



General Discussion: Technology and the Labor Market

Chair: Janice Eberly

Chad Syverson: Thanks to all the panelists for those interesting presentations. I don't have enough time to ask all the questions I want to ask, so I'll ask two quick ones.

First, I'm not going to argue with the idea of the importance of pure research in the university model. I totally agree that theory and empirics tells us that's important. I also wouldn't disagree with the notion that professors should be paid more, of course. But I'm not so sure the particular example of the AI industry pulling people out of academia is clearly an inefficient outcome. It might be efficient, especially temporarily. It's quite possible, I think, that researchers' marginal products could be higher in the industry now. The industry players have complementary inputs that the universities don't have. I wonder what you think about that possibility.

And second, Laura, I totally agree with the notion that who gets the surplus is really key. As technology changes, we start thinking of data as a new type of capital. Usually, when we think of traditional capital stock, that surplus goes to the owners of the firm. Do you think that's true dollar for dollar with data capital,

General Discussion: Technology and the Labor Market

2

too, or is it more tied to the worker? I think it's more tied to those who work with the data. I got the impression you were trying to impute the size of the data capital stock by using the wage bill paid to people working with the data. It seems there's a natural connection there: some of the workers have human capital that's complementary to data. As a result, they're going to get some of the surplus from the data, not just the stockholders. I wonder about your thoughts on that, too. Thank you.

Şebnem Kalemli-Özcan: Thank you so much for insightful comments. And Chad asked my question to Laura, so I'm not going to ask that.

So, my only question is to Ufuk. Ufuk, you painted a very depressing picture, right? The talent goes away from university to industry. And even in the industry, there is this large-takes-all issue, right? So, the large firms take the R&D, small firms are disappearing, they are not even starting. So, you have a story of financial constraints there. You showed us for very small firms (that Europeans call microenterprises) with employment one to nine. These guys were going on their business with a lot of debt. And then the Federal Reserve starts its tightening cycle and they start being delinquent and default. I mean, this may sound like monetary policy is creating this. But an alternative interpretation that I like is actually, monetary policy is helpful with creative destruction here, right?

Why did these firms start at all? Finance 101: they shouldn't have started with debt, right? You start with equity, especially if you're a transformative entrepreneur, which is what we want for growth. So, they shouldn't be there in the first place. So, my question to you is, what is the role for policy? Do we really worry about small firms not even starting? Or are funds not allocated properly? If R&D is not allocated properly, the same with the

funds. So, what is the role for finance? So, is this a question for antitrust or is it a question for bank regulation?

In our work, we are matching the Federal Reserve bank lending data to Census data, where we can look at all small-medium enterprises, which are the firms defined as less than 500 employees by the Small Business Administration. And there, actually, the picture is much more complicated. So, my question to you is, how do we see this from these other policies, right? Better allocation of funds to productive firms?

Ufuk Akcigit: Thanks for the excellent questions. Chad, my concern isn't simply that industry salaries are rising quickly and creating a pull factor. My worry is that we are not having enough discussion about how universities can keep up with the speed of change. By the time we realize the full implications, it may be too late. Universities are slow-moving institutions.

That raises practical questions: Should we make it easier for researchers to work in multiple settings—joint appointments, consulting, and flexible affiliations—as well as better research infrastructure? What kinds of research freedoms and opportunities should we offer to retain talent? These are fundamental issues.

I'd also emphasize something else. When the same scientists move to industry, their academic publications tend to drop substantially, and they start patenting more. That implies a shift from an open-science model toward a more proprietary knowledge model, which matters for spillovers.

Consider the arc of someone like Geoffrey Hinton: a large part of the field grew through open science and cumulative progress. In industry, speech and disclosure can be more constrained—Hinton himself has described leaving in part to speak more freely.

General Discussion: Technology and the Labor Market

4

The broader point is that open science and spillovers are core public goods, and universities are central to producing them.

That's where taxpayer support is most defensible: basic research with broad social returns that are different from the incentives inside private firms. Of course, universities also have inefficiencies—no one denies that—but the social contribution of university research is not the same as private R&D.

Sebnem, I also agree with you that crises can create opportunities for creative destruction. Before COVID, we worried that low interest rates were propping up zombie firms and muting churn.

What I'm highlighting is unequal access to resources. I'm not claiming every small business is efficient; rather, the right answer is that we often don't know, because we haven't studied small businesses with enough detail.

We are now using granular data (for example, supervisory and linked administrative sources) to understand how small firms finance their growth. The high-potential young firms are typically the most reliant on external funding. So when access to financing tightens at exactly the moment they are trying to scale, the burden can fall disproportionately on the higher-type firms—depending on the mechanism you estimate.

That's why the discussion shouldn't be framed as "small versus large." There are small firms that are one year old and small firms that are 20 years old. The 20-year-old small firms are unlikely to grow, but young firms can. Our models and policy frameworks need to incorporate that heterogeneity—especially when thinking about financial frictions. Thank you.

Laura Veldkamp: I'll tackle Chad's question of how do we measure data and where does the value come from?

So, in the interest of time, I gave you just the hint of what this does. Specifically, we examine two classes of workers: data managers, who possess database skills but typically lack domain knowledge, and analysts. And then there are people who work with structured data. We call those analysts. So, you know, those are people who in the context of finance know a bunch about the finance industry in addition to having data skills. So, think of this as a data value chain. We go from raw data to structured data, structured data to knowledge and so forth. So, the data management hires help inform us about what's the size of the structured data set. Consequently, we are identifying the exponent primarily based on the data-to-labor ratio in AI versus non-AI firms.

Now, you hinted at a broader question of where does the value of data come from? And I'd say, just like having more workers makes the marginal product of capital high, having more data workers makes the marginal product of data high. However, data is generated in a fundamentally different way than traditional capital. So that's entirely separate. That's the last part of the discussion, which says that fundamentally the data comes from our transactions. So, it's the consumers that are producers of data. And that are in effect selling it to firms. Currently, we essentially coerce consumers into selling their data because they lack the option to obtain essential goods and services without also engaging in the data sale. So ultimately, the value does come from consumers because they're the generators of most data.

Tiff Macklem: Three excellent presentations. Thank you. I learned a lot. A question for Larry and then a bit of an unfair question for Laura.

Larry, you went to tremendous pains and efforts, very impressive, looking at the historical time series for the United

General Discussion: Technology and the Labor Market

6

States. I'd be interested to hear a little more about the cross-country evidence. Obviously, I know something about the Canadian data. In Canada, post-secondary education is cheaper than in the United States. We have higher rates of post-secondary education, and we have a lower wage premium and less inequality. So, at that level, it looks like it works. But there are some things we wonder about. For example, we have fewer people with advanced degrees in Canada than in the United States, and we kind of wonder if that is because the wage premium is lower so it's not worth it? Are we losing out? So anyways, interested in the cross-country evidence.

And Laura, my slightly unfair question, because I don't think there's much data, but I think there's a lot of data that you can speak to. I mean, one of the things we're hearing about in AI lately is it's shifting from machine learning to reinforcement learning. You put a lot of emphasis on data, the value of the data, the data-labor ratio. Reinforcement learning doesn't need as much data. So, is there any evidence that the return to data starts going down? You start to hear that they're kind of reaching the limits of the data. And I wonder, is it going to be all the same? Is it all about data in the future? Or is the AI itself going to become more important?

Ayşegül Şahin: I have a quick question about the fascinating data from Germany. You find that firms are actually retraining workers that are exposed to AI, but this could be actually specific to Germany. Is it possible to look at other episodes and what happened to typists or what happened to phone operators in Germany versus in the U.S.? Because in the U.S., we spend much less effort in skilled redevelopment, so maybe this is a more optimistic view because we are looking at Germany.

Larry Katz: That's a very good question. The U.S. and Canada is actually an excellent example. And we did a book on this in the 1990s comparing the U.S. and Canada. And you essentially see similar changes in technology and labor demand, at least at the college and non-college. And Canada kept up the supply growth of education and created much greater access. There's a lot of work people like Phil Oriopoulos have done in Canada on access to college. And essentially you did not see the same growth in inequality. I think once you start thinking about the post-graduate market, North America becomes a much more integrated market. And obviously top Canadians have labor market opportunities at the scale if one includes U.S. market. If you look more broadly at the cross-country patterns, you find that countries that did not have the sort of slowdown in educational growth that the U.S. did, did not see nearly as large an increase in inequality in the 80s, 90s, and 2000s. The data going into the 2010s and 20s look a little bit different but have a similar flavor.

There's an interesting question that if you break it up, the return to post-college training is what's really gone up in the 21st century and increased dramatically. And if you look at the U.S., we've done a great job of expanding college in some sense. On the cheap, middle quality and lower quality schools. So, the number of slots in elite state universities and private schools have basically stayed the same over the last 30 or 40 years, even though the population and the international student flow have grown so the share of Americans in a cohort who go to these top universities had declined. But we flooded the market with less selective colleges.

So, I think of the college supply as having these two groups. There's some drop in demand in the middle, and there's been a huge increase in supply. So, you start seeing falling wages at the

General Discussion: Technology and the Labor Market

8

bottom of the college group. Whereas we have a lot of great data showing that the person on the margin who gets into a state research university still gets a tremendous rate of return. Demand remains very high for top-end colleges' skills. And we've done much worse in the last several decades at expanding access to those skills than we did in the mid-20th century.

Laura Veldkamp: I'll start with the question, unfair question for our Canadian friend. So, at some level, I agree 100 percent. As consumers, we have missed out on a significant amount of value because the initial data points had the highest value. But what's done is done. Now, does that mean there'll be no value left? Absolutely not. Why? Because the world changes, right? So, if the world stayed exactly as it was and there were no new products and no changes in trends or patterns or anything else, then, yeah, we will have missed most of the value and the value of data goes to zero. But every time a firm wants to introduce a new product, the marginal value of data goes back up again, right? Because we don't have data on that new product. It's new. Or every time we get some new regime or new habits or new fads or new trends.

Ayşegül's question. Yes. So, we have questions specifically on retraining, and we're going to look at that. And I agree that that may not be the same between Germany and the U.S. That's not actually the piece of it that I'm most interested in. The most interesting and potentially portable insight is that firms treat labor as a resource allocation problem under uncertainty—essentially a portfolio management issue. I'm going to need some workers. How many workers? How many workers will I need, you know, doing this function, producing this product at this plant and so forth? And AI is really good at prediction for those uncertain needs. So, it might help firms achieve a more efficient allocation of workers and adapt better to coming shocks. And those are two

very different reasons why firms reallocate workers. And so, part of the challenge is to tease the two apart. But I agree with you. You know, the retraining bit could be quite country-specific, but we have the data to tease that part out to break it out separately. Thank you.

Alan Blinder: I want to just put on the table what I take to be the big macro question and the big micro question. I get asked all the time, as probably almost everyone in this room does, is AI going to raise productivity? And I say, sure. Then the next question is, how much? So I'd like to hear three people that know a lot more about that than I do. What I say is it is anybody's guess, but I'm looking for three better-informed guesses.

The second question is the micro question of regulation. Laura, at the end of your talk, when you were talking about making sure consumers have options to protect their data, et cetera, in my head that said, to get that done, there has to be a lot of regulation and maybe heavy-handed regulation, which goes against a lot of trends. That's why I'm asking the question.

Chad Jones: So, I think the take that the three of you had on AI in the labor market was less pessimistic than the take one sometimes hears. And I think I probably agree with you, but I wanted to try to push the pessimistic argument to get your response to that. So, think about how much AI has improved over the last decade, and then imagine those same kind of improvements over the next 10 years. What would the labor market look like in that case?

First, there is an argument I hear from young people going into the labor market. It seems like programming jobs may not be the same as the past, and that's a big part of STEM jobs. Also, people going into consulting or professional services in entry level positions. AI models are already good at putting together slide

General Discussion: Technology and the Labor Market

10

decks, the models in 10 years seem like they'll be even better. So, does that put pressure on an important part of the labor market? Another example that is very prominent in San Francisco is Uber drivers. So, you can already ride around in Waymos and they're great, and they would be taking over the market if they were in greater supply, and you know, I think it's only a matter of time. Those jobs are probably going away.

So, I think I another version of this is, if we if we looked at agriculture in, you know, in 1900 in the United States, we would say, tractors are coming, they're going to replace farmers . They did replace farmers: the share of employment in in agriculture plummeted, but there was manufacturing. If AI is able to automate most cognitive work and if robots keep getting better, is there a manufacturing? Is there something else there for all of these workers? So, just curious about your thoughts.

Larry Katz: So, I think what we're seeing, Chad, is that with AI there are certain tasks, it's clearly automating, and there are certain jobs—certainly autonomous vehicles could be a huge impact on number of drivers—although things get more complicated. Think about autonomous long-haul trucks. You could eliminate long-haul trucking jobs for humans with the right infrastructure, but you would also then increase commerce across geographic areas, creating all sorts of new industries, potentially increasing the need for short-haul trucking because of greater commerce between places, which it appears in the short run is more difficult for AI to do, and we would then need more short-haul drivers. The net effect's not totally clear. There are new sorts of tasks and new industries being created as you create commerce for food products that may not have been able to be traded. It will clearly change the composition of demand and it's certainly going to displace workers in certain traditional jobs.

But I think the worst-case scenario is not that AI is super productive because then I think it's going to create all sorts of new jobs and industries. It's what Daron and Pascual Restrepo have called so-so automation. If AI is just a little bit better than workers and cheaper, it displaces a whole lot of them, but it doesn't create a lot of broad productivity benefits. And that's where I think you get dramatically poor results. Some automation technologies from the 1980s to the present have hollowed out a lot of jobs with only modest improvements in productivity. But if AI ends up really being bang, bang revolutionary in terms of taking a whole bunch of tasks and making them much more productive, and this is what we see in some settings, such as customer service agents and software coders, then it's eliminating certain types of jobs, but it's creating, as in other ones, jobs that combine the cheaper processing of that data with new ways of using data. So, pure information processing disappears, but then customer relations and other collaborative sort of work with AI seems to expand.

So, I believe it could potentially have tremendous displacement, but it also has a tremendous ability over the long run for creating new types of work and commerce. I don't know what the net effect is going to be. It's still open. And there's also a lot of low-hanging fruit. AI can learn today by basically being trained on what existing workers do. And it can learn what good workers do. It's less clear that once we have gotten to that level, that we're actually pushing out the frontier further, rather than mainly having AI help a lot of people get up to the level of what's already in the trained data. That's an open question still.

Laura Veldkamp: Alan, I have no idea about the productivity number. I can tell you we did estimate a time fixed effect for the AI firms in our sample, and it was increasing in every year in our

General Discussion: Technology and the Labor Market

12

sample, but it's still going up. So, I don't know where that leaves us.

About heavy-handed regulations. So, I'm glad you asked this. No firm should be disadvantaged by a requirement to post two prices. Here's why. Because you might say, well, having to deal with private transactions and non-private transactions is costly. Well, then you should assign that cost to the private transactions and make those transactions more expensive, just like anything else. If I produce a pen and this pen is more expensive to produce, I charge more for it. If a firm prefers not to engage in private transactions, I am not suggesting we cap prices; as an economist, I wouldn't support that. I'm not going to do that. They could charge a million dollars for the pen privately, and then nobody's going to buy the pen, right? Either you're going to buy it and give up your data or you won't. But then you could say, well, then what's the point, right, if they set these really high prices for private transactions? The point is, if there exists a price, then first, another firm could come and say, you know, some people like pens and like privacy. I sense a business opportunity here. I can do a lot better than a million dollars, and there's a clear business opportunity for them to enter.

And second, consumers learn about the value of their data. One of the reasons, right now, I think we're not getting very well compensated for our data is that I have no idea what the value of my data is in a pen transaction or a coffee transaction, because I never get to see what the price would have been if I didn't give them my data. So, you know how sometimes you go to a foreign country and you land and, like, the exchange rate is, you know, 2,632 to 1, and you're hungry and you get lunch and you're like, you know, 12,000 sounds great. Sure, I'll take it, right? And you realize, two days later, you just got totally ripped off because you

had no idea how to transact with this. Even though we pay with money, we are often like uninformed tourists who don't realize we've been overcharged because we never see a comparison between the two prices. So, just getting to see what the with and without scenarios are might help us be better informed consumers.

And then, Chad, the future of labor markets. AI is scalable. In theory, it can do all our jobs, but the power is not available, right? And it's power hungry. So, there are going to certainly be some jobs for which, yeah, AI could do a better job and make fewer mistakes, but you know what? It just costs too darn much, and we're a pretty efficient source of energy to solve that problem. We have a comparative advantage. Everybody's talking about absolute advantage. The fact that AI performs a specific task more efficiently is an absolute advantage, but as economists, we should focus on comparative advantage. So, we've got to think about what the input costs are. And maybe one job of the future might be a data producer. Maybe a low-skill job is somebody who allows a firm to track all of their keystrokes or track their retina, what they look at and what captures their attention, or see what they see during the day and sign rights to that data over for a wage. Just a thought.

Ufuk Akcigit: On productivity, I think there are two components.

First is the near-term effect within existing technologies and tasks: which workers become more efficient, and which tasks become obsolete. That's the standard discussion, and it's important.

Second are the medium- and longer-term effects: new technologies and new tasks that we can't yet specify. The nature of innovation is that, if market forces are allowed to operate,

General Discussion: Technology and the Labor Market

14

entrepreneurs—often young and unexpected—develop new applications and new firms that take the economy in new directions. Over the long run, that dynamic is a central driver of productivity.

So, while the task-replacement discussion is essential, we should not ignore what happens to entrepreneurship itself. Entrepreneurship is what brings future firms, applications, and technologies. And here I think the picture could be improved.

What we increasingly observe is that, instead of pushing genuinely new ideas, many entrepreneurs build technologies that are close complements to incumbents—hoping to be acquired. We also see concerns consistent with “killer acquisition” strategies. In earlier work, we showed that when the same inventor moves from a smaller firm to a large incumbent, their compensation rises sharply but their measured innovativeness tends to fall, and concentration rises.

That suggests we need to preserve the conditions under which market forces and entry can drive productivity, so young entrepreneurs can set the direction.

And to connect this to universities: research itself has changed. Decades ago, economists often wrote alone. Now, we work in teams, which requires resources, but can generate more ambitious, data-intensive research. The shift from a solitary model to a lab-like model took time, and it remains more advanced in the United States than in many places in Europe.

AI is moving even faster. That creates a challenge for how academic research is organized and funded—what infrastructure we need, and what funding model will support it. The future path of productivity will be shaped both by market-driven entrepreneurship and by the basic research engine in universities.

As for point forecasts of productivity effects: speaking as someone who works on productivity, I can't offer a tight confidence band. The uncertainty is enormous. Thank you.