Ep. 26 Productivity Puzzles podcast transcript

Artificial Intelligence and Productivity

Episode release date: 29 June 2023

Host:

- **Bart van Ark**, Managing Director of The Productivity Institute and Professor of Productivity Studies at The University of Manchester (BvA)

Guests:

- **Erik Brynjolfsson,** Jerry Yang and Akiko Yamazaki Professor and Senior Fellow at the Stanford Institute for Human-Centered AI (HAI), and Director of the Stanford Digital Economy Lab (EB)
- **Tera Allas**, Director of Research and Economics at McKinsey & Company, UK (TA)
- **Lea Samek**, Economist at the OECD Science, Technology & Innovation Directorate (LS)

BvA: Will artificial intelligence rescue us from the productivity demise? If humans cannot get productivity up, can intelligent machines bring us the productivity revival? Or is it still the same old story of people and machines having to work together? We are going to find out. Welcome to Productivity Puzzles.

Hello, and welcome to Productivity Puzzles, your podcast series on productivity brought to you by The Productivity Institute. I am Bart van Ark and I'm a professor of productivity studies at the University of Manchester, and I'm the director of The Productivity Institute, a UK wide research body on all things productivity in the UK and beyond.

Welcome to the June 2023 episode of Productivity Puzzles on artificial intelligence and productivity. Say technology today and everybody will talk about the one big thing: artificial intelligence, or in common speak, AI.

While certainly not the only digital technology that has come along in the past few decades, AI perhaps speaks to our imagination more than all those before because it directly impacts on the daily activities of many of our listeners to this podcast.

While the machines of the past were mainly replacing physical power, and more recently pure computing power was about replacing routine tasks, AI is about replacing activities initiated from the human mind. And the arrival of ChatGPT and all recent follow-up applications show that AI is developing at a record speed and makes us all wonder if what we can do cannot at some point be done faster and better with AI.

As we will discuss on this podcast, AI will impact on the majority of today's occupations, and could potentially affect up to half or more of the tasks and work activities we do.

Of course, AI is also a potential driver of productivity. But as we have seen before, potential is not the same as realisation. There are many predictions and projections out there, but the impact of technological change is notoriously difficult to predict.

Questions are whether AI will be fully exploited, what the drivers and the barriers for adoption are, and how long it will take to see the effects emerge.

So, to work ourselves through these complex questions, I am fortunate to be joined by a fantastic panel today, which will give us different perspectives on AI, productivity, and related topics such as innovation, jobs, and industrial strategy.

First of all, I welcome Professor Erik Brynjolfsson from Standford University, where he has a large number of roles, including associations with the Graduate School of Business and to the Economics Department.

Erik is one of the world's top researchers studying the effects of information technology on business strategy, productivity, and performance. And he is the author of bestsellers such as The Second Machine Age and Machine Platform Crowd.

He's also the director of the Stanford Digital Economic Lab, and associate of the Stanford Institute for Human Centred AI and the Stanford Institute for Economic Policy Research.

Quite a record of activities, Erik, but it's great to have you on and thank you for giving us some time today.

EB: My pleasure, looking forward to it.

BvA: Our second speaker is Tera Alas, who is the director of research and economics in McKinsey's United Kingdom and Ireland office. She leads McKinsey's research on the economy, growth, and productivity, and brings together expertise and experience in strategy, corporate finance, economics, and public policy.

Tera also works very closely with the McKinsey Global Institute, which is McKinsey's research arm. And she is chair of The Productivity Institute's advisory committee. Great to have you on, Tera.

TA: Thanks, Bart, really glad to be here.

BvA: Looking forward to hearing from you, including some of the recent work that McKinsey has been doing on AI.

Our third panellist is Lea Samek, who is an economist at the OECD in the Science, Technology and Innovation Directorate. Her work encompasses a wide range of innovation and industry policy related topics, including jobs

and skills in the digital transformation, and the diffusion of artificial intelligence. Lea, welcome to this podcast.

LS: Hi, Bart. Thanks, excited to be here.

BvA: So, we have a lot to talk about in this podcast, including what AI exactly is and what it means for productivity. Some of the projections of its impact, the drivers and the barriers for diffusion of AI, the implications for jobs, and what it all means here in the UK.

So, we should probably get going and benefit from what our panel has to bring us on this.

And Erik, I would really like to start with you, to talk a little bit about what are we talking about in the case of AI. A couple of definitional issues because this is one of those basket terms.

EB: Uh-huh.

BvA: So, what kind of AI should we focus on and what kind of AI is most important from the perspective of studying productivity?

EB: Well, there have been several waves of artificial intelligence. I started working with expert systems in the last century. Then around 2012 there was really an explosion of interest in the category of machine learning, especially deep learning.

Geoff Hinton famously scored the highest on ImageNet challenge, recognising images using deep learning techniques, and very quickly everybody else switched over to that. So, for the past decade there's been a ton of interest in that.

And then most recently, generative AI. Using these technologies not only to understand and read images but also to generate text and images, and now increasingly sound and computer coding.

What's different from the current wave compared to earlier waves is that these technologies can learn to solve problems from looking at information. Whereas most software, previously humans had to write step by step exactly what they wanted the machine to do, you know, which could generate trillions of dollars of value in many kinds of software. Now, with machine learning, the machine learns from the data and instantiate, you know, hundreds of billions of parameters to make much better predictions and understanding than we did in the past.

BvA: Why is this the big game-changer for productivity?

EB: It's a huge game-changer because it affects so many of the tasks that humans can do. You know, large language models are generating emails,

generating computer code, generating fiction and non-fiction stories. There's not many of us that don't use language in some part of our jobs.

At least 55, 60 per cent of the US workforce are primarily information workers or knowledge workers. So, all of those jobs are going to be affected. Especially people who work a lot with language, like, say, lawyers or some of the people in marketing and advertising. And for that matter, professors and teachers.

And the gains are quite significant. I'm sure we'll get into this more, but we're already seeing significant productivity gains in a lot of areas: writing, coding, call centres. And so, if you multiply that large section of the economy by significant productivity gains, that translates into a big impact over the coming decade.

BvA: I think what you see a lot is that AI is a very important technology application, but it's critical to be connected with all our technology applications to create productivity. You know, software, hardware, data collection, even bioscience. It is a very complementary type of technology. Does that, sort of, impact the way you're thinking about the potential for AI to generate productivity effects?

EB: Yeah. So, it is part of a complementary system. I think it's fair to call it a general-purpose technology, like electricity or the steam engine before it.

I was visiting DeepMind there in London a few years ago, and I was struck by a slogan they have there: Our mission is to solve intelligence and then to use that to solve the other problems of the world.

And it sounds very ambitious, but I think there's a lot of truth to it, that if you can solve, or at least improve the ability of machines to be intelligent, it opens up a wide set of new things. Arguably, it's the most general of all general-purpose technologies.

I think we're still in the relatively early stages of it, but it's something that I think we'll look back at and see that we're at the dawn of a staggering set of improvements in our ability to solve many of the problems that face the world today.

BvA: Tera.

TA: And one of the things that's really exciting, I think, about this latest leap in technology is that this is not just about cost reduction, it's not just about automating what people do. But there's a lot of creative energy around…if you just think about from a business perspective, you can use this tool to generate more revenue and not just lower costs.

So, whether it's in marketing our sales or it's in pricing, or it's in your innovation functions, you're going to be able to grow the top line as well as reduce your costs. And so, that should give a huge boost to productivity.

BvA: So, Lea, at the OECD, you're looking at this broad range of different technologies, and you're looking both at the effects on productivity but also on the effects on employment.

We'll talk more about employment in a minute, but there is this, sort of, major concern about the employment effects of this. Do you see that AI is really a different technology application compared to all these other technologies that you're looking at?

LS: I think this is a bit of a tricky question, right. Because I think on the one hand what we've already discussed just at the beginning now of the podcast is that previous waves of technological progress were, kind of, primarily associated with automation of routine tasks, right. So, for instance, we have computers that can do, kind of, record keeping, or calculations. We have industrial robotics that, you know, can do a programmable…things like welding or packaging.

Now, there is this, kind of, new change happening where the game's really changing, right. So, recent advances in AI mean that non-routine cognitive tasks can also increasingly be automated. And so, the distinction between tasks has, kind of, evolved. You know, we're no longer talking necessarily about routine/non-routine, because it's really impacting different tasks within the same occupations.

I think a lot of work previously focused on how AI impacts on certain occupational groups. But actually, what we're seeing is that the, kind of, skill mix that is demanded within occupations is changing because the tasks that are used are changing.

BvA: Erik, to wrap this section up, and then we'll go and talk a little bit more about the numbers here. But you co-authored a piece with Brookings recently – which is in the show notes of this programme – and there was one sentence that I particularly found interesting, and that is that the criticism of large language models is that they're merely stochastic parrots.

Can you talk a little bit more about why that criticism arises and why, I think, you don't agree with that?

EB: Yeah. Well, it was a clever title of a paper, but it really didn't grapple with what these technologies are doing.

The idea is that…what large language models do in some ways is very simple, they predict the next word. So, given a few thousand previous words, what is the most likely next word. And in that sense, it's, sort of, a prediction and repeating back part of what it's read elsewhere.

The thing is that it in practice generates a lot of new content. And in order to make that prediction effectively, the model has to take into account a lot of data, has to, in some sense, have insight into, you know, what a mammal is, or geography, or chemical properties, or many other things in order to make a meaningful prediction.

And in fact, I have the opposite reaction, that these are remarkably powerful and make insights that I wouldn't have expected. And when I talk to the top AI researchers, they're also surprised at how well they are able to create insightful prose.

And all of us here are economists and businesspeople, I think ultimately the question is less whether this is, quote, truly intelligent or not, and more does it create value. And there the answer is an unequivocal yes; they're already being used in many applications to create value.

And so, in terms of its impact on the economy, on jobs and productivity, that question has been answered, and it's only likely to create more value in the coming years.

BvA: Yeah, and there was another sentence from the Brookings report I'd like to quote…a report which I can highly recommend, by the way. And that quote is: Economic value depends not on whether they are flawless technologies but on whether they can be used productively.

EB: I think it's a very important point. These models are often flawed in the sense that they will give false information, they will hallucinate, or confabulate. They're far from perfect. And many times, they make mistakes that no human would make, and that can be disconcerting.

So, when people are using them, they need to be very careful. And it's often a good idea to keep humans in the loop and use them to augment people, not simply replace them.

But ultimately, the question, I think, is can they create value that wasn't there before. And there, if they're used correctly, the answer's a strong yes.

BvA: Okay, so let's talk about the effects on productivity more from a quantitative perspective. Because there are some really courageous attempts out there, and some of you have been involved with those courageous attempts, to provide some kind of ballpark estimate on the effects of AI application on productivity.

But it's not easy, of course, to translate some known effects from very individual use cases to something that looks exciting at the more aggregate level, whether it's for firms, or for sectors, or for the economy.

So, I'll just give you some numbers. Goldman Sachs came out with a report earlier this year, said that up to 2030 we would probably get about a one and a half per cent point productivity gain from the application of AI for the next ten years. But at a very wide range, between, sort of, 0.3 and 2.9, which already tells you something about how uncertain these estimates are.

McKinsey, which is the other report we will talk about here, and which is referenced in our show notes, shows an incremental effect of the latest generative AI of between 0.1 to 0.6 percentage points per year through 2040. But combining that with all other technologies, could also add between 0.2 and 3.3 percentage points, depending on the speed of adoption and worker replacement to other activities.

And then Brookings, the report that I mentioned earlier, that Erik, you have been involved with, is talking about something like between 0.9 and 1.8 per cent over the next 20 years as well. Sort of, you know, what you want to look at.

Now, you can…all these studies are in our show notes, and you can study them in more detail.

But let's unpack it a little bit, Erik, maybe you can talk us a little bit through how you and others get to these estimates, and where these wide ranges in terms of the bandwidth of these estimates are coming from.

EB: Yeah. Well, I think it starts with what we already see. So, with some technologies…I don't know, Blockchain or 3D printing, there's a lot of speculation about what it could do in the future. With these technologies, we're increasingly seeing that there's already people putting them into work.

So, there was a nice paper by Kalliamvakou, who looked at software coding and found that they coded about twice as fast. Noy and Zhang for a writing test, also about twice as fast. Anton Korinek looked at 25 use cases for economists, like us, and found ten to 25 per cent increases.

And, you know, the one that I'm most familiar with is a paper that I wrote with Lindsey Raymond and Danielle Li, where we looked at call centres. And there, the nice thing was this was not just in a laboratory, this was rolled out in a real company. They've been using it for a couple of years and they use the technology to augment the workers, not replace them.

And what we found was that there was an average of about 14 per cent productivity gain for the users who were using it versus the ones who weren't. Strikingly, it was about 35 per cent for the least experienced workers. So, those who were just beginning to use the technology got the biggest gains.

And this goes back to what I said earlier about how the technology captures some of the tacit knowledge from the most experienced people and makes it available to the less experienced people, so they benefited the most.

We also saw a big gain in customer satisfaction. Also separately in customer sentiment, the number of happy words versus the number of angry words in the millions of transcripts we looked at. Much more happy words.

And even employee turnover improved. People were less likely to quit, they apparently liked working with it more. So, in all these dimensions there was improvement.

Now, if you take these individual case examples and you multiply them by the share of the economy, you quickly get some very large numbers. Much larger than the ones I published, so I think maybe we're lowballing it.

On the other side of it though, of course these gains don't translate into the bottom line instantaneously. There will be some period of adoption, and that could take years. You know, we'll probably talk later about the productivity J-curve in the adoption. But if you spread those gains over say a decade, that's how we get those kinds of numbers that we had in the Brookings piece.

I should say, it's certainly fair to realise that there's a wide range around those. It may be that we are able to adopt things faster, I think there's reason for that. It's also possible that additional barriers will emerge, 'cause there are a lot of things, as we touched on earlier, a lot of weaknesses with these models as well.

TA: So, one of the ways to think about it may be the, sort of, numbers that people quote out there is to think about it on three different levels. First of all, what are the tasks that could be automated. And secondly, what tasks do make up a full job, and then what are the jobs in the economy. And so, the examples Erik has just given are actually a really good illustration of that.

So, at the level of a task of a coder writing a software programme, maybe the gains can be as much as doubling of productivity. But even coders don't spend their whole time coding, they spend a lot of time communicating and talking to other people, talking to their users, et cetera.

And then at the level of the job, that's the other example Erik gave, about customer service operators, maybe the gains are more in the, kind of, order of magnitude of ten to 15 per cent for that particular type of job.

And then you look at all the jobs in the economy and, you know, the flip side of more than half of jobs could be augmented or partly automated is that

another third of jobs, roughly speaking, don't really have any potential for automation using AI.

And so, when you add it all up, you get to a, sort of, more reasonable number. But I fully agree also with Erik that the adoption parts are really uncertain.

EB: I really like your framework there, Tera, because you're right, we do a lot of the analyses down at the task level, and a typical occupation might have 20 or 30 distinct tasks. Radiologists, I was just looking at, have about 27 distinct tasks. And then you could roll those up to those occupations. But then further up to the level of process, where often many different occupations have to interact to deliver value for the customers or for the society.

BvA: So, Tera, the McKinsey report that was recently published shows that the effects are really quite differentiated by function and by industry. Can you talk a little bit about that? What kind of functions are particularly impacted in terms of productivity improvements, and what kind of industries are particularly impacted in terms of productivity potential?

TA: Functionally, it's really driven by the fact that the McKinsey report was looking specifically at generative AI. And we've talked earlier on about how it's all about language understanding and language production as well as, of course, images and sound and so on.

So, if you then put that into the business context, a lot of the applications for that, sort of, functionality are going to be in sales and marketing. So, generating marketing copy, interacting with customers. So, customer operations, for example, augmenting the way that your customer service operators in call centres can help their customers find solutions to their problems.

A huge benefit to software engineering by using co-pilots, which can help you write code much faster and better. And by the way also, the coders are happier because they don't have to do a lot of the boring stuff. And so, there's also enormous potential in R&D and innovation, where you can go through a lot more, as it were, alternative innovations quite quickly, and generate creative ideas that maybe a human even couldn't have come up with.

And so, when…so, those functions make sense from the point of view of what generative AI is really, really good at. And then when you overlay the functions and look at in which sectors are those particularly important, you end up with the sort of sectors that are mentioned in the report, like banking and finance, but also of course software generation itself.

And by the way, education as well, as Erik mentioned earlier, is possibly going to be quite upended or changed by the use of AI, both by teachers and educators, and scientists as well as the students.

BvA: The OECD has also done research along, you know, the areas where we see the biggest impact. Are these results aligned with what the McKinsey report found?

LS: Yes and no, because what we actually find, depending on whether we are looking at firm level data or online vacancy data – so we really get a feeling of how the workforce is supposed to look like in these big AI firms – we actually see that AI is, kind of, permeating every kind of sector.

However, at the same time, we do see that the highest share of AI adopters is really within the information and communication technologies. Which, I think, you know, is not surprising, as Tera was saying, you know, given that AI applications most often really originate in the ICT sector. But then we also see a huge proportion of those jobs actually in professional and scientific services.

And the interesting part is that we see this regardless of what kind of data source we are using.

So, we recently published two different papers, one which was a study of colleagues of mine where they tried to, kind of, widen the scope by exploiting confidential and representative firm level data. And so, they basically sourced information from official firm level service on the use of ICT from 11 countries. So, this is a representative study.

But at the same time, our team recently published work where we tried to use a very novel approach, and we tried to identify and characterise AI firms based on different sources of large microdata. And so, this was an experimental study, not looking at the 11 countries but only at the UK for now, where we combined data on intellectual property rights. So, here the idea is to really identify the firms that are developing AI and innovating in AI. Then matching that to website information of the firm.

And then we are also linking that to online job postings. And again, here the idea is, you know, if tasks are changing because firms are adopting AI, then this should be reflected in the skills that are mentioned in those job postings.

And so, the kind of, you know, promising result here is that regardless of what source we used here, we, kind of, came to the same conclusion.

BvA: Interesting, that's really interesting work. Now, before we leave this topic of the productivity impacts, Erik, you know, I always have to ask this question about the measurement side of this, right. So, any sceptic will say,

well, how can you ever measure the output of these kinds of new technologies.

So, you've done a lot of work on that, and you're using the term, in the Brookings paper, silent productivity growth. Can you talk a little bit about that?

EB: Yeah. I mean, this has always been a problem, is measuring productivity. We all know there are issues with measurements, some things are easy to measure, some of them are worse. And I think unfortunately, in many ways our measures are getting worse on average, because more of the economy is becoming digital, more of it's becoming intangible, and those are much harder to measure.

The set of activities that are being affected by these generative AI, and more broadly other types of AI, have two particular weaknesses in particular. One is that many of them take time to have a real impact on other measures of the economy.

So, Tera mentioned education and, you know, as we do a better job of teaching people algebra and calculus and critical thinking and other skills, they may not translate into greater output until years later. Or as an engineer designs a better aircraft or is able to come up with new biological substances to help with healthcare, these are all things that will show up years, sometimes decades later in the productivity statistics.

The second problem is that some of them won't show up at all. When we create digital goods or entertainment goods, if generative AI helps us make a better movie or a more interesting poem, or a better non-fiction or fiction book, that doesn't directly show up in our productivity statistics. The sales may be just the same. Unless we do some kind of effective quality adjustment, which is very difficult to do, we're going to miss those gains as well.

My team and I at Stanford have been working on an alternative metric we call GDP-dash-B that measures the benefits not the costs of goods and services. It seeks to measure the consumer surplus people get and that would help a little bit. It would give us a chance to see if you really like some of these goods better than the other ones, then that would show up in GDP-B as an increase in consumer surplus.

So, we can come up with some alternative metrics, and digital technologies make it easier to use those new metrics, but we do have to take all these measures with a grain of salt. And my belief is that we substantially underestimate productivity growth with the current measures.

BvA: Well, these are our best estimates for the time being, and as we discussed, are based on a variety of assumptions. But one key assumption is the speed of adoption by companies and organisations across the economy to

adopt AI in combination, as you already discussed, with other technology applications and innovations.

So, Tera, the latest McKinsey study, again quoted in our show notes, sketches a series of adoption scenarios creating a time range for the potential pace of automating current work activities. Can you explain that a little further?

TA: So, the McKinsey report adoption scenarios I think are really interesting in three different ways.

The first one is if you just look at them, the range is absolutely huge in terms of how quickly or slowly we think that these technologies might actually be put in place and start replacing or augmenting human labour.

The second point is what's interesting also is to compare to what we said in 2017. And based on that, the, sort of, point at which maybe 50 per cent of human labour might be automated, and now that number has jumped to 60 or 70 per cent. So, this use of language, understanding of it and generation of it, is a really big, kind of, boost to the degree to which machines can now take over or at least augment a lot of the work that humans do.

But the final point is if we look at actual adoption, we're, sort of, tracking towards the bottom of the scenarios that we said in 2017. In other words, adoption so far has actually been slower than we thought back then.

Now, there are lots of reasons, and lots of different things that go into both the estimation of adoption but also what actually happens on the ground. And one of the things that of course as economists we think is that adoption is driven by costs and benefits. People will look at these technologies and say, it'll cost so much of time or hassle, and here's the benefit I can get. If it's worthwhile, I'll go ahead and adopt.

But I think there's a lot of nuances that we need to put around that in the real world. And the first bit of nuance that's super important is that the actual costs and benefits are often really different from the perceived costs and benefits, and maybe even risks.

So, one example that I give of this is that in our UK work we have asked businesses why it is they're not adopting digital technology. And a fifth of them, so 20 per cent, say, well, we don't think there's any business case for any kind of digital technology. So, obviously from their point of view, the benefits are not sufficient for them to research the solutions.

Another important point is that a lot of AI adoption I think is now invisible. So, another survey that was done actually earlier this year, in April, the ONS asked businesses what kind of AI they're using, and 80 per cent of respondents said, not applicable. And if you think about what we just said

about the ubiquitousness of AI and the, kind of, general-purpose technology nature of it, it is applicable to absolutely every business.

However, a lot of those businesses don't know they're using AI. They are almost certainly using search engines, which are almost certainly using AI in the back end, but they're just adopting it. And in the future, they will almost certainly be using, you know, Office, Microsoft products, which very soon will have AI embedded.

And so, a lot of the adoption I think will come through, sort of, this invisible…it's packaged up in a piece of software or in a service that you are buying, and you don't necessarily have to do very much to adopt it.

And so, I think that will make adoption faster, 'cause obviously the hassle of it and the skills you need to adopt are going to be less.

BvA: Erik, I also think that in your work you are arguing that actually, despite all these barriers to adoption, there is actually some specific opportunity that AI adoption could go faster.

EB: I think that's right. I mean, as Tera has pointed out, there are a lot of things that happen between a technology affecting things at the task level versus at the economy level. And whether it's a steam engine or electricity, we saw years or decades until they were adopted.

And it's reasonable to believe that there's a productivity J-curve, where at first you even potentially have a negative effect on productivity as people have to re-engineer and reskill their workforce.

That said, as you were pointing out, Bart, I have a different view about this wave of technology. I think it's going to happen significantly faster. And the reason is that we've already built out a lot of the infrastructure.

So, if you want…if companies want to use ChatGPT, if individuals want to use it, they just have to tap into their intranet infrastructure. And it famously went from zero to 100 million users in about 60 days. I think it's over a billion users now. That's because we already had much of the preconditions.

This summer, Microsoft and Google are incorporating these tools into their office productivity suites. Thousands of other companies are using cloud services and software services to make them available. These were parts of the infrastructure that we built out at great expense and great effort over the past decade or two, now we're able to leverage them.

And the other thing, and we touched on this a little bit earlier, is that you don't need a lot of special skills to use these technologies, you can interact with them in English. In fact, my son does coding, and he was thrilled that he could type in English style language requests, and it would point him to

the relevant Python libraries and GitHub repositories that he needed to access without having to become an expert in some of the subtleties of some of those tools.

So, this is something that…it opens up possibilities for people to use these tools much more effectively and much more quickly than in the past. It still will take some time to do some business process re-engineering, and some reskilling and some rethinking, but not nearly as much as with earlier waves, in my view.

BvA: So, we're not starting from scratch in that sense, it's moving very fast.

Lea, at OECD there is a lot of work around, sort of, intangible capital. I mean, we've already discussed this whole issue of AI doesn't go on its own, it needs a lot of complementary assets. What kind of complementary assets are we talking about?

LS: In the study that I just mentioned earlier, which, you know, is a cross-country study based on surveys, my colleagues looked at AI users. And they found that especially the large firms, they are the ones that tend to be on average much more productive than any other firm.

But they also found that this is not necessarily a…it's not necessarily reflecting the use of AI alone. In fact, what's really a critical role here are complementary assets. And so, these include intangible assets. And intangible assets, you know, it can be human capital, it can be cloud computing, it can be software.

But what they find is basically that AI use is much more likely in the presence of digital infrastructure and also other digital capabilities. So, you know, these would be cloud computing services and ultra-fast broadband connections, for instance. And so, this is just what, you know, Erik and Tera were mentioning, that this is something that's actually most of the time already set up.

Obviously, there are, you know, cross-country differences as well, regional differences. Firms, you know, have differences there within regions of course. But it also highlighted the importance of the workforce.

And so, in this particular study they looked at ICT skills in particular, and they looked at a proxy basically with ICT specialists and, kind of, ICT related training, and this was really key for adoption of AI.

BvA: Yeah, so let's move now to, sort of, the other side of the coin, which is the jobs and the types of work. We already touched a little bit on that, but let's face it a little bit head-on, let's talk about actually the AI workers themselves. You know, it's a relatively small group but they're critical in terms of their skills and statistics and computer science and machine learning. They're actually the ones who maintain those AI systems.

So, how important is this group of workers, what are their characteristics, how dependent are industries actually on these…on this critical but relatively small group of AI workers?

LS: Yes, in fact, exactly as you are saying, the AI workforce is actually still relatively small. It only represents a very small share of overall employment in OECD countries.

And, you know, obviously this share depends on the definition of AI that is used. You know, whether we look in terms of potential workers through vacancy data, or whether we look at employment statistics, you know, what country we look at, what year we observe.

But the number really varies and is still usually less than one per cent of workers in all of these studies. So, we're really talking about a very, very small proportion here.

But those workers really differ from the overall population. We have colleagues in the Employment, Labour and Social Affairs Directorate that recently published a study where they looked into their characteristics in particular, and they found that two in three of them actually have at least a tertiary degree, and this approaches 80 per cent when we look at, kind of, the top ten AI occupations.

They found that, you know, half of those workers earn above the 80th percentile in the earnings distribution, and actually less than 40 per cent of those workers are women.

Now, we also see that the highest share of AI workers are in the ICT sector and in professional and scientific services, so just as I mentioned earlier when we also looked at that firm level data here. But we know that it permeates really all industries.

And, you know, this is what we see all the time. I mean, AI is applied in healthcare, you know, for medical imaging analysis, for patient diagnostics. We see it in retail for inventory management, we see it in manufacturing for process optimisation. I mean, the list really goes on and on and on.

So, the demand for AI workers is not only not limited to specific industries, but we actually also see that it's expanding across sectors as firms really start to recognise the potential of AI to, kind of, drive innovation and competitiveness.

Now, you mentioned skills, and that's a bit of my passion, I have to say. And, you know, these workers are usually endowed with the skills that are needed to design, develop, and maintain AI systems. So, they create algorithms, they train models and so on.

So, of course, when we look at the skills that AI employers list in their online vacancies, we often find, you know, that a strong foundation in maths and computer science, strong expertise in machine learning, anything related to big data, is really important. You know, we also see programming languages always popping up, especially Python actually. So, these are all in high demand.

But what we also find is that socio-emotional skills are invaluable. They always come up; they are always hand in hand with these technical skills. So, we're talking about, you know, communication skills, teamwork skills, problem solving skills. And I think, you know, this is something to bear in mind, because it's really not this, kind of, you know, techie person that's necessarily working in these AI positions. AI employers really require, kind of, a broad skill mix from the AI workers.

BvA: So, does every firm who wants to be serious about AI need those, kind of, AI workers, or at least a few of them to get it going?

I'm a little bit puzzled, because earlier on we talked about how easy this technology is to apply and to implement. But it seems that when businesses try to implement it into their business models, they still need to do a lot of additional stuff in order to make the technology really work.

LS: I think we need to distinguish there between different types of firms, right. So, there are firms that are purchasing, kind of, off-the-shelf AI technologies, right. These are probably the ones that we would not capture here in this analysis necessarily.

What we would capture is based on the data that we are using, the workers that are really, you know, kind of, helping to maintain AI, or the ones that are developing AI. Particularly when we rely on the IP data, so the data where we, kind of, you know, read through trade mags and patent information. And so, I think, you know, we need to make this difference.

And in this matching work that we have done where we combine different data sources, we actually were able to look at the firms that are mainly innovating in AI. You know, the firms that are mainly, kind of, having AI at the core of their business, and the ones that are, you know, really looking for AI talent. And so, these, kind of, skill mixes always jumped out, so it really seems to be that a bundle of those things are needed.

But yeah, there are always firms, you know, which are purchasing off the shelf, and I think this is becoming more and more and more, for sure.

EB: I think this is one of the things that's really changing, is the way that the tools are getting much easier to use. Probably most of your listeners have had a chance to work with ChatGPT and see how easy it is. But many times, as you mentioned earlier, people don't even realise they're using AI when they're doing tasks.

My wife was recently setting up a new website and she wanted to come up with some copy for the website. And she had written some stuff, but then as she went to log in, it said it would be happy to suggest some copy for her. And it went ahead and wrote some copy, and she was, like, wow, this is really good, this is better than what I wrote.

It didn't say it anywhere in there but I'm certain, judging by the quality of what was written, it was using a large language model in the background and was generating some pretty useful copy.

And not only was she taking advantage of the tool without any special skills, I wouldn't be surprised if the company, Shopify, also didn't have to do a lot of their own engineering, because there are so many off-the-shelf tools.

I'm here in the Bay area in California, and there are literally thousands of startups that have addressed various different verticals using these tools, trying to make it as easy as possible for companies and end users. And that's one of the reasons I'm optimistic about the adoption curve.

TA: Maybe just to add to that, Erik's example of a website is a really excellent one. 'Cause if you think back 15 years, you did need somebody who knows how to code in Java, Script, and very…had quite a, kind of, deep knowledge of how the internet worked in order to be able to create even the most basic of websites. And now you have all these tools which are literally drag and drop and you need no skills at all. You need to be able to use a mouse and you need to be able to open a browser, and that's pretty much it. And all the rest of it is, sort of, designed in the background.

And so, I think something like this, as Erik has already actually suggested, as has Lea, is going to come into the programming of AI itself. You know, the meta-AI that's going to help people use AI more efficiently and effectively will exist. And so, that hopefully means that the shortage of, like, very specific AI machine learning, programming, et cetera, skills is going to be…it's still going to be acute probably, but slightly less acute and won't be quite so big a brake on adoption.

BvA: Let's talk for a minute or two about the, sort of, broader effects on the workforce. Again, there are, sort of, lots of estimates out there, but I think the consensus is that – as I mentioned in my introduction – the majority of occupations are going to be impacted one way or another, including our own.

Lea, you referred to that earlier, that if you look at, sort of, the tasks and the work activities of people, then there's not so much the question about some occupations will completely go away, it's more that, you know, tasks in virtually every job will begin to be affected one way or another.

And what I find interesting about these stats, in the past we were talking very much about, you know, digital technology impacting, kind of, routine tasks. You know, computing power was, sort of, impacting these tasks in the middle, routine types, not manual tasks, not, sort of, abstract tasks.

But now it seems it's shifting a little bit to, sort of, think about, sort of, mechanical tasks are being replaced by AI, thinking tasks, you know are being replaced by AI. And then I think you mentioned earlier, Lea, not to come back to that, social skills, empathetic tasks, interpersonal relationships, and those tasks become more important.

Tera, maybe I can start with you, how is this, sort of, affecting people's jobs? You know, it seems that there is a very different call on the kind of things they are expected to do in their job than was the case in what they were doing previously.

TA: I think you're right, Bart, that this is really going to change the way people operate. And if…you know, I always think of it in terms of comparative advantage, what are the things that humans are better at doing and worse at doing. And they don't have to be better than the machine, they just need to be comparatively better at something, and the machine will do the other stuff.

I know it's a simplistic way of thinking, but I think humans still will maintain an edge over machines for the time being in unpredictable environments with unpredictable tasks, whether they be physical or they be emotional or they be cognitive. Because humans still have this, you know, better ability to, kind of, recreate and ask new questions and, you know, think of new ways of doing things.

But there are many, many other tasks which we used to think were, kind of, human domains, such as language, which machines can actually now do better than humans. Or at least in, sort of, certain measurements.

I think what's really interesting about this new wave of generative AI is that in terms of the jobs and people it affects, we are used to thinking of automation of skills biased as…you know, it'll be particularly bad or particularly challenging for the less skilled and the poorer paid people. But actually, generative AI has the biggest impact, relatively speaking, on the highest paid people.

So, you know, if you look at that incremental impact of automatability, with and without generative AI, the big hit is in the top income quintile. And so, there's a bit of a levelling effect then, you know, that everybody's jobs are going to be impacted rather than just previously where we thought that it was mainly going to be the lower paid and lower skilled jobs that would be impacted.

LS Yeah, I think, you know, it's, like, the areas where AI is currently making the most progress – you know, like generative AI – they are associated with medium or high skilled workers now, as Tera was saying, right. Then, I think it should not be forgotten that these are also the type of workers that usually rely on abilities AI does not currently possess, you know, which are these socio-emotional skills.

BvA: I want to spend the last couple of minutes on talking a little bit more broadly around, sort of, strategy with regard to AI. Also, you know, we have a large UK audience, always talk about the UK, it's great to have Erik on here who's in Silicon Valley and can give, sort of, the US perspective on this. But does a government need an AI strategy and does the UK have one that you think is effective?

TA: I think most governments will need to have policies to make the most of, but also avoid some of the pitfalls or risks around AI. The UK actually has several different publications in that space, including an AI strategy from 2021, and then there was a White Paper on regulation, and various action plans and various other announcements in between.

I think what's quite good about the '21 paper is that it looks at it quite broadly. It's not just about making sure we have the science and the leading-edge businesses that are maybe generating some of the software or the services that might be brought up from our AI base. But it also looks at making sure that there's adoption across sectors and regions, and that the regulation is fit for purpose.

The intentions are good. If I look at what has been announced and what progress has been made, it seems to me like quite a lot of, actually, progress has been made on both the regulatory side, where the White Paper is very clear about, for example, rolling out legislation to be really restrictive about the use of AI. And it's a, sort of, pro-innovation regulatory stance.

And then, quite a lot of progress has also been made on the more scientific front, specifically just making sure that the funding is there for leading edge research and science.

Where I see a little less is focused on this adoption piece, and in particular on the, sort of, more broad skills. So, not the high-end skills that Lea was talking about earlier, which are definitely also needed, but the, sort of, broader skills base in the population, in the workforce, and in the six million businesses we have in the country, to actually just have digital literacy and be comfortable with these sorts of tools and be willing to experiment.

Because, you know, if we want aggregate productivity to rise, it's not enough for us to have a few superstar companies. We need the average, often not very sophisticated, small or medium-sized business to also be benefiting from these technologies?

BvA: Erik, give us a bit of the US perspective, if you can. Because there is a sentiment on this side of the ocean that on the one hand the US doesn't regulate anything, it just happens. But on the other hand, we have, you know, a Biden Government, we have the IRA, and it feels like there are just huge, massive plans and investments into those technologies and that there is a strategy. How do you assess these two different perspectives on the US role in AI strategy?

EB: Well, I was just in Washington, and I had a chance to meet with folks in the White House, the Council of Economic Advisers, the National Economic Council, the incoming Secretary of Labor, some senators. And there is really a sea change in those two attitudes, where people are much more interested and concerned about understanding what AI can do and thinking about what may need to be done to regulate it.

And it's not just among the folks in government, many of the leaders here in Silicon Valley are, as you know, calling for regulation. And part of it is that they seem some very significant benefits, but also some potentially very large harms.

The harms I'd put in three broad categories. The first is around deep fakes, misinformation, disinformation, persuasion. And those could be threats to democracy, or fraud, and I think we all have to be very careful about those, and there are some things we can do to address those.

The second is around economic disruption and jobs. I don't see mass unemployment, but I do see a lot of rearranging that needs to be addressed, and training and safety nets.

And the third is around existential risk really, like large, potentially dangerous uses of the technology to create biological weapons, or cyber-attacks.

All three of them I think are worth paying some attention to, roughly in that order of when there are likely to be issues. Given the nature of this podcast, let's focus on the second one though, around economic disruption. And I think there, we're making some good progress in creating a bit of a roadmap to understand where the technologies are most likely to be affected.

In 2017 I wrote a paper in science, in the American Economic Papers and Proceedings, with Tom Mitchell and with Daniel Rock, where we looked at 18,000 distinct tasks and tried to evaluate where machine learning was going to affect most of them.

More recently, Daniel Rock has done a similar exercise specifically for generative AI. And as Tera was mentioning, many of them, higher paid jobs, are the ones that have more tasks that are likely to be affected.

But with this kind of a roadmap, I think we can give some guidance to government and to, for that matter, CEOs and executives, about which parts of their nation, which parts…which regions, which industries, which occupations, which parts of their company are likely to be affected first and in the biggest way.

And I say affected, because it could be affected in a way that augments the workers and creates more productivity, it could be affected in a way that automates them and drives down wages. And to some extent, that's a choice. But it all starts with having this, kind of, a roadmap, using these technologies to see what's affecting…

I started a company called Workhelix that gives some guidance to companies and governments that want some help in creating that kind of a roadmap. We also have a number of papers coming out of the Digital Economy Lab that describe the methodology for doing that.

LS: Yeah, just to build on that, the OECD has developed some principles on AI, which are really there, you know, to, kind of, give recommendations and guidance to governments. And in relation to assist growth, AI systems should be designed to be human-centred and transparent, you know. There should be accountability, they should be safe and secure, and of course, the OECD encourages that there is cooperation and collaboration among different, kind of, stakeholders.

But I think the important thing is now to move from principles to the implementation. And countries are at very different stages here of developing and implementing different national AI strategies and policies.

And the OECD has a database where it's collecting, basically, different national AI policies and AI strategies and initiatives for over 60 different countries. And things that always, kind of, pop up that are really the, kind of, main priorities are, you know, to improve financing in AI R&D institutions and projects promoting AI uptake by businesses. You know, kind of, equipping the population with the skills for developing and using AI, and, you know, fostering a fair labour market.

But then there are also, kind of, growing priorities, which are all about, kind of, data sharing. You know, either through different incentives, which is…I think the UK are doing, but also through more centralised, accessible repositories.

And so, I just wanted to briefly bring that up, that I think, you know, across countries there is a lot going on.

BvA: Yeah, and I think that comparative base is really important because there are different ways to go about this, right, and it's a tricky balancing between jobs and productivity. And the worry, of course, is that we regulate so much that it will become a productivity barrier.

So, the different strategies, there is this, sort of, roadmap, sectoral strategy, there are strategies that are much more focused on privacy and ethics issues, and also approaches that are more pragmatic. And I think that's what some of you have been talking about, that, sort of, make impact assessments all the time, just to identify the risks that exist, either for jobs or the risk for actually slowing down productivity growth. So, it will be really good to see productivity being part of that story.

Well, a lot more to talk about but we're definitely through the time, and we covered a lot here in this podcast. So, Erik, Tera, Lea, thank you very much for helping us to work ourselves through this complex topic.

I'm sure that you may want to read a bit more about it. Now, we have a lot of show notes here referring to the Brookings report, to some OECD work, and to the McKinsey report. So, please go to our show notes, read things in more detail, and of course stay tuned for any other technology innovation topics and productivity that are definitely going to come up. But thank you for joining us.

EB: Thank you, Bart, it was a real pleasure.

TA: Thank you.

LS: Thank you.

BvA: Our next episode of Productivity Puzzles will return to the issue of regional and local differences in productivity performance. Something we have addressed in earlier shows focusing on the UK.

But this time we will look overseas, and we'll discuss some really interesting work done by researchers at the Blavatnik School of Government at Oxford University, looking at how some troubled cities and regions in Germany, France, Australia, and the US, have transformed themselves.

The work looks at cities like Leipzig, Duisburg, Lille, Pittsburgh, and Newcastle in Australia. We'll hear about what those places did to regenerate themselves and make their places more productive and better places to live.

And for this episode in July, we'll also be joined by my colleague Philip McCann, professor of regional economics here at the University of Manchester.

You can sign up for the entire Productivity Puzzles series through your favourite platform to make sure you also don't miss out on any future episodes.

If you'd like to find out more about upcoming shows or any other work by The Productivity Institute, please visit our website at productivity.ac.uk, or follow us on Twitter and LinkedIn.

Productivity Puzzles was brought to you by The Productivity Institute, and this was me again, Bart van Ark at The Productivity Institute. Thanks for listening and stay productive.

End of transcript