

Tim Phillips [00:00:00]:

Today on VoxTalks Economics, what will the impact of AI be on the jobs we do? Welcome to VoxTalks Economics from the Centre for Economic Policy Research. My name is Tim Phillips. Every week we bring you the best new research in economics. So remember, subscribe and follow us on Instagram as well at VoxTalks Economics. By automating non routine tasks, AI may have a profound effect on the jobs we do and even whether those jobs are going to exist in the future. So how much should we fear and how much should we welcome this change? In the second of our VoxTalks from the Chicago Booth School Economic Experts Conference 2023, I've been speaking to John van Reenen of the London School of Economics and MIT, who is of course an expert on the economics of innovation. When I caught up with him, I began by asking him how much we really know at the moment about how AI will affect our working lives.

Tim Phillips [00:01:16]:

John, your colleague Canice Prendergast talked about how when we're looking at the application of artificial intelligence in the economy now we're at the wild west stage, how much do we know and how much do we still need to find out about the impact of AI?

John van Reenen [00:01:32]:

Well, we really are in the wild west stage, and we've been there for a while, actually, because with any new technology, it's very difficult to predict and know what its impact is going to be on the economy, on productivity, on people's wages. So we really are trying to do a lot of guesswork and so much is changing every day, it's hard to figure out. We can learn some lessons from the past. I mean, we have a lot of experience of some forms of AI. It's been around for a little while. It's changed a lot. And we have experience of other technology, some history, and we can use that as an insight to what might be happening now and in the future.

Tim Phillips [00:02:07]:

There is some enthusiasm about AI because we have been so worried about the productivity slowdown, particularly where we are in the UK. That has been going on for more than a decade now and people think, well, maybe AI is going to get us out of this. Is it a realistic hope in the near future from what we know about how productivity changes?

John van Reenen [00:02:29]:

That's a great question. And it's absolutely true, which people often forget, that not just in the UK but around the world. The last 15 years or so, the period since the global financial crisis of 2008-2009 has been a pretty bad one for productivity. Productivity being how much output you can get for every hour of human effort you put in. So it's a good measure of how things are improving, how the size of the economic pie is increasing. That generally grows from year to

year, but it has slowed down. And in the UK, for example, as you mentioned, it slowed down from something like two and a half percent a year in the dozen years before the financial crisis to about 0.5% a year in the dozen years after the financial crisis, even pre COVID. That's a huge fall since the longest period, certainly since the Second World War. So the question you asked was, is AI a way of pulling us out of that? AI, like many other technologies, it holds up some promise for doing that. If you dig into where productivity growth comes from. The lesson from history is it really is a story of technology. It's not really about just throwing more people or even throwing more machines and buildings. It's really about learning how to do things differently. So AI has a lot of promise. However, in practice, it takes a long time between a new general purpose technology and something that can be used not just in one industry, but in many industries and many firms. It usually takes a long time between the invention of that technology and the impact that it eventually has on productivity and therefore people's wages and incomes. Back in history, electricity, if you think that came around, around the kind of 1880s, it took another 20, 30, 40 years before that really started impacting on productivity. It wasn't until people could figure out ways of using it. For example, building factories, which were lighted and open 24 hours a day, running production lines like Henry Ford did 24/7 with the division of labor, that actually helped really supercharge that technology into really having massive effects on productivity. My best guess, of course, I might be wrong and things might be different, and we can talk about that, is that it's going to be similar with AI. AI is not, I think, going to be the immediate panacea for our slow productivity growth. I think it has some promise in the future, it might help us. But certainly if I was like the British Chancellor of Exchequer, I wouldn't be banking on it to rescue us over the next five years.

Tim Phillips [00:05:05]:

One comparison that I've seen often made to the introduction of AI is that industrial comparison over the last 20 years, the introduction of robots to automate manufacturing processes. AI automates different types of processes. Is this a valid comparison?

John van Reenen [00:05:25]:

Well, it's true that there's some automation, but I don't think the comparison is very close for a couple of different reasons. First of all, AI is a much more widely used technology than robots can be. So robots are used in many areas, but they tend to be used a lot in manufacturing. Think of a big car plant being replaced by robot arms. AI can be used almost anywhere, anything where there's human intelligence used, which is most of the jobs that we do. AI has the ability to influence that. So it's much of a kind of wider based technology. The second reason it's different is the type of tasks, and therefore people and jobs that it affects. One of the things that robots have replaced is very repetitive, routine, manual types of tasks. The car worker working or putting liquids in bottles, moving things logistically from one thing to somewhere else in an Amazon warehouse. Those are the kind of routine, manual tasks which robots replace. With artificial intelligence, it's a much wider range of tasks. There's actually tasks which are much more skillful to use, like pulling together information on a legal database in order to make a legal

brief, speeding up the ability to code, which is an IT programmer type of skill, getting information in order to write a scientific article like I have to do, you can actually get ChatGPT or an AI to try and pull these things together. Now you've got to be careful, because sometimes ChatGPT makes things up.

Tim Phillips [00:06:53]:

Yes.

John van Reenen [00:06:55]:

I asked ChatGPT once to put together a bunch of papers that I'd written, and it put some together and it also put some together, which I hadn't written. I was very impressed by that, but unfortunately I hadn't written them. Well, not written them yet.

Tim Phillips [00:07:07]:

Now you have to ask it to write them for you.

John van Reenen [00:07:08]:

Exactly.

[Voiceover] [00:07:14]:

In September 2022, we spoke to Anton Korinek on how autonomous machines might mean there is not enough work to go around. Listen to the episode called; Our Workless Future. We also spoke to Anton in March 2023 on his research into how economists can make use of AI. In the episode, AI is Reshaping Economic Research.

Tim Phillips [00:07:46]:

You've done a lot of work on superstar firms in the recent past, firms that have managed to use innovation and technology to create improvements in productivity and have been massively rewarded for this. Do you think that AI will reinforce the power of those superstar firms? Will it undermine it? Do you think we might have a new class of superstar firms? Do you think we might have a new class of superstar, I don't know, employees or superstar countries because of AI?

John van Reenen [00:08:17]:

It's true that one of the things that you really notice, if you just look at the industrial landscape of economies over the last 30 years, has been that these growth of these very, very large firms, I

mean, in the digital area, we know them the best the Apples, the Amazons, the Microsoft's and so on. A lot of their success has been built on being able to come up with new products, new innovations, that's shifted lots of activity towards them. My sense is that AI will probably reinforce those effects. You can make arguments going both ways. I mean, some people say that because AI is software, it can be copied by many people. Think of Amazon web services. You can start up a new business much more easily. That reduces the entry barriers to many small firms. So I think there's some truth in that. However, one of the things where these digital superstars get a lot of their power from is the ability they have to amass data. So you think about Google. When we put a search request in, Google gives us an answer, which is often a very good answer, but it collects data on that search and that enables it to create a better algorithm to improve its search in the future. And that network effect that's sometimes called is a way of cementing Google's power together, because it's very difficult for another search engine to come in and get access to the same kind of data to build that same powerful search algorithm. AI is built on data. So the reason that things like ChatGPT can be very good at predicting about what you want is it's able to search the open web, amass huge amounts of information, and then give you an answer to what you need. Now, the algorithm is clever. That's one of the breakthroughs we've had. But the underlying raw material is still data, and the companies which have the largest data set are going to be the companies which can best use that in order to improve the algorithms they do. And think about recommendation algorithms on Amazon, or think about ways of personalizing advertising. The access to data is going to be very important for that. In addition to that, who can hire the smartest, brightest people to work on AI? Well, it's going to be the largest superstar firms. So the superstar firms help the superstar workers who are creating this. So I think there's going to be both an effect on, certainly on the very, very top, with people who are able to create new types of AI applications, new types of algorithms, who own shares in these very successful companies. That's going to be a force for creating very wealthy superstar individuals, which we already see. That's the other side of the coin. We'll have to see. I mean, it's a choice, with all these discussions, we should remember technology isn't destiny. Technology helps create more wealth. It creates what economists call surplus. The size of the economic pie increases, but who gets the bigger chunk of that? How that pie gets divided is a question for society, how we set the rules of the game in terms of competition policy, how we help people get the skills they need, how we redistribute or don't redistribute the benefits from that. So it really is all up for grabs.

Tim Phillips [00:11:17]:

I mean, even within these firms that are made up of what we would call knowledge workers. Over the years, I've been able to spend quite a lot of time with these firms, and you see that within them, there are different hierarchies of knowledge worker. There are some people who you would truly say were visionary. There are many people in the middle whose job is really organizing information and organizing people. And then there are a lot of people who do well, for instance, a lot of coding within those firms. I would imagine AI is going to involve a reorganization. Do we know how it's going to affect people at these different levels?

John van Reenen [00:11:57]:

Yeah. So I think take your example of that type of a firm, a digital software oriented type of firm. We think that people who are doing the kind of coding work, it might be engineers, it might be relatively, from an economy wide point of view, relatively well paid people, they're likely to see a lot of pressure on their jobs, because the more straightforward part of coding and the translation of that coding can increasingly be done by an AI algorithm. Now, that has to, of course, be checked. But I think a lot of those coders, if you like, are going to see their jobs or demand for their jobs going down. As you go further up the firm's hierarchy, so people who are trying to manage the coders, of course there's less coders, so maybe you need less of them. But you certainly, as you go further up, you're going to need more skillful people who are able to see what the machine is missing. What is the AI not doing right or doing wrong? How do you coordinate across a large number of different actors in other parts of the economy? At the top, the apex of these organizations, there's going to be types of skills which are very difficult to completely automate away, even with AI. I mean, maybe sometime in the future we get to ask for general intelligence and we can even replace the CEO. My guess is we're a little bit away from doing that. Of course, those guys are also going to be putting a lot of pressure to stop that happening as well. They have much more power to do that than people further down the hierarchy. So I think you will see this removal of some layers of the hierarchy, more wealth going to the people at the top. I think, though, the people who are still there, who still have their jobs will have to be doing a kind of wider range of different jobs. So there will be a sense, I think, in which there is for the people who are still in those firms who do have these skills which are not easy to automate away, that they will actually have more decision making power, there'll be some more decentralization and have very kind of rich and interesting jobs, at least until the AGI comes on.

Tim Phillips [00:13:48]:

Well, this brings us to the simultaneously fascinating and scary implication of AI, the creative destruction implication. You've just edited a book about creative destruction. There is always resistance to the destruction part of this by people who stand to lose. Do you think there will be resistance to AI, effective resistance to the spread of AI, both within organizations and between them? Where does that come from?

John van Reenen [00:14:17]:

Well, thank you very much for the plug, for the book. I have to say, by the way, Economics of Creative Destruction, available from Harvard University Press by myself and Ufuk Akcigit. So innovation by its nature has these two faces. Often the second face was ignored, right? So focus, and economists are guilty of this, on the positive side, the creation of new ideas which get commercialized, and that's how societies grow, that's how we become rich, that's how the humans have been so successful. But the darker side is the fact that the older ideas gets made obsolete and destroyed. The older jobs, the coders in routine coding, those jobs go. That

process of creative destruction is the beating heart of market economies. But it also has this real tension overall. New jobs get created, new, old jobs get destroyed, the overall number of jobs. And I'm not so worried that we're all going to be unemployed. And that's not...

Tim Phillips [00:15:08]:

You're not worried about that?

John van Reenen [00:15:09]:

I'm not worried about that because the history of technology is that new jobs get created, old jobs get destroyed, and we manage to find new ways of doing things, new demands for doing different things. But the people who do lose their jobs or who see their wages under pressure will feel angry about it, rightly, they feel upset about it and will resist. I think the way to think about that is to try and create societies which enable it, what we call a just transition, a way of helping people who are losing out to benefit from the overall increase of the economic pie. And one way to do that is actually to try and find new roles for them. Retraining, reskilling, alternative jobs, either in the same company or in other companies. There's this concept in Denmark called Flexi Security. One reaction would be, say, okay, we want to smash the machine, stop the AI to try and prevent this happening. Even if it was desirable, which I don't think it is, it's impractical to do that because if one company does that, another company actually adopts it and outcompetes. Or if one country says, okay, we're not going to adopt new AI, there's going to be plenty of other countries which will and outcompete the country which doesn't adopt it. So I think that there's lots of calls for AI taxes and things like that. I don't think those are the way to really deal with the issue. The best way to do it is through this flexi security idea, which is to create a society where people can get retrained and move jobs and new roles to get new tasks. So I think that's the best way to do that. Now, that's not always possible. There's going to be people who do lose out, people who are in their 50s, like myself. It's harder to retrain, it's harder to do things which are different even with our lives expanding and new types of technology. So I think you have to think of having a redistribution as well. So you have to have a system whereby through the tax system and through the Social Security and benefit system, you compensate people who are going to lose out from some of the turmoil they have. And if you don't do that, what's going to happen is just as we're experiencing in many parts of the world, a lot of political reaction against change as a whole. And because that, I think, is not desirable, we got to embrace change. The best way to do that is to have flexibility to move people around, but also to have generous compensation of people who are losing out from the process.

Tim Phillips [00:17:26]:

But John, the lesson surely of 30 years of technological change is that the institutions that ensure a just transition have been weakened. The policies that ensure that just transition outside places like Denmark have not been put in place. What makes you think it's going to be different because of AI?

John van Reenen [00:17:47]:

Oh, well, I mean, I don't think it's preordained, but I think many countries have done this better than other countries. It's not just the Denmark's or Sweden's, but also Germany, for example, has actually been an example of a country with all we know is that it's had problems and issues, of course, but it's managed to have very low unemployment, as low unemployment than America or Britain does, for example, and hasn't had the same huge increase of inequality that you've seen in the United States, and to a lesser extent in the United Kingdom as well. Part of that is the retraining aspect. Instead of doing what happens in many countries, I think broadly in the United States or the United Kingdom, if you haven't gone to university, forget it. They actually have a whole system of training and apprenticeships and a more equal education system, which actually enables people to get access to the kind of skills they need. They have stronger unions, those institutions are still there, and so they've spread out the benefits of new technologies more equally than has happened in other countries. Now, that's not to say that system is perfect, but I do think that unless we think about ways of changing our institutions to allow for that, then I think it's going to be an increasingly polarized and unhappy society that we live in. We can see that happening. I'm not saying that's just because of technology, but technology has played a part in it, both directly through what's happened to people's jobs and wages, but also indirectly through what's happening in social media and the way people get information, the way political campaigns are being run. That's led to much more polarization. So I think that is an area where we have to think of the regulation of AI. And that's not saying we stop it, but we have to think of ways of dealing with it and regulating it in a way which is appropriate for our societies.

Tim Phillips [00:19:29]:

Which leads me to the last question. You are a researcher, but you're also an educator. Does education need to change to create the next generation of people who are going to be in the labor force alongside AI so that they're not replaced by AI.

John van Reenen [00:19:46]:

No, absolutely. My profession as an educator has been profoundly influenced and changed by AI in a number of ways. And one of the positive things now is that it's much easier to generate material and spread that out to lots more people. I was giving a lecture yesterday to 1000 people in a virtual PhD program on climate change and innovation. You can't get a thousand people into the biggest lecture hall at MIT or LSE. So there's lots of ways in which technology is enabling material. But on the other hand, we have to think about how education becomes more flexible and get away from this idea that people go and do a college degree for four years or three years, and that's it. We have to think on the supply side of education as ways of providing tools and materials for people who can learn throughout their lifetime, structuring that helping people learn is really important and also the ability to learn the kind of skills which won't be so

quickly replaced. Here's another example which keeps us all, I won't say up at night, but certainly we don't sleep quite so easily. How do you test people? How do you test that they learned things? Students can use ChatGPT, every time you set them a take home question, you get these essays back. Is it written by the student or is it written by a clever algorithm? And I don't think we really know the right answer to that. One solution is to say, let's just get them in the classroom and we put firewalls around the classroom. And when we test people, they have to sit in a room and maybe even verbally give the answer. So going back a few hundred years, I think outside a narrow part of what maybe we do, I think that's not really desirable or feasible. One idea might be to say, well, if you think about ChatGPT, it's not so much about the answers you get as about the questions you ask. And if you use those types of generative AIs, the first question you ask doesn't give you a very good answer. And then over the course of time, you refine your questions and this gives you a much better useful answer. So it's the sequence of questioning which is very much like a research program, actually. It's the sequence of questions that you ask. Eric Pan Jolson, my coauthor at Stanford, is very keen on that. Of course, the problem with that is that then you ask ChatGPT to generate the sequence of questions. So you have to think even more creatively about what you could do. These things will pose challenges, but I think with imagination we can think of ways of adapting to those technologies and coming up with other ways of testing students.

Tim Phillips [00:22:05]:

John, let's hope your productivity stays high, because there's a lot of questions that do need an answer. So thank you very much for talking about it today.

John van Reenen [00:22:12]:

Well, thank you for inviting me.

Tim Phillips [00:22:21]:

John's book is called The Economics of Creative Destruction. It's published by Harvard University Press. And he's also a podcaster as well. He's one of the hosts of the new Innovation and diffusion podcast at the London School of Economics.

[Voiceover] [00:22:41]:

This has been a VoxTalk from the Centre for Economic Policy Research. If you enjoyed this episode, remember to subscribe. You'll find us wherever you get your podcasts. Next week on Voxtalk's Economics; Why are female politicians more often the victims of violent attacks?