

Peterson Institute for International Economics

AI + work: Understanding AI's impact on the labor market

March 10, 2026

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

[00:03:48.26] - Peter Orszag

Hello everyone. Good afternoon. I'm Peter Orszag. It's an honor to be back at Brookings for today's discussion on artificial Artificial Intelligence in the Labor Force. Whenever I return to Brookings, especially with the Hamilton Project, I do have a deep sense of back to the future. It was almost exactly 20 years ago today, April 2026— 2006, sorry— when a group launched a new initiative right across the hallway, invited a young senator who had recently been elected to his first term to speak at the event. That young senator was Barack Obama, and the group was led by Bob Rubin and Roger Altman, and I was honored to be the first director. Thus began the Hamilton Project, and 20 years later, it's great to see it going strong under Aviva's leadership. But we're not here to talk about history. We are here to speak about the future. So thank you to the Hamilton Project, the Yale Budget Lab, and the Peterson Institute for International Economics for putting together this event about how AI is reshaping—not just our economy in the abstract, but also the day-to-day realities facing firms and workers. Let me start with one macro point about the labor market and shocks, which is that in general, the labor market can adjust well to small shocks, even if they happen fast.

[00:05:07.23]

If one factory cuts a shift or a new software tool automates a single task, for example, firms and workers can usually adjust relatively smoothly and rapidly. Similarly, history has shown that the labor market can also respond at least moderately well to large shocks that unfold slowly. This is the oft-cited history of past technological innovation, as a new technology was adopted relatively slowly and the labor market had time to adjust and create new types of jobs. We've all heard the story about ATMs and agricultural mechanization, and factory electrification, all of which followed a similar pattern of relatively gradual diffusion, innovation that occurred, and new jobs coming up that helped the labor market to adjust. But that history should not give us much comfort for a different type of shock, a large shock that unfolds fast. This combination of a large shock that happens fast is one where I doubt the labor market adjustment would be as salubrious as past innovation. So a key question is how deep and how fast the AI revolution will

take hold in the labor market. As Jed's paper, which anchors today's discussion, makes clear, we're still in the very early innings on this issue.

[00:06:30.42]

His analysis reviews the early and still quite fragmented evidence on how AI is showing up in the labor market. It highlights how much uncertainty remains about which path AI will ultimately take and how its effects on employment and wages will net out once we move beyond the first round of case studies and anecdotes. In the face of this substantial uncertainty, and with significant humility in doing so, given the wide range of potential outcomes and the high risk of being wrong, I want to articulate 3 bets to you today. My conviction level around each of them is only moderate, as it should be given the substantial uncertainties involved, but still strong enough to articulate them in a setting like this. First, my bet would be that although AI has not affected the labor market all that much to date, it will do so increasingly over the next 5 years. In my experience, the technology is improving at such a rapid rate and its implications—its applications are so potentially widespread that I would take the side of those who believe the labor market effects in the future will be quite substantial. A second bet I would make is that AI will accelerate a trend we have already been seeing, a widening spread in corporate performance.

[00:07:47.00]

In collaboration with McKinsey, a team at Lazard has examined the return on invested capital, or ROIC, for U.S. companies over a 20-year period. In 2003, the firm at the 95th percentile of that distribution in the United States had a ROIC that was 10 times the firm at the 10th percentile. Today, that gap has grown to about 25 times. Similarly, David Autor and co-authors have documented that over the past couple of decades, a small group of superstar firms has captured a disproportionate share of profits and investment, even within narrowly defined industries, while many competitors have struggled to keep up. To be sure, there is a very real possibility that the leaders of today will become the dinosaurs of the future, and new entrants will become the frontier firms as the technology unfolds. But that will only be the case if the leading firms lack the culture and orientation to continue innovating in the way that startups do. I suspect that will not be the case. And if the frontier firms meet the moment, AI could deepen the divide in corporate performance by allowing the best-positioned firms to get more value out of their existing advantages in scale, data, and talent than others.

[00:09:02.34]

For thinking about the firms that may win in this environment, I find Andrew McAfee's notion of the "geek way"—even though I don't like that title—quite helpful. He highlights companies that

treat culture as an operating system built around trial and error, rapid iteration, and a willingness to overturn legacy processes when the data change, exactly the mindset that serious AI adoption requires. For incumbents, the central challenge is therefore not acquiring one more tool. It is building the organizational agility to use AI to redesign workflows and roles. When firms get this right, AI becomes a way to shorten decision cycles, break down silos, and continually reallocate tasks between humans and machines. A final bet I want to touch upon is the narrative about human judgment in an AI-rich world. It's increasingly common, or perhaps somewhat common, to see the argument that as AI adoption grows, what will matter most is timeless human judgment and broad intellectual range, while the technology merely handles the numbers and patterns in the background. So my third bet is that this view will increasingly be proven wrong. AI is already changing the train on which judgment is exercised by aggregating significantly more information, comparing across many more use cases, and quantifying risks that humans systematically misperceive.

[00:10:32.15]

I think I accidentally touched the computer. Sorry about that. In my own life, I've increasingly relied on AI not just for factual questions, but as one input into my own judgment on a wide array of topics to complement the human kitchen cabinet I use for important decisions. I don't use the AI in isolation or in a definitive way, but then again, I don't take advice from humans that way either. I'm now 2.5 years into my tenure as CEO of Lazard, and I've used AI tools to assess how to make the next 2.5 years effective and successful. The tools, not a human, came up with the framing of contextual alpha, that in making business and investment decisions, the narrow alpha perspective of the specific math involved in a business or investment decision is too narrow. And the broader macro geopolitical and regulatory factors that are the contextual part of contextual alpha often define whether the decision is sound. Lazard has always prided itself on what we now call contextual alpha, and that term came from an LLM. Not a human being. I've also used the tools to assess my calendar in terms of how I'm spending my time across clients, internal topics at Lazard, interaction with our investors, and thought leadership along with travel, and for things like suggesting additional client invites for salon dinners or conferences.

[00:11:55.47]

These all involve judgment and not just mechanical computation. I'm finding the tools increasingly impressive on medical and legal issues. Never as a substitute for professional advice, but as a supplement. I've used it, for example, to translate complex risk information into clearer probabilities. As pretty much anyone trained in statistics who has interacted with the medical system will tell you, many medical doctors are famously bad at Bayesian statistics. So having a system that forces clearer probability estimates can be a very useful supplement to

their expertise. The broader implication is that judgment is already not a purely human preserve, one to which AI cannot contribute. Instead, AI can already help make judgments more rigorous, more consistent, and more widely accessible across an organization. And in my view, that will only go in one direction in the future. In practice, these systems are already capable of informing a meaningful share of what used to be labeled elite judgment. And for now, the best decisions likely emerge from combining human values and experience with machine-scale evidence drawn from millions of data points and patterns that are often invisible to the human eye. With that framing of my three bets in mind about the labor market, dispersion in organizational culture, and judgment, Let me hand things over to Jed.

[00:13:25.19]

His paper lays out the empirical evidence on AI and the labor force far more systematically than what I've just done. And later in the program, we'll return to what those findings might mean for policy. Jed, thank you.

[00:13:50.22] - Jed Kolko

Thank you, Peter. And let me add my thanks to the Hamilton Project, Yale Budget Lab, and Peterson Institute for bringing us together today. So I would like to spend a few minutes talking about the state of research about AI in the labor market. It is early to do something I rarely do, which is use a sports metaphor. We are still really in the first inning. And that has implications for how we interpret the wealth of research that's already come out about AI in the labor market. So my sense is that this surge of early research that we've already seen about AI in the labor market, some of the authors who will be here, who are here today and will be on the panel in a few minutes, this surge is still first inconclusive. Second, what it has found, I think, are pretty weak signals about what to expect about the future. And then finally, are really covering only one corner of what's a vast AI landscape for research. Now, none of this is intended to be criticism. The work that's been done has been excellent. It's been creative. It's been heroic given how young AI as a technology is and how long it takes new technologies to show up in labor market and economic data.

[00:15:11.59]

So the point of my explaining why this research is, to this point, inconclusive and a weak signal for the future is not to criticize, but rather to inspire and hopefully encourage much more research into what the effect of AI is and will be. So first, let me talk about why the research that we've seen so far is still inconclusive, why we don't have a clear picture of AI in the labor market. So the basic method that much of the early AI labor market research has followed is to

compare or predict labor market outcomes like the level of employment or changes in unemployment in more versus less AI-affected jobs. That means you need some kind of measure about how much each kind of job or occupation or task is affected by AI. And there have been a bunch of measures that have already been developed that capture how exposed an occupation is to AI, which could mean a good thing, how amenable that occupation is to being augmented, or it could mean something a little more threatening. What's the likelihood that that occupation could be automated? Another way this is measured is by current usage of AI in different occupations.

[00:16:36.33]

There are multiple AI exposure and usage methods. They disagree somewhat, and that's one of the things that leads to some of the inconclusiveness of early research. Two of the most prominent pieces of research have not necessarily conflicting conclusions, but at least conflicting interpretations or emphasis. This. One, and you'll hear from the authors shortly, that employment fell for young people in AI-exposed jobs. Another focused on the fact that unemployment rose less in AI-exposed jobs. So again, you can reconcile these findings even if the interpretations are somewhat different. I think at its core, it is still very hard to isolate the effect of AI from other factors that are affecting the labor market and the economy. Two reasons for this. One is the timing is still unclear. So the introduction of ChatGPT in the current era in late '22 is still too recent to account for some of the changes in the labor market we've seen that date back to that point or before. I'm going to call this issue with timing the macroeconomist objection, that the time series don't line up as you'd expect. But the other is the microeconomist objection, which is there might be omitted variables, which is a fancy way of saying that these occupations more exposed to AI might also be special in other ways.

[00:18:12.44]

For instance, The occupations that rank highest on these AI exposure measures for the most part are many of the same occupational sectors that were predicted to be more amenable to remote work, jobs that could be more easily done from home. And in fact, when you look across sectors, the correlation between AI exposure and ability to be done from home is something close to 0.9. Now, that doesn't mean you can't disentangle, but it does mean that you have to do some extra work in this sort of analysis to really try to tease out, is it AI that is driving some of the differences that we see correlated with labor market trends, or might these AI exposure measures actually be picking up some other factor? That doesn't mean you can't figure it out. It just makes some makes some of the early research a little less conclusive. My second point is that this work, even if it were more conclusive, would still be weak signals about the future. And

this is really inherent to what we see at the early stage of any kind of technology diffusion. LLM diffusion so far is really too recent for its impact to show up broadly in economic data.

[00:19:33.20]

Past tech revolutions, took years or decades to appear in data. And AI is advancing so rapidly that today's conclusions, even if conclusive, could soon be obsolete. The occupations that are the most AI exposed today could be very different from the occupations that are most AI exposed in future generations of AI capabilities. Um, and we're still at the point where research tools and definitions are still nascent. We haven't settled on a single way of describing AI exposure. We haven't settled on augmentation versus automation and other measures that describe occupations from each other. So there's still a lot of shaping to be done in how we talk about and research AI in the labor market. The third point I want to emphasize though is the work that's been done so far looking to see at Which occupations or tasks are most exposed to AI and what that means for labor demand is only one of many important questions in this vast landscape about what AI might mean to the labor market and the economy. There has been initial research in other areas, but much room for more work there. So, three additional areas beyond the labor demand studies that have proliferated and gotten a fair amount of worthy attention.

[00:20:59.42]

The first is on productivity, looking a bit more under the hood, often in a very detailed or more case— or almost case study way to understand how AI is affecting productivity within organizations. Most of these studies find benefits with some notable exceptions, though the studies are focusing on early adopters of AI who may be very different from the median adopter of AI years or decades from now. Team and collaboration dynamics tend to have been a bit underexplored in this early research. A lot of the focus has been on individual productivity. So, there's rich area for understanding more about how AI is affecting productivity in the workplace. There's been even less emphasis on what AI might mean for labor supply, how much people want to work. Which, of course, you have to look at in addition to labor demand if you want to understand anything about employment or wages or inequality or almost anything else we care about. AI could improve or hamper job search. It could change the relative value of leisure. Autonomous vehicles could make commuting much more pleasant, increasing the number of hours people have for work in a day. It could make leisure so much more worthwhile that people want to work less.

[00:22:17.12]

I think historically, the way technologies have affected labor supply has been underappreciated. One of the most dramatic technological changes affecting the labor market, of course, was the pill, which dramatically changed women's formal labor force participation. But that is often overlooked in the examples of ways in which technology has changed the labor market. The third area where I think there's room for more emphasis is on transition dynamics. Even if it is possible that there will be plenty of jobs, lots of new tasks that are created, short-run dislocation might still be possible because of AI, even if the long run turns out to be relatively manageable. There are lessons we can learn from history. So if we look at how much change there's been in terms of occupational turnover, In recent years, the most recent period since 2019, we have seen an acceleration in how quickly the occupational mix is changing, which is one measure for how much technological innovation and other deep factors are affecting the labor market. But there have actually been periods in history where there's been even more occupational churn—the 1910s, the 1940s and '50s. Saw much bigger dislocations than we've seen in the most recent years.

[00:23:39.19]

And other research that, again, written by folks you're about to hear from on the panel, have looked most recently at how much change there has been in the occupational mix just in the LLM era. And it looks comparable to what we saw in the PC era in the 1980s and the internet era more, more recently. I think there are other ways in which past disruption offers really important lessons for thinking about AI in the labor market. Aside from wherever we think AI might bring the labor market eventually, one is that geographic concentration matters. Some of the most difficult-to-manage disruptions have been geographically concentrated, whether you think of the mechanism of agriculture, which dramatically affected rural economies, or the China shock, where the effects of trade with China were quite concentrated in particular regions, especially in the Midwest. One difference with AI is some of the most AI-exposed jobs are more geographically dispersed. So it may be that once we start thinking more carefully about transitions and the factors that make transitions more or less manageable, We can point to lessons from history on dimensions like geographic concentration. The other point is we know from other disruptions, both recent and distant, that there's much room to improve policy for managing disruption, whether it's making health insurance more available and affordable during periods of non-employment or shoring up the unemployment system.

[00:25:17.31]

There are unsexy solutions to some of the sexy problems that AI brings that are, again, lessons from past disruption. And finally, adaptability really matters. It is not just who is most exposed to the effects of AI, but who can manage that transition. That may be something that is about an

occupation, that certain occupations that are more AI exposed also happen to have people who are more adaptable. That's what some other particularly compelling research found recently. But even within occupations, adaptability may depend on where someone lives or their demographics or what their past work experience is. Someone who is in an AI-exposed job who has worked in that field for 20 years may be less adaptable than someone else who is also in the same AI-exposed job, but has had a much more varied work history with experience adapting and transitioning. I want to end with some principles for future research around AI. Again, because the field is so new and there's so much room still to set the terms and the language for how we think about AI in the labor market, I want to end with 4 principles. The first, just to underscore the previous point, is to think comprehensively about labor market effects, not just about labor demand and which occupations may be more exposed or affected, but also some of the under-the-hood effects on productivity, impact on labor supply, and especially a focus on transition dynamics.

[00:26:59.41]

The second point is really to encourage people who are doing this research to think about how they can contribute to the collective data infrastructure. There's incredible public benefit when you come up with a taxonomy or create a new AI exposure measure to make that public and reusable. Make it easy for other researchers to merge that in with their own datasets so we can replicate and compare different AI studies. We need to make research useful for decision makers. Policymakers, educators, workers who are trying to navigate the job market all are eager, if not desperate, for insights about how AI is going to affect the critical decisions they need to make. So the more that research can be translated into plain usable language and tools that decision makers can use, the more valuable that research is. And then finally, just a plea to be broad-minded about historical analogies. There may be ways in which AI is unprecedented. The long-term effects of AI in the labor market might look unlike anything we've ever seen, but the transition dynamics will in some ways be familiar. The transition lessons may still apply even if the long-run impact differs from any kind of technological shock we've seen before.

[00:28:20.16]

Thank you for your time. So glad you were here today. I want to turn it over Um, to Ben and the panel.

[00:28:49.21] - Ben Casselman

Thank you, everybody. Thank you, Jed, for that. It was fascinating. Thank you to Brookings for hosting us here and to all three organizations that helped put this together. I usually start these

sorts of panels by sort of talking about how important this issue is and how important it is that we all pay attention to it now. And, you know, this is a really urgent issue that's not getting enough attention. I think I'm going to skip over that part here. I think there's pretty widespread understanding that we need to be paying attention to this. But I am glad to be having sort of this specific discussion with this specific group of people today, because as much as AI and the AI conversation has become ubiquitous recently, it has sometimes seemed to take place almost completely divorced from any evidence. That's not true of the people on this panel, but the debate overall seems to have been dominated to some degree by kind of multi-thousand-word works of speculative fiction rather than by rigorous academic research. And as I think we heard from Jed a moment ago, and as I think we'll probably hear at length in the panel here, there's some good reasons for that.

[00:30:02.44]

It's hard to do serious data work in an absence of data. And this is all very early, and it's all taking place, you know, very slowly. But, I mean, evidence is going to be really important here, to sort of state what I hope is the obvious. It matters whether this is a transition that plays out in very localized geographic areas, localized parts of the economy, localized demographic groups, versus if it cuts across the entire labor market. And it matters whether this plays out over 2 years or 20 years. And so we're gonna need evidence in order to follow what's actually happening here. And so what I hope here is that we can at least lay an evidence-based groundwork for what we're seeing, or if in the absence of evidence, at least to sort of agree on kind of where the areas of uncertainty are and where we ought to be looking for more evidence. So I'm going to dispense with sort of detailed bios of all, of everybody, given our times here, and you know, You all have Google or ChatGPT to fill you in on who everybody is. But just quickly moving down the panel here, Martha Gimbel is at the Budget Lab at Yale.

[00:31:08.43]

Nathan Goldschlag is at the Economic Innovation Group. And Bharat Chander is at Stanford Digital Economy Lab and the Institute for Human-Centered Artificial Intelligence. I have many, many questions, but if you all have questions, I think you all have note cards which you can fill out and pass up, and I'll try to distill those into brilliant pithy questions that I'll take credit for. To move down the panel though with sort of a first set of questions here, and Bharat, maybe we can start with you. You have done some of the work that Jed referenced that sort of looks at the early evidence of impact. And I wonder if you could describe sort of what you think we are seeing and kind of what that maybe signals about where things might sort of play out over the next little period of time. Sure.

[00:31:53.55] - Bharat Chandar

I wanna start by Saying thank you for the invitation to be on the panel with these wonderful researchers who are doing some of the highest quality work on collecting the data so that we really know what's going on in the labor market today as it relates to AI. And I think that's really important because I think this issue is becoming increasingly polarized and that's not very good for the science. So these sorts of discussions where we can all have a conversation about the data and where we still need to be doing more, I think is going to be really useful. So part of the reason that I think there's a lot of uncertainty here, and maybe that's also feeding some of the polarization about the discussion around this topic, is it's really hard to estimate the effect of AI on the economy. So we don't have a world where we can see a world without AI and compare it to a world with AI and compare the outcomes there. So we have to make the best that we can with the data that we have available and the variation that we have available to try to tease apart different explanations.

[00:32:53.58]

I do think that we probably have a lot of common ground. All three of us have papers that basically show that the aggregate effect of AI on the labor market right now is quite small. So, you know, we are not seeing widespread job displacement from AI. I think where the uncertainty lies is on specific groups of workers. So there are some work— there's some papers that look at online labor markets and what's happening there. They potentially find some disruption. But I think a lot of the discussion is focused on entry-level workers and specific jobs that are more exposed, such as software development, more administrative roles, data entry, customer service, things like that. In our work, we found our preferred estimate is we find a decline in employment for those types of roles. So there's a relative decline in the most exposed entry-level jobs compared to the least exposed entry-level jobs. That really starts to decline around the beginning of 2025. And it's been increasing since then. So since we first started collecting the data, it's now 16%. That's a 5%age point increase then compared to when we started in July. So I think in the context of the broader literature, all the papers that are out there, my sense is that the aggregate effect of AI on the labor market is quite small.

[00:34:12.00]

There might be some impact on entry-level workers. I would say there's something like a 70 to 80% chance, in my view, that it's having some appreciable effect on the entry-level labor market and some of these specific jobs. I think a lot of the uncertainty around that is specifically around, you know, can we actually estimate the effect of AI, or are there other things in the economy that are going on that are also, for whatever reason, affecting these workers? Jed suggested teleworkability is one example of that. We tried to test for that in the paper, but There's no

perfect way to do that. So that's, that's kind of my assessment of where things are right now. Now, going forward, I think as more data comes in, that would lead me to update my view about how much of this is being driven by AI. We don't have any ambiguity about the fact that COVID affected the labor market because it was such a huge shock. So I think it's important that we continue to collect this data and track it over time and see if these changes actually start to continue to accelerate the way that is suggested in our early research, but I think there's still a lot of uncertainty about what's gonna happen going forward.

[00:35:15.49] - Ben Casselman

Thanks so much. Nathan, you have been publicly a little bit more skeptical that we're seeing much of the impact so far, and so I wonder sort of what you think we're seeing, and kind of where you agree and disagree with what Bharat just had to say.

[00:35:28.07] - Nathan Goldschlag

So Bharat, I think, So I'm also grateful to be here. It's a great group, really fun discussion. It's fun to be part of this conversation because it's just moving so fast and there's so much to react to and it's just so interesting. So I agree on the point that we generally, we do agree. I mean, so the studies are consistent in finding that there's no aggregate impacts, sort of. And, you know, the team at the Digital Lab and Brett's team they— one of the things that they bring to the table is much bigger data, right? So the Current Population Survey is a relatively small survey. It's really hard to focus in on really narrow— narrowly defined subsets of the data to find these sort of effects. But the data that they brought to the table, you're able to do that. So that was a big contribution. I think that part of this is a little bit related to, I think, a misguided sense in the media that there's this, you know, job apocalypse among young college grads. So a lot of this comes through these— this data from the New York Fed. It's aggregations from the ACS data showing like 7% unemployment rates, these giant unemployment rates for recent graduates.

[00:36:41.09]

But when you look at that data, you know, my former colleague Connor O'Brien has done some analysis of these things. It's very noisy, just like the CPS, it was very noisy. And so, you know, we're not sure if the unemployment rate for recent physics grads is 4% or 11%. So that's either the top of the labor market or the bottom, right? So that's a big gap. But when you zoom in on just in the general question of what is the labor market pattern we're trying to explain, you know, if you think about, okay, so these young college grads and young non-college grads, you sort of see the unemployment rate for young college graduates sort of inching up a little bit quicker than the unemployment rate for the non-college grads. When you zoom out to think about those

that are not participating in the labor force, think about something like an employment rate, you get, you get a slightly different story where young non-college individuals are exiting the labor market because they can't actually, uh, find any meaningful employment. So they're not, they're not remaining in, in the searching mode. So the unemployment rate as a measure requires that you either be employed or be looking.

[00:37:45.28]

So if you stop looking, then you fall out of that measure. And so, you know, when you sort of zoom out a little bit and think about the broader health of the young labor market, you have to have something that's consistent with both the college and non-college, which isn't on its face, obviously, an AI story, right? Because we know that from the exposure measures, for whatever they do or don't capture, they sort of correlate with BA, bachelor in skills. So like higher skilled workers tend to be more exposed. And I think even in the fantastic appendices that Brett's team has put together, you sort of see a little bit more of explanatory power among the occupations that tend to have lower education levels, which again is sort of not obviously consistent with that. So I think in broad sense we agree. And I don't think he would disagree with this framing it this way either, that, you know, so even if we do sort of march forward in time, we don't see any appreciable measurable difference between the exposed and the unexposed up through 2024, say early 2024, and then it starts to show up.

[00:38:54.41]

The difference, I think, with COVID is that COVID sort of march happened and then the world was different. Right. And all of those things manifested quickly. But now, if you're sort of differencing from, you know, November of '22, you know, we're sort of— lots of things have changed between then and now. And so, fundamentally, I think I totally take the point that these technologies are probably going to have really important impacts on the labor market. It's just like, are the tools that we're using and the ways that we're using the data up to the task? I'm not so sure.

[00:39:30.07] - Ben Casselman

So Martha, I guess my question for you at the start here is why is this so hard? Like, I think that there's probably a fair amount of appreciation for why it's difficult to know where things are going. Yeah. But the fact that it's so difficult to even establish where we are now, I think would be maybe surprising to many people. And so I'm curious, sort of what is it that's difficult to measure here and why is this so difficult?

[00:39:57.29] - Martha Gimbel

Yeah. I mean, I think part of it is there's such a difference in adoption, even amongst people who, like, quote unquote, look similar, right? And I think people have probably experienced this in their daily lives, right? There's economists who are using a ton of AI very quickly. There's economists who are following along. There's economists who haven't even started yet. And so you're then being stuck with data which is categorizing people into these, you know, sort of very broad groups when that really may not reflect the sort of beginnings of technology where there's still so much difference across people and across firms. I think the other thing is that in some ways we are expecting both too much of the technology and of the data, right? It would be astonishing for a technology that only emerged a couple of years ago and for many people only started being really, really useful, let's call it in the last, you know, year, right? Not that it wasn't useful before, but that, you know, people really started using it at work to be having mass labor market effects, right? That would just be unusual. It takes time. It takes time for firms to adopt.

[00:41:14.18]

It takes adapt, it takes time for workers to adapt, it takes time for customers to adapt. Things just don't happen instantaneously. And so part of it is, right, that we're trying to— I don't want to say torture the data because that makes it sound like we're all committing some sort of econometric sin and I don't think we are. But trying to wrest some signal from the data that's not perfectly set up for this question. But also in some ways we're asking too much of what technological change looks like to have a definitive answer of like, this has happened already and therefore we are good to go.

[00:41:48.04] - Ben Casselman

But let me, let me live with that, that point for a moment here, because I think that if you were to follow some of the popular press coverage, maybe including in my august publication, uh, you would get, uh, I don't mean that ironically. It is an august publication. Thank you. Um, but I think you would sometimes get the impression that a lot is happening very, very quickly. We see public layoff announcements that get blamed on AI. We see, uh, claims made by, you know, leaders of the AI companies that say there are sort of huge impacts. And I, I could imagine some people in the audience here sort of being skeptical and saying, well, maybe your measures aren't picking up on it yet, but that doesn't necessarily mean that there aren't huge disruptions happening here. Like, to what extent do you think we are just failing to pick up on something that is really happening versus this is actually playing out more slowly than maybe some of that coverage suggests.

[00:42:44.01] - Martha Gimbel

I do think if you were seeing like a tsunami that we would be able to pick it up. I do want to focus in on one thing that you said about, you know, these kind of CEO statements. I do feel really—the CEO statements have driven a lot of this coverage, and I do feel really strongly that if you're looking at ways to track the the impacts of AI, CEO statements are—in the labor market. CEO statements are possibly the worst way to do it for a couple of reasons, right? One, there's selection effects in who issues a statement when there are, for instance, layoffs or hiring, right? Your hardware store down the street does not issue a statement when they hire or fire. There are selection effects in who gets picked up, right? If you say, I'm laying off 50% of my workforce, or I think someone said 40% recently, due to AI, everyone's gonna really jump on that. That's a huge story. Other people may not be making that statement, that's not covered, right? There's an old cabaret song of like, this many nuns did not have this thing happen to them today, making fun of headlines in august publications.

[00:43:52.54]

And that's true, right? If something doesn't happen, It's not covered in the same way. And then also, of course, there's incentives for CEOs, right? The CEO is not going to get there, up there and go, listen, I have badly bungled the macroeconomic situation for the last couple of years and now we're going to fire people. But don't worry about it. You all should keep giving me your money and it's going to be totally fine moving forward. Like the corporate communications team would like, I can only imagine, would just throw their bodies in front of that, right? You're going to say, You know, like, look, this is, you know, we're making productivity enhancements for the future, et cetera. The other thing is I think it's hard for people to know. You know, I often say that hiring at Budget Lab, we're small, has not been impacted by AI. We're just, you know, having people do more with what it is. I don't know that, right? Like, I can't go back in time and look at the version of Budget Lab with AI and without AI and how many more junior staffers are How many fewer junior staffers did we hire?

[00:44:50.07]

And so I think it's really hard for people to tell. And I do think this technology feels very personal and very scary to people. And so there's a real rush to the sort of apocalyptic version as opposed to the kind of boring data like, well, within these standard errors, we see these effects.

[00:45:13.23] - Nathan Goldschlag

Kind of thing. The other point to make about those announcements is that gross flows are enormous. You know, so the economy, every few months you add 100,000 jobs, you delete

100,000 jobs. Underneath that are flows of 7, 7.5 million, right? So, you know, when you read some headline, even though it might be in the tens of thousands, in terms of gross flows and the sort of things that we're thinking about in the aggregate effects, is just not going to move the needle.

[00:45:43.45] - Martha Gimbel

I mean, as an example of this, Meta laid off 600 AI researchers like 6 months ago or something. They laid off AI researchers so they could hire different types of AI researchers. And it was 600 people. And to be clear, this is very, very bad for those 600 people. I'm not trying to minimize that. Although maybe they found jobs quite quickly. I'm looking at Bharat, who is more plugged into Silicon Valley hiring than I am. But like in the context of what we had 1.7 million layoffs last month, like if that doesn't measure. And the headlines were like, even Meta is doing AI-related layoffs. It's like, that's not what this was. I just think people are really primed to expect the apocalypse here.

[00:46:27.09] - Bharat Chandar

Yeah, I largely very much agree with what's being said here. I just don't take the layoff announcements as very informative. I think it's way more informative to look at the data. So I'm with Martha and Nathan on this one. I will say, I don't think we have to be in this binary where it's either all due to AI or like not due to AI at all. So we should have some amount of humility about how we interpret those statements. But I also think they're mostly uninformative about what's happening in the broader labor market.

[00:46:56.19] - Ben Casselman

What measures should we be watching here that if we all gather again in a year from now,—and we will presumably know somewhat more or in a couple of years. What are, what are the measures that are going to help us distinguish that over the course of, of the next year or 2 years? And Bharat, if you want to start on that.

[00:47:13.08] - Bharat Chandar

Yeah, I mean, the thing that we're building out, I think we're going to launch this in April, but we want to build these trackers where we can follow employment over time for different age groups and different occupations. And I think we want to track two things like one, are there— is there reduced hiring into certain jobs? But also, and this is what we've seen with all prior technologies, are people shifting their employment to other jobs instead? And I think that's what we want to be tracking and following the employment series over time, both at the aggregate

level, but then at the more granular level about how are people adjusting to this technology is, I think, going to be really important. So that's what I would want to be tracking. I just want to be following the data over time. Are we going to end up in a world where people are going to have a hard time transitioning to other work? Or are we going to be able to see people actually going into other jobs where they can find good employment, good wages, etc.? I think there's a lot of uncertainty around that. But I think that's among the things that we want to be tracking, both in both at the granular level using things like ADP, but also in the aggregate data that comes out from the government.

[00:48:16.55] - Nathan Goldschlag

So I agree with all of that. I might be a little a little less positive about ADP as the solution. I think there's a lot of things that you want to get out of the federal statistical system because it's a nationally representative frame. And so I think, you know, so there's two layers to this. One is I agree 100%. We want to compile lots of suggestive evidence about the time series of what's happening in the different cuts of the data. And to the extent that these occupation measures are informative, we're going to sort of see those. For sort of people in those sort of certain types of highly exposed occupations start to switch around to different jobs. Like, those are all well and good, and I think you could do much better, I think, by using some administrative data at the Census Bureau and within the federal statistical system to get, you know, the granular data that something like the Current Population Survey does not get us. So between now and then, I would hope that we would make additional investments in the federal statistical system to get the granularity we need using administrative records that we have.

[00:49:19.47]

So that's sort of like the acquisition, you know, sort of piling up all of the suggestive evidence. I think there's a second track that I hope Bharat and your team will be engaged in, which is the more serious econometric work. And that, again, you're using microdata within the federal statistical system where you observe adoption at the firm level, and you can sort of, sort of, you know, work through and control for all the things that Jed pointed out and sort of do much more detailed analyses and get a whole lot closer to, you know, something that we all feel comfortable about, a causal estimate, right? So that's a separate track. It's going to be much more boring, but I think it's also going to be really informative. But then at the same time, absolutely, we sort of accumulate all those things that are, you know, on its face, you know, we have a pile of evidence that's either consistent with or not consistent with mass changes in the labor market.

[00:50:12.06] - Bharat Chandar

If I could just add one quick thing to that. So totally agree. We need to build out that infrastructure in the census. And I also think arguably we should be doing both things. Like, it shouldn't just be people like me who have access to ADP. So insofar as we can make some kind of like private-public partnership there, that would also be great.

[00:50:29.23] - Ben Casselman

If you're playing the bingo game along here, we've had a call for funding for the statistical system and a cabaret reference. Martha, I wanted to ask specifically about some of these exposure measures, because a lot of this, a lot of the analysis here is based on these measures of AI exposure. Jed referenced that some of these tell different stories. You've done some work on this.

[00:50:50.14] - Martha Gimbel

Yeah, I mean, I want to be clear, we've all used the exposure metrics. I think we're all sort of aware of the various issues here. I think there's a couple of things to keep in mind with the exposure metrics. Different people have come up with different measures of exposure to AI. Two things. One is sometimes that gets turned into measures of probability of automation. Your job is exposed, therefore you are going to lose your job. That is not what the exposure metrics are supposed to be looking at. It's supposed to be like, will your job, how you do your job, be affected in some way? Way. And so I think there's been, you know, some at times miscommunication on that issue. The other thing to keep in mind, right, and we did some work on this at the Budget Lab, is the exposure metrics don't all agree with each other. And I will shout out Jed for originally sending us down this research path, but, you know, on the aggregate, yes, there is agreement, but a lot of it is driven by everyone agrees that HVAC technicians are not exposed. Everyone agrees that plumbers are not exposed.

[00:51:59.19]

But as you move to higher and higher levels of exposure, there's more and more areas of disagreement. Because this is hard, right? You're trying to guess what the path of technology is and people are looking at different types of technology. And so I think we need to be really humble about what the exposure measures can actually tell us. To go back to your earlier question about like what are you looking at? You know, I think it's really important, one, for us not to be— decide ahead of time where the impacts are going to be and make sure that we're looking more broadly. If we're just looking at young workers or we're just looking at software engineers, we may miss something that pops up someplace else. But I am also particularly interested in some of these measures of occupational change that are not tied to a specific

exposure metric or a specific usage metric, they won't necessarily tell us like, this is AI. It'll simply tell us something's changing quickly in the labor market or abnormally in the labor market, and we should dig into that. But I think it's important to kind of have these alarm systems, kind of, for lack of better phrasing, set up that just says, something has tripped the alarm and we should dig into this a little bit more.

[00:53:14.15] - Ben Casselman

But do you want to jump in on that? Because I know that you've found that these exposure measures, as much as they may disagree to some degree, do have some predictive value that you're—

[00:53:21.26] - Bharat Chandar

Yeah. So a couple of thoughts about this. One is that it's true that some of the exposure measures don't exactly line up, but some of that is by construction. So, for example, in Jed's blog post, I think he references one from Felton et al. 2021, but that's like a pre-LM AI exposure measure. So it makes sense that it doesn't have the same— it doesn't give the same predictions as what we would get in a post-LLM world. Number 2 is that— so my perspective on this is that we should absolutely look at the different exposure measures and how they correlate with things like demand, in part because we have predictions about where we're going to see impacts, and we should see if those predictions actually bear out or not. Now, if these exposure measures were completely useless, then they shouldn't predict anything. Now, they might predict things even if it's not a causal relationship, but if they were completely useless, they wouldn't predict anything. So in that sense, insofar as they're predictive of things, they might be useful. They might not be perfect at predicting those things, and they might push us towards not seeing effects when there might actually be effects there.

[00:54:31.49]

But in that sense, I actually view them as pretty useful, actually, insofar as they can actually predict real-world economic outcomes.

[00:54:43.27] - Ben Casselman

To what extent— we know with the disruptions, whatever disruptions we see, we expect to see some level of job loss in specific places over time. I don't think anybody on this panel is disagreeing with that idea. We would also expect to see some level of job creation, new jobs that come out of this. New firms that come out of this. To what extent are we measuring that now? How can we go about measuring the creation side of the creative destruction here?

[00:55:10.33] - Nathan Goldschlag

So I spent a number of years working on sort of measuring business dynamism. So I have very strong opinions about this. So there's a couple, there's a couple different sources. So we sort of have the gold standard, really slow data, the Business Dynamics Statistics. These sort of tell you new firm starts, job creation, job flows. So it's the flows of jobs and firms and establishments, establishments being physical locations of a business and a firm being, you know, some entity that owns more than one establishment. So you can sort of catch all those things. We have sort of higher frequency measures of that, the business formation statistics. So there, you know, so we're sort of taking these applications to start a business. So in order to start a business, You pay taxes, you file for an employer identification number. So, you know, there's a program that takes all those applications, tries to guess which ones are likely to become employers, right? That's a much faster statistic than when it hits all of the business frame data. And so there, you sort of, you know, something that the entrepreneurship literature has been chewing on for several years is that, you know, for 40, you know, 35 years, the percent of new businesses has been falling, startup rates have been falling.

[00:56:20.56]

Pretty consistently. And then post-COVID, there was this boom in entrepreneurship. And it was like, well, is this real? Is this a bunch of Etsy things because of COVID Turns out it was real. It was sustained. That measure was picking up real activity. So there's a research team at the Census that sort of took the applications and some of the write-in data and went looking for AI in there, right? And so you can use— and this is just a plug for the idea of using administrative records in a creative way to get at sort of universe-level frame statistics. And so for that, you could sort of pick out, you know, what do we see? We see, and this is, you know, I think it ran through 2024. I expect that this would, if they updated this analysis, it would continue to go up. But yes, there's a huge surge in business formations associated with AI businesses. And if you go pick up the data that came out just a week or two ago, right, you sort of saw throughout 2025, is sort of another increase. So, you know, if you— if I'm drawing, I don't have a slide, but if I'm drawing with my hand, you sort of, you sort of have this pre-COVID and then a post-COVID bump, and it remains there high, right?

[00:57:24.46]

And then in the past year or so, another surge, like a second surge. And if you look at the industries that are sort of surging along with it, sort of professional business services, the information sector, right? These are noisy measures, Right? But they are— they're a weak

signal, but they are a signal of what I'd expect to see going forward. And so, yes, I think the one thing I think we all would agree on is that we're in for a period of lots of change, right? The question is, like, how fast is that going to happen? How much, you know, is the labor market going to be able to adjust? Are people going to have to flow through unemployment? How long are they going to spend in unemployment? You know, geographic distribution of these shocks. You know, all of that, well and taken. This could be a period of change, and a period of change, I think, you know, the first thing that comes to mind is dynamism, and I think a lot of that's going to be facilitated by entrepreneurship.

[00:58:18.45] - Ben Casselman

There's been a lot of understandable focus on young workers, entry-level workers. Young workers also have more time in their careers to adjust, which is a positive. I have a— looking through these questions, there are a lot of questions about older workers. I can't— I will make no assumptions about why those questions have come. But I mean, to what extent are we starting to see, or should we at least be concerned about this sort of moving up the hiring ladder and starting to play out in groups that may have less time and ability to adapt? Bharat, I don't know if you—

[00:58:55.13] - Bharat Chandar

Yeah, so we're not seeing that so far in the ADP data or in the CPS. So, I don't want to make very strong predictions about the future because I think the focus of this panel is more about just understanding the data as it is right now. We can also have a discussion about future impacts, but we're not seeing it right now.

[00:59:18.26] - Ben Casselman

What about margins other than employment? You know, when I talk to young people today, a lot of what I hear about is not just about finding a job versus no job, right? It's about finding a good job, a job that they thought they were going to get when they graduated. I also hear a lot of concerns about wages, of course. To what extent are there margins that we can be looking at here other than just that employment/non-employment binary?

[00:59:44.47] - Nathan Goldschlag

I think there's some really interesting ones. I don't have strong predictions about how— what I expect the magnitude of the change along these margins. I think one thing that Brett's team found was that there was no wage effects, right? So we sort of had headcount effects, but but no wage effects yet. That's not to say that we won't, right? Because if there are productivity

effects and it changes the marginal productivity of workers in a certain sort of way, you can—that'll be reflected in wages. But I think one interesting, you know, sort of subthread of the economics literature on what AI is going to do is the theory of the firm, right? What is a firm? Why do we have firms? Well, you know, so it's to lower the transaction costs. So, one of the things that AI might do is to lower transaction costs, so you might see more self-employment. So, instead of the traditional job, you might see an increase in self-employment and sort of contractor relationships between individuals. I don't know that that's going to happen. I think it's very interesting. And, you know, we do have a statistical infrastructure that might be able to pick that up.

[01:00:46.55] - Ben Casselman

Also, not something that our safety net system does a very good job of dealing with, of course.

[01:00:51.07] - Bharat Chandar

Brett, did you want to— Yeah, so just— Just one small thing. We did update the data through December recently, so we're hoping to post that soon. So there's maybe a little bit of a wage difference now, but for entry-level workers, but I'm not sure we would say statistically that different yet.

[01:01:05.44] - Ben Casselman

So we're basically out of time here, but I want to just quickly run down the panel with two questions, one of which came in here, which I think is fun, which is, did you or how did you use AI in preparation for your comments today? But then also, if you had a, a one, one dream data set or piece of data that we could be tracking here that will be useful going forward, what would that be?

[01:01:29.40] - Martha Gimbel

Martha, why don't you go? I did not use AI, which maybe I should have. I am testifying before Congress tomorrow. Everyone should tune in. And I did use it for that to plug my testimony in and to ask it like, if you wanted to ask me mean questions, what are the questions that you would ask me? And I would like to say for the record, it was really mean. So that was useful. Dream data, I mean, there's so much of it. For me, I would really love more direct connection between some of the firm-level data and the employment data so we can kind of see more directly usage, et cetera, at the firm level and how that's flowing through to employment headcount. Et cetera.

[01:02:15.08] - Nathan Goldschlag

So my comments are also boutique, artisanal. No, no AI in these. No M-dashes. Most of my code is not artisanal anymore. It's a lot of AI. So yeah, no M-dashes. So in terms of the ideal measurement, two things come to mind. I think linking the firm-level use measures at the micro level to the employee-employer data is important. I think the second thing, and I think this would, you know, a lot of people come to me, they say, well, how do we measure prevailing wages better? Any lots of different questions. I give the same answer and I'll give it here is that we need to get occupation on the employee-employer data. Right. So, you know, we have employee-employer linked data, millions and millions of jobs. Right. But we don't have occupation on that. And so something like ADP, the value of it is it's employee-employer data. With occupation on it. The downside is it's not nationally representative, doesn't pick up the young firms very well. You know, there's some downsides to it. You know, it's hard to measure individuals longitudinally, I think. I don't think you get the full employment history of an individual over time.

[01:03:21.25]

But that's a solvable problem within the federal statistical system. There's a couple different ways we can accomplish this, but suffice it to say, if you got occupation on something like, you know, the employee-employer data, you could have detailed— because the data is thick enough, you could have really detailed public tabulations of what's happening in really fine segments of the labor market?

[01:03:41.09] - Bharat Chandar

Unless you count Google searches as using AI, I actually also didn't use AI. Despite the fact that I try to use AI as much as I can, I don't use it for a lot of writing just because I want to make sure that I actually know what I'm talking about. And yeah, that data that we should be tracking. I think another thing is we really don't have a good idea of what's going on in the rest of the world. So I'm working on a project that's trying to collect labor market data from other countries and also collect some new AI adoption data as well. But it would be really useful if we knew what was going on outside the US as well.

[01:04:14.56] - Ben Casselman

Well, thank you all so much for being here. Thank you to our panel and to Jed as well, and to Peter, who has taken off. But thank you all so much for coming.