, their benchmark capabilities grow substantially. So that becomes, for us, a driver itself.

[07:25:51.55]

So basically, improving the productivity of the model. And then finally, scaling. This is the standard in the literature, at least in the early ancient ages of GPTs, GPT-3 and GPT-4, people were focused on scaling. So the scaling also is the other component. So we project the ECI, this benchmark, on time, release time, and model size. And this contribution is going to be separate for us. OK. So, let's start with prices. I'm going to fix the AI technology as of '26 and the friction as of '26. And I'm going to show you what happens with our 3 different scenarios relative to the background. The background is 1.1% growth in GDP. So, as you can see, the gains are pretty small. They're basically 0.1%. They're very Asimov-esque in terms of AI predictions, right? Okay. Now, let's move on and focus just on frictions. Now, we freeze everything else and we just ask what happens if you just continuously get better at using the AI we have today. Same prices, same AI capabilities. And this already gives us 1% growth over the next 5, So, the top left is basically real GDP growth. It doesn't come at much cost in terms of labor share.

[07:27:09.17]

That's— we'll see, you'll see— is a feature of most of the results we have today. So, this is great news. So, but then, of course, it's going to run its course. Why? Because once you actually max out what you have today, then you're going to basically stay there. So, now, let's bring in everything except algorithmic progress. So, what I have is price, that's what you already saw. It's not very strong, it's pretty mild. Then on top of that, I add scaling, that's the orange one. You can see that actually, you get more mileage, but it doesn't even— within the next 10 years, it doesn't even catch up with the friction only. But if I combine all 3 of the forces together, that's where things get more interesting. Now, I'm projecting up to 2.5% growth by 2035, in 10 years from now. What's happening to labor share? It's not budging by much. It is declining, but not by nearly as much as you would have expected, even though the share of automation is going all the way to 12%. So, now, let's bring the big guns. This is now where we bring in the algorithmic progress.

[07:28:12.29]

So, these are the 3 scenarios we have. So, the number we have as the benchmark is the number we see in the data. Then, we try our best within— so, this is 4 years of data. on models, right? Remember, this technology has been around for 4 years. We have 4 years of data, right? So, I basically try to find the period of time where this growth has been minimum. That's my conservative estimate. And even with that estimate, by 2035, these numbers have never been witnessed in any frontier economy in the world, right? So, this is where we find ourselves. Even more interestingly, Look at the labor share. Even then, it's not budging. So, what does this look like? It's essentially TFP gains. If I do a growth accounting of the results that we see as of today, and I have to caveat everything I'm telling you, these numbers are changing. Again, we get more data, we are adjusting. I just wanted to use this as sort of a progress report to all of you. This is essentially TFP growth. Labor is happy, capital is happy, everyone's happy. or just becoming richer over time. So the last slide I want to— so we have an additional extension of this framework.

[07:29:20.10]

Already we have some results looking across different countries. There we basically focus across countries on where they are for demand for skills depending on where they are in the process of structural change and importantly on their skill endowments, which we haven't talked about today, right? And so hopefully we get the type of access to the kind of data Filiz was talking about earlier to be able to be able to speak to how these scenarios are going to unfold

across different countries. I'm going to stop here. I think hopefully you've already gotten the big picture of what we did here. Thank you so much.

[07:30:02.47] - Jacob Funk Kirkegaard

Well, I mean, there's no reason to go on if you have capitalists happy, laborers we just get richer. I actually hope the discussion is not going to pick that one apart. But we've come to the discussion part of our last panel. So we're joined by Rodrigo Flores, who comes from GW University School of Business. Please, floor is yours.

[07:30:25.46] - Rodimiro Rodrigo

Okay.

[07:30:28.14]

All right. Here's the learning process. Keep going. On that direction?

[07:30:35.14]

Okay.

[07:30:37.22]

Okay.

[07:30:38.30]

I know we're all tired, even the computer. So many slides that we didn't see, but the appendix is amazing. All right.

[07:30:58.25]

Thank you.

[07:31:00.54]

Well, thank you so much to the organizers for giving me the opportunity to learn from these 4 papers, as Daniel was saying. So all of them fascinating, really. I would say these 3 papers have 3 things, big picture, in common. One is that they are all talking about the diffusion. of this new technology. Also, they are all working with fascinating, unique data. So that you are going

to— I mean, you already see it and I'm sure that you agree with that. And the third thing, and just also a little bit motivating where I'm going to be building some bridges or bringing some of the discussion, they are all focused on developed countries. And it's not a complaint because it's where the technology is. But I think I'm going to try to push a little bit on thinking a little bit some potential questions about the developing country as well. So big picture, John, unfortunately not here anymore, but I had the pleasure to to discuss a little bit with him my thoughts. He's going to be looking again, these 3 papers, diffusion, but in different lens, right? John is going to be talking about these firms and asking these 750 CFOs, how is going on this process of adoption?

[07:32:40.49]

He found, he showed us that the adoption is being broad, that the gains are modest. I mean, although All of them seems to be quite positive about the benefits of AI, with very minor effects on overall employment, although very consistent with what many of you also show us about being the clerical workers the most affected. Then we are going to go— I mean, I switch a little bit the order. You will find out why. But then Daniel was telling us about, or thinking more deeply about this, what limits the diffusion or the automation, right? Thinking a little bit more deeply on the choice of the development and the development of automation, right? So in the previous models, we had automation to be, or the decision of automating, as a payment of a lump sum cost. While here, Daniel, in a more realistic way, I mean, in the other paper that you didn't see, he was thinking on the case of these 40% of the firms that John was telling us that actually pay for hardware, develop the technology in-house. In this other paper, he's going to be thinking on these other 60% of the companies that are actually paying a rent and really more deeply thinking on which are the tasks that they should automate.

[07:34:15.46]

And finally, on the network side with HANA, we're going to be seeing, again, a very interesting dataset with 29 technologies, 17 countries, where AI is going to be spreading widely and are going to be finding very different patterns versus previous technologies, right? The distance is not going to be so relevant, the global value chains, et cetera. So, let me go into the first paper. And well, of course, my comments, take it as probably some thoughts for your follow-up project, right? In the case of John, Also, kind of quickly summarizing, we were saying that we saw already a big adoption. 60% of the firms said that they already adopted. But for next year, I mean, 85% of the firms will be— well, by the end of this year, they should be 85% adopters, right? Also, very interestingly versus previous technologies, we see a lot of small firms adopting in his case, right? Of course, he is looking at, or he's interviewing large firms indeed, but that is

also quite promising in terms of this new technology, right? Very much associated to how they invest, right? The smaller firms are more intensive on this rental side that we were talking about.

[07:35:50.53]

AI arrives also in terms of of this expenditure as a service that I already mentioned. And in terms of the sources of the gains, are going to be very minimal on the capital deepening side and very much associated to the TFP. This is the narrative or his reading, right? The jobs are going to be not much reductions and mostly are going to be concentrated in these clerical workers that are going to be suffering the most, no? On this graph that I'm plotting for you there, I mean, I have this very interesting narrative that the paper has, that given that they ask the CFOs what is their perceived productivity gains that they attribute to AI, They report this 1.8% for 2025 and 3% for 2026. Nevertheless, when they are asked about the actual growth in output or the output actually on revenue and in labor, they increase the real measure is 1.6% and 1.8%, right? I mean, which is very, Interesting that that is coincidental, the number of 2026 with 2025. But I mean, it's something that happens, but that is going to be part of the thoughts that I'm going to bring to you, right?

[07:37:29.35]

The first thought is going to be whether CFOs can really identify the effects of AI in the way they are asked, right? Because the CFO at the end of the day doesn't have the counterfactual. I mean, she observes how much they invest on it, right? And she may be actually quite excited about the decision they made, right? And so kind of in a way, the interpretation that we are giving it to this AI-attributed growth growth can be a little bit too much, right? And so that is something that in the way for 2026 really is relevant to analyze what are the realized outcomes that we observe there and trying to disentangle whether this is an ex ante optimism or really there is something that can be pointed out as an actual productivity paradox, right? Remember that what we have seen in the productivity paradox, at least with previous technologies, is something that doesn't happen from one day to the other. And it's because it's a complex process of learning from the organizational perspective rather than something that could happen from one year to the other. But so it's something that there is a space to keep learning from that.

[07:38:55.01]

Kind of related with that point, Right? Well, the result of that debate that I was saying is that then the reported AI productivity gains are going to be larger than the implied gains. And so, yeah, I mean, also including in the next wave some questions to get a little bit more granular on hours saved. For instance, we have this interesting work from Hulum, right? That with a similar survey

in Denmark, he finds out that workers report to have all of them half an hour extra, right? I mean, or savings of half an hour. So it is not at all that the CFO is wrong, right? She sees, again, she sees the benefits. It's just that what at the plant level the manager managed to do with each of these half an hour of all of the workers and how that is converted into productivity is a function that is more complex. So learning a little bit more about the redeployment process, the quality or the output units, even the hours saved could help us to actually get more granular in these measures of of productivity on the firm side, right? Rather than telling me what is my productivity gain, that is, it is a very complex question.

[07:40:25.50]

I think that here, John also opens another very interesting area to understand better, because in the accountings that the paper does, it is also said that the capital deepening explains only 15% of the gains, and most of the gains are going to the TFP. And so here, I mean, many, many things to work with, no? For instance, rather than value added, they have revenue. So it would be interestingly that, given that they have the operating cost, to at least work with the revenue minus the operating cost. To separate a little bit more these 2 things. It is something that I think that more work is going to be seen on this because remember, here again, versus previous technologies, in case of robots, for instance, the technology was capital, right? So in this case of AI, how should we account for AI? Is it capital or is it just a service? So, in the case of the paper, it's taken as a service. And well, let me skip the last point to just jump to some policy implications. So, what I was telling you about the accounting measure, the labor redeployment to understand a little bit better, to try to talk about the productivity paradox, I think is something important to keep in mind.

[07:42:04.34]

And let me jump to Daniele, where I was already kind of sketching what is, I mean, a big contribution here that is moving from those models that we're having these fixed costs to pay, right? To really think on a more deep way on the firm perspective, how So they picked a model for each of the tasks and defined which tasks are going to be automated and with which model, after minimizing this cost, interesting cost function. A second big contribution is that he's going to be taking very seriously how to discipline this model. doing this extraordinary effort on going and serving real workers, how they observe or how they evaluate these products of AI doing what would be their work. And so, I mean, with all this machinery, he's going to find 6.1% of task time And interestingly, well, he already showed us that it's not going to be the price of the technology or even the frictions, right? Removing the labor frictions, which I think that there is a space there also to learn a little bit more. But even removing these 2 things that is important that the model has, the gains are only 2%, right?

[07:43:48.05]

while it is mostly the progress on the technology what takes us very far away. And this, as he was mentioning, I think that it also gives some explanation to why Acemoglu and Restrepo, in the first paper a couple of years ago, they give such a kind of small prediction of the benefits of the technology, right? Because of course, the technology that we were seeing in those days, had a very different progress to what we are working with nowadays. So actually, I'm going to jump into the policy implications, if you guys don't mind, because I see that I have 1 minute. OK, so frictions. OK.

[07:44:42.46]

Sorry. Yeah, jumping ahead made me lose the track. So yeah, I think that there is a lot to learn, as I was saying, about the labor frictions, right? Something that is a little bit interesting or striking on the model is that one of the results is the no change in the labor share. So I think that that invites also to more reflection on that side, and we can can keep chatting offline. I think that— yeah, sorry. Yeah, I think I mixed some of my comments from the other paper. I'm sorry. So, Yeah, so I got a little bit confused with the stuff. Oh yeah, this is a— yeah, I see the problem. This is a previous version of the slides as well. So I sent one last night. Well, let me just say that, yeah, that I think that the model is fantastic. I think that there is some things to keep thinking about, the frictions of the labor frictions that he considers on it. I think that, as I was saying earlier, what happens within the firms with the adoption and how that is going to impact on the dynamics that you are seeing, in particular, in terms of labor share.

[07:46:31.12]

I mean, the fact that everybody is winning in this story is kind of striking on my side. But I mean, it's something that is worth to keep analyzing, right? And so, I mean, jumping very quickly to the last paper. In the last paper, I think that it is, again, a fascinating dataset. I should wrap up. Okay, so I think—

[07:46:55.44]

Excellent. Okay, thank you so much. I think that the diffusion here is amazing. Let me just, given that I have 2 minutes, let me jump to open questions of these 3 papers. I would say one of the questions is whether will capital gains become productivity gains and what it has to happen for that to be true, right? I mean, we have some lights on that, but there is still so much work to do on that. What are the adoption frictions and how fast do they fall? How is AI financed? We are seeing that more and more is more of a renting type of technology. So what are the implications

of that? And probably that is an explanation of the wild adoption, right? Does AI create or relocate talent? We are seeing, and let me finalize with that. I mean, I think that some of these papers touch on the fact that one of the most automated occupation, right? I think Daniele and Hannah as well, in both cases tends to be clerical workers. So I think that that used to be or is an occupation that was very labor absorbing in the developing world. And so what happens once it's redeployed by AI is going to have important implications.

[07:48:30.12]

Okay, thank you so much.

[07:48:43.43] - Jacob Funk Kirkegaard

Well, if I can ask the last panel and scene-setting presenters to join me. There is one last opportunity for the audience to ask questions to the authors of these fascinating both datasets and papers. Once again, of course, I recognize that it is late in the day. But there is microphones in both the— at least, yeah, in the front as well as in the back.

[07:49:31.59] - Audience

Okay. Matthias of the OECD. So my question is for Daniel. So I was really struck in your model where you showed towards the end that algorithmic progress rather than cost basically reductions of those models can lead to really substantial growth rates in your model. And you also mentioned earlier that your model does feature kind of the standard weak links setup. And I was just wondering what in your model makes it that algorithmic progress, at least over the projection horizon to 2035, can overcome this traditional weak links setup where other tasks where AI cannot actually lead to improvements will become bottlenecks. And if that's something that maybe is happening mostly during the near-term transition period as cost-effectiveness leads to a broader application of AI, or if that's something that also holds up kind of for longer projections into the 2040s and 2050s, assuming that algorithmic progress keeps pace?

[07:50:41.34] - Danial Lashkari

That's a really good question. So we live— our model is a model of weak links. We are living in the same world that Chad Jones and Chris Tonetti in their paper, recent paper, they've depicted for us. What is happening here is that there is a share of tasks, 60%, 55% maximum, that we are allowing because we're only looking at language model tasks to ever get automated. The pace of progress is so fast in terms of capabilities that you're burning through all of that in the course of 10, 15 years. And there's some progress that is going to come out of that. Now, if you take the elasticity of substitution, which we take to be 0.5, to be 0, the amount that— you're

going to get a smaller growth, but we think of 0.5 to be a reasonable number. In their paper, they take number closer to 0.25, but you should still have— you live in that world. What is happening again is that the pace of progress we are observing in the past 4 years is incredibly high. Now, I have to caveat— I didn't caveat this result sufficiently. We're already going back to the data and trying to look for some evidence of potential decreasing returns in— So what is happening in the background is that these small models are becoming so capable, if we project forward the trends that we've seen in the past 4 years, that essentially in a few years' time, you're going to be running a small model off of your laptop to run an entire company, right?

[07:52:13.00]

So potentially there has to be some kind of physical laws that prevents us from being able to go too far in that direction. It's just very hard to look at the current data and see vestiges of that, evidence of that. So we are working on that to see whether there's a sense in which we can bring that conceptual point into the framework. But that is what is happening.

[07:52:33.05] - Jacob Funk Kirkegaard

Thank you. Any further questions?

[07:52:47.43] - Audience

Hi, I was curious about the—to Daniel again— why the labor share is staying relatively constant even under 50% automation.

[07:52:59.22] - Danial Lashkari

That's another thing that we are trying to understand. So, what is happening as of this calibration, and this is by no means something that is intrinsic to the model, it comes down to the quantitative setting that you're looking at, is that what's happening is that because when you commit to a model you go all in, the only margin of adjustment that you have in the— let's say the price, so what is the elasticity of substitution? That is the key point here, right? Is the elasticity of substitution between AI and labor in the aggregate less than 1 or greater than 1? It turns out most of the action is going to be at the task level.

[07:53:37.06]

Excuse me.

[07:53:38.45]

So what is happening right now is that when the price of compute changes, at the model level, there is not much more you could do. So all the change comes from the change from one model to the other. So when the price of compute What is the impact of that changes? What is the impact of that on your labor demand or your compute demand? It's going to depend on what new model you're going to choose. Because the impact of model size, the new model, the change in that margin is pretty mild, then the change in your labor demand, capital demand is fairly muted. That is basically as of today what is happening. So that if it was a different scenario where you were— you hadn't maxed out the possibilities within the same model, then you would respond much more strongly to a change in the price. I don't know if that intuition holds any meaning to you, but that is my best understanding of what is happening currently in the calibration that you're seeing. So, that becomes crucial, right? Because if the elasticity is greater than 1, that's where you're going to see the share of income going to labor decline pretty fast.

[07:54:50.27]

Right now, we're basically in a capital loss world. Interestingly enough, there's nothing baked into this, by the way. It is just numbers as they come out.

[07:55:00.10] - Jacob Funk Kirkegaard

So we have a question in the back and in the front. We'll take both of them and then we'll hand it over to the panel.

[07:55:07.13] - Audience

Hi, Samuel with the IMF. A question for Filiz and the survey that you're running at the OECD. I was wondering if you can go into more detail about who were surveyed?

[07:55:18.48]

Because I was wondering about this personal versus in-the-job use of AI. I was wondering if these are purely workers or you're also surveying managers and other people. And related to the fear of AI and the use, I was also wondering if this applies only to workers or you also see managers who are more averse to using AI.

[07:55:41.34]

for different tools, for different uses.

[07:55:45.34]

Thank you. I have a question for Daniel as well. So one speed brake, you know, on this very fast growth projections you have is the one that was mentioned, you know, by Jones, that some of the AI tasks are complementary with human tasks whose productivity doesn't increase as much. But another one is is also potentially the energy, right, that's required to operate these models, especially the more sophisticated one. And if everyone starts, you know, adopting them very extensively, then, you know, those—

[07:56:22.48]

I mean, this could put very big strains on energy supply.

[07:56:28.23]

So maybe that's a factor that you don't consider in the model, you know, that could reduce those projections? Just would like to have your take on this.

[07:56:37.05] - Danial Lashkari

That's a very important point, and we are hoping to go in that direction. Right now, this framework is very simple in that— along that dimension. As I mentioned, the compute sector is just basically some capital. You're saying that there's some energy that goes into it. But again, think of this as the rosiest view of AI you could ever find. We happened upon it. We didn't design it, but that's where we are. You don't need compute. If algorithmic progress is what wins the day, you don't need energy, you don't need much because in fact, if anything, if you look at the contribution of additional capital to the growth we're finding here, it was tiny. What you're doing, what's happening is that you're doing more with less. So everyone is happy, even our environmentalist friends are happy in a world like that.

[07:57:27.43] - Filiz Unsal

Thank you. Yes, so we have this distinction between upper managers and middle managers and junior managers in the data. So what we have seen is that in terms of the adoption rate, yet we don't see that much of a difference. I mean, of course they are more exposed, right? Upper managers, overall managers, they are more exposed category, but we don't have really much of a difference in terms of their adoption being higher or smaller in the adoption of the core tasks. So of course, as I mentioned before, we just got this data a few weeks ago. We will do more slicing and dicing. But so for now, for the graphs that I presented, that 20% number is pretty common across all the occupations. We have Matthias here, one of the co-authors. I don't know if you have anything to add, Matthias.

[07:58:29.10] - Audience

I think the only thing I would add is that we can look at the sentiment questions by occupation. That's actually something that we're doing now and we haven't shown here and there are interesting patterns across across occupations to be explored. We will also— let me just iterate that, reiterate that— make the microdata available to the public in a couple of weeks from now. So a lot of dicing and slicing can be done by everyone here in the room. So, yes.

[07:58:58.46] - Filiz Unsal

Go do it yourself.

[07:59:04.41] - Jacob Funk Kirkegaard

Very good. Well, we are very close to our witching hour. So unless there are last burning wishes to comment from the panel—

[07:59:17.38] - Filiz Unsal

Maybe I actually have a question for Hannah, if you allow.

[07:59:21.01] - Jacob Funk Kirkegaard

Absolutely.

[07:59:22.04] - Filiz Unsal

Yeah, just because it's very interesting you mentioned that global technology diffusion is actually not that common in AI compared to other technologies. We do observe these global spillovers in technologies. So what does this imply for countries in terms of how they should develop their own AI global supply chain? Is that the conclusion or is it jumping a bit too much?

[07:59:50.47] - Hannah Wei

Yeah, I think one of the caveats for our current results is that we only measure AI adoption in the early days from before 2022. So before the birth of the widespread of ChatGPT. So our conclusions are only referring to the early days.

[08:00:06.20]

And firm networks do still have a very positive influence on adoption, but it's just comparing to other technologies. Maybe it is still too early to have a firm conclusion on whether they affect AI so significantly. But with more data, as I also just heard of today, that there's global data also from Lightcast that's possible to extended to the most recent year. So we can also extend the analysis to include the GPT and also CloudWave going forward.

[08:00:39.50] - Jacob Funk Kirkegaard

Very good. Well, when you ask a frontier model how, as an organizer, you ensure that you wrap up a conference on time, one of the things that it always tells you is one of the first things to do is you give the chair of the final panel the right to give the closing remarks. So that is what we will do right now.

[08:01:02.59]

And therefore, on behalf of PIIE President Adam Posen, it is my great privilege to thank everybody here, everybody online, our co-organizers with the OECD, the IMF, and the World Bank for what I certainly consider to be a flawlessly executed—

[08:01:26.20]

thanks to our always performing PIIE event staff— flawlessly executed and enormously enriching conference. I mean, we all hope that you leave here with some AI questions answered, but far more importantly, you leave here with knowledge of new data set and hopefully with a host of new questions that needs to be answered on this issue in the future so that this conference has helped influence forward-looking research agendas in the coming years. And with that, by the powers vested in me by the printed agenda, I declare meeting adjourned. Thank you very much.