

Newsclip [00:00:01]:

It returns! The thing with almost human brain... Electro the robot.

Tim Phillips [00:00:09]:

Today on VoxTalks Economics, Daron Acemoglu on how big tech is increasing inequality and undermining democracy.

Welcome to VoxTalks Economics the Center for Economic Policy Research. My name's Tim Phillips. Every week we bring you the best new research in economics. So remember to subscribe and follow us on Instagram at VoxTalks Economics.

You might have noticed that digital technology, and AI in particular, is getting a lot of hype at the moment. So in a special episode of VoxTalks this week, I'm going to be speaking to one prominent economist who's on a mission to prick that bubble, arguing that this cult of techno optimism is concentrating power in the hands of an unaccountable super wealthy elite while threatening the livelihoods of the rest of us. Last week, Daron Acemoglu was in London to talk about his new book, written with his colleague at MIT, Simon Johnson. It's called Power and Progress. It argues that we need to reshape the path of innovation if we want to harness it to create shared prosperity. So when I caught up with him, I first asked if he had any idea that AI was going to become the thing that everyone was talking about at the moment, exactly at the time that he publishes a manifesto for change.

Daron Acemoglu [00:01:38]:

Yes and no. We started writing this book more than three years ago. I wrote an article in Washington Post which was entitled: The AI We Should Fear is Here. A reference to the fact that it's not super intelligent killer robots, but how we are using digital technologies today that we should be concerned about. But genuinely speaking, I am impressed by some of the things that Chat GPT or GPT Four does. On the other hand, I'm very concerned about the direction in which it's being used.

Tim Phillips [00:02:06]:

I suppose you better explain what the book is first of all, called Power and Progress. Explain what power and what progress?

Daron Acemoglu [00:02:14]:

Power is about who wields social power in society. That's everywhere. Do you get more wealth, more status than me? That's the question of power. But the one that we really are concerned with is who has the power to set the direction of future technology and how we distribute the

gains from it. And progress is about technological progress on the one hand, and whether we are truly progressing as humanity.

Tim Phillips [00:02:41]:

It's a very directly worded book. You've got quite strong opinions on this. And the strong opinions, there are a lot of things that we're doing with technology aren't going in the right direction at the moment.

Daron Acemoglu [00:02:51]:

Oh, I failed and I was trying to hold back. Look, I think what we are trying to achieve is twofold. One, we want to counter our narrative, and that requires being direct. And that narrative is a form of techno optimism, both in the tech world, Silicon Valley and beyond, and in academia, that somehow automatically anything that improves our capabilities will bring benefits for most people. Quality of life will improve, better products, higher wages. We want to add a note of caution to that because in the past that has sometimes worked out. We are fortunate, enormously fortunate to be living today rather than 300 years ago but there was nothing automatic about it. And second, we are trying to broaden the conversation about who gets involved in having a say over the future of technology because we firmly believe, again based on our reading of history that democratic control of technology is critical.

Tim Phillips [00:03:56]:

Your reading of history in this goes back 1000 years. One of the things that's, quote from the book here, the extent to which productivity gains are shared between capital and labor are not inescapable givens. They are societal choices. What comes out of this is that the choices have been made in different ways at different times, with different outcomes. Quite often, though, it has been to the benefit of a small elite. And that's what you think is going on now?

Daron Acemoglu [00:04:24]:

Absolutely, that's what I think what's going on now. And to understand that and also to evaluate why in general within academia, for example, in my profession of economics, including some of my earlier writings why there is a sense that things are going to work out for the better. We need to understand the different forces. And what we call productivity bandwagon in the book is critical for this because most of us earn our living via selling our labor, work. So if there is going to be shared prosperity there must be some process via which technological improvements increase wages. That's the productivity bandwagon. And the reason why economists believe in it, and I did, is because we think technological change increases average productivity. How much output you can produce with a given amount of labor. And when that's higher, firms go and hire more labor or want to hire more labor. And that pushes up labor demand and pushes up wages. What we argue in the book is that there are two big assumptions for that. These are

not always satisfied. One is that average productivity increases. We produce more, but marginal productivity of labor may not improve. And the reason for that would be if we use technology for automation. We are very productive in having a factory that doesn't have any workers. That's superb for average productivity. Big number divided by zero. But it's not great for the contribution of workers to productivity and employers are not going to rush to pay more to workers. And second, institutions, even if workers are more productive, employers may try to keep more of that for themselves. So bargaining negotiations, what the general norms are central for who will benefit from technological progress.

Tim Phillips [00:06:16]:

What changed your thinking about this? When did you start to think that your earlier work was painting too rosy a picture of technology?

Daron Acemoglu [00:06:25]:

Well, a lot of it was a process about 15 years ago I started working on models that have automation in it. And that's one turning point. And second is I've always been interested in institutions. I've always approached the institutional aspects from a lens of who wins, who loses, issues of power. And some of my work reflected that already in the labor market, but I hadn't put everything together. For example, we assume that labor relations are not coercive. But if you look at history, the modal form of employment is not a free market system, but it's one in which employers have power, sometimes violent power, over their workers. Slavery, feudal relations, servile labor, debt peonage, all of them are, even today, not unheard of.

Tim Phillips [00:07:20]:

There's a sense of urgency about this book as well that things have to change. We'll get into that in a minute. But you do think that what you call the oligarchy today, the people who are doing very well out of technology changes, are bad for society? In what ways are they bad for society?

Daron Acemoglu [00:07:42]:

So let me try to unpack that. We talk of a vision oligarchy. So it's not the economic oligarchy that I'm concerned with. And what does that mean? If you start from the premise that there isn't a preordained path of technology which is so important for the techno optimistic view no, we choose the path of technology. We choose the direction of technology. Then the question is what determines that direction of technology? Clearly, economic forces matter. But Simon and I argue it's not just economic forces. It's also the vision of the powerful actors. Do you think you should use technology to empower people or disempower people, to make them more democratically amplified or sideline them? Those visions matter. They have always mattered. And they matter even more today because a small group of people have an oversized voice,

over debates on the direction of technology. And that's what we call the vision oligarchy. It is remarkable how homogeneous the background and the worldviews of the most powerful people who are shaping the future of technology are.

Tim Phillips [00:08:47]:

I assume that they would say that they are doing good for the planet and that they have an insight that perhaps is not widely shared.

Daron Acemoglu [00:08:56]:

Well, that's the techno optimistic view. Here is a corollary to the techno optimistic view. And I'll come to that because this corollary is actually very important for thinking about the broader set of issues. You can envisage a world in which progress happens. Thus there is a small group of very talented people engineers, scientists, geniuses. They design technologies and the best way of society to progress that those technologies to be used unquestioningly by the rest of society. That is a vision that's very common in Silicon Valley. It has its roots both in corporate power and in how ideas going back to Alan Turing have articulated what we think about computers, autonomous machine intelligence. Those two ideas of corporate power and autonomous machine intelligence meeting the genius sort of approach gives us a very specific vision. And of course, if you ask these people, they say you should all be amazingly grateful to us. Our contribution is beyond the billions we earn. That's why we have such status. But even that's not enough. I'm questioning that. Of course, there are some very, very smart people in Silicon Valley, and the creativity in that ecosystem is amazing. But it still doesn't mean that every decision they make is the right one. That's where democratic control comes in. And Simon and I are taking a different view. We are departing from autonomous machine intelligence. We argue what we should want from the machines is not their intelligence. That was completely the wrong way to frame the question. We should want machine usefulness, machines being useful to humans. Why does that make sense? You need a different way of thinking about the problem. Again, it's an empirical matter as well as a philosophical matter. But my view is that humans are amazingly versatile. They're very skilled, they're very adaptable, they're very creative. Their real strength is their diversity of skills. And there is a way to amplify and improve those skills using technology. That's what machine usefulness should do. Calculator, it doesn't replace me. It's not intelligence. It's amazingly useful. Wikipedia it distills in the right way, in my opinion, human wisdom, and presents it to us in a digestible way that we can actually go and check. It's not authoritatively pushed down our throat. So those are the kinds of things that we should strive for, not this autonomous machine intelligence. That vision is leading us astray.

Tim Phillips [00:11:19]:

Your argument about that vision is that it's all about replacing humans?

Daron Acemoglu [00:11:23]:

It's about replacing humans and disempowering humans, most humans, not the people who are designing the technology. H. G. Wells got the gist of it in the Time Machine in 1895. He said, technology is not just about human dominion over nature, it's about human control over outer humans.

Tim Phillips [00:11:40]:

And of course, what happens in the time machine is we get two entirely different races of humans a master race and a subject race.

Voiceover [00:11:53]:

Daron and his co authors have published many articles on Vox EU on the impact of technology both at work and for society. For example, in March 2023, industrial Robots on workers, winners and losers. In November 2022, automation and polarization. And in November 2021, dangers of unregulated artificial intelligence.

Tim Phillips [00:12:21]:

However, there is again that techno optimistic view of the future, that we could be freed from work. If there was just the money around that we could live, then why would we have to work? Why do you want to protect jobs so much?

Daron Acemoglu [00:12:36]:

Great question, and it's a difficult question, so let me take another one first, which is that what's wrong with techno optimism? And what's wrong with being just alarmed about killer robots? That they both pacify us? Yeah, I think we should be concerned. You're right to pose the question that way, we should be concerned. But that concern should be translated into a positive agenda. An agenda in which people have better knowledge, better understanding and a voice over the direction of technology. Part of that is the question that you ask why do we care about work? Well, three reasons. First, unless we completely change our institutions, that's where we earn our living. And if we destroy work, it is going to really go towards a much more two tiered society. Allah HG. Wells. So I'm really worried about that. No, I don't believe that we're going to build better redistributive systems that could make up for the lack of income growth for the poor if we get rid of work. Second, I actually think that the direction of research in AI is distorted. It can be better if we amplify human capabilities. So even leaving philosophical and redistributive issues just from the pure productivity point of view, we are better off trying to complement humans, creating new tasks for them. Let's think about this this way. Just let me make two arguments in that regard. One is everybody in Silicon Valley and many people in the US are talking about these amazing technologies. The number of patents in the United States has increased five to sixfold over the last 40 years. We have a new gadget every day, but

productivity growth is as low as it has ever been. We're not making good use of it. And when you think about it, it's very difficult to increase productivity just by automating. Even on the most rapid automation. The US. economy is going to automate perhaps 3 / 4% of tasks and occupations that it performs. Imagine how much better than humans machines need to be for that to translate into 1 or 2% productivity growth at the aggregate level. Instead, if we make humans a little bit more productive, that's a big, big gain. And third, I actually think that a society in which people don't work would be a dystopian one. How do we define our social networks, our sense of worth and belonging? In fact, I think and Michael Sundell is onto something here in the tyranny of merit as well. Economics has gone a little bit too much down the Rawlsian path of thinking about justice just in terms of consumption. Of course consumption is important. So there are aspects that Rawls of course, had it right. But we should also add to it a contributive sense of justice that people perceive and are recognized for contributing to society. If it's not work, what is it then? Are we going to be contributing to society by playing games in the virtual world? So I think there is a conundrum there.

Tim Phillips [00:15:30]:

Look at it from the other end of the telescope and I'm a Chief Executive. If I am a tall techno optimist, I'd be looking out for, for example, that machine that could program really well because as soon as I have one, I have as many as I need. Programmers don't move around. They need to be trained. They need to be recruited. The problem with making machines that just make humans more productive is you then have to employ humans, and that's a really difficult thing to do.

Daron Acemoglu [00:15:58]:

Absolutely, you got it perfectly right, Tim. That is the bias of the corporate world. Once you go down the rabbit hole of shareholder value maximization, thanks, Milton Friedman. Then you start thinking, well, why should I try to make humans more productive? They are troublesome. They will bargain for higher wages. They will make demands, and I have to pay some of that. So instead, let me automate and that vision, that dream has been with businesses throughout the ages. Even during the Industrial Revolution, self acting machines were very attractive to employers. Already in 1950, American magazines were talking of the workerless factories because of the numerically controlled machinery that was being rolled out at the time. And office automation was hailed as something that was going to create high productivity because you could get rid of your white collar workers. Industrial robots did that for blue collar workers. So that dream has become partly a reality. AI threatens making that more of a reality. And if you are the manager, that might be a good thing. Not for the workers, not for society at large, and again, not for productivity, because we could do better with technology by complementing humans, by creating new tasks for them.

Tim Phillips [00:17:04]:

One of the ways in which we might realize your vision of what a good future would be is by giving more people a voice in the direction of technology. And one of the quotes that stood out from the book was from Norbert Wiener, who was at MIT, where he was talking about robots are effectively slave labor. He was saying, any labor which competes with slave labor must accept the economic consequences of slave labor. And that seems to be the direction of travel that you're worried about.

Daron Acemoglu [00:17:34]:

I am.

Tim Phillips [00:17:35]:

So the first thing I guess we have to do is to find a way to give affected workers a voice. And that's very different to how it was, say, in the 1960s or the 1970s. We don't have that kind of union organization anymore. How are workers going to get a voice now?

Daron Acemoglu [00:17:54]:

Absolutely. I think that's a critical question. It's not just workers. Democratic control means society at large. By the way, thanks for bringing Norbert Wiener. I think he was really ahead of all of us writing about these things in the late 1940s. He did not use the term machine usefulness, but we credit him with those ideas as well. So it's great to remember him, but how are we going to do that? By the way, Norbert Wiener also understood that workers had to have a voice for that to happen. Yes. I don't think old style unions are going to come back. And if you look at the history of unionism in the UK and the US. It's not one of great success.

Tim Phillips [00:18:34]:

No.

Daron Acemoglu [00:18:34]:

Why not? Because I don't think with a few very important exceptions, I don't think union leaders understood that they had to bargain not just over wages, but about organizational work, training and the direction of technology. And that was very difficult in the US. It was very difficult in the UK. Because much bargaining was a study at the establishment level was very conflictual, but it took place in the Nordic countries in Germany. But even those models are going to be challenged because they were centered on blue collar industrial workers, and there are fewer of them. So we need to find new organizational modes. But there are models. You know who gets the issue of technology? Writers guild. This writer's guild strike is exactly about the right issues. They are talking about AI, data, creative work. Their concern is, can the studios take old sitcoms and old shows, throw GPT four or GPT eight at it, whatever it is, and come up with boring,

canned, but still profitable shows? And if they do that, who gets the profits? Now, the problem is, if they do that, they're not creating productive stuff. Human genius is about creating new things. You're not going to achieve that by taking old reruns of Friends and making them slightly better. You need new creativity. We can use the AI tools for that. But control of data negotiations between creative workers and studios going to be very important. More broadly, let's translate that to the service industry in Amazon warehouses. Who controls workers? Worker autonomy, I think, is part of the problem. When you have ceaseless monitoring and a complete control over worker time by the bosses, that's not going to lead to a good environment.

Tim Phillips [00:20:27]:

One of the topics that comes up regularly when we are talking about slowing down the adoption of new technology is taxing it. Robot taxes were in vogue at least discussed five years ago. Could we have AI taxes to make them less attractive as a replacement for humans? It makes humans comparatively cheaper.

Daron Acemoglu [00:20:50]:

There are a number of issues there, so let me take them up one by one. First of all, anytime you're talking about slowing down technological change, you have to be very careful. Our political track record of blocking technological change is not good, to say the least. It's awful. Why Nations Fail, the book that I wrote with Jim Robinson, had many examples of disastrous blocking of technology because of interest groups. We don't want to repeat that. That having been said, let me ask the following question what is the cost to humanity if the adoption of large language models and development of large language models is delayed by six months? I don't think anybody could make that argument that it's huge.

Voiceover [00:21:35]:

An open letter signed by Elon Musk is calling for a pause on the development of more powerful artificial intelligence systems, citing risks to society...

Daron Acemoglu [00:21:45]:

That's why I was one of the early signatories of the call for halting the training of large language models for six months until new understanding regulatory framework is developed.

Voiceover [00:21:56]:

Calls on developers working on systems more powerful than OpenAI's, newly launched GPT Four to pause their work for six months to make time for safety.

Daron Acemoglu [00:22:08]:

I love technological progress. There's nothing more that I enjoy and cherished than science. But AI is going too fast and it's on a wrong path and we just don't have the regulatory framework or understanding of where it's going to be able to respond to it. So that what worries me. So reluctantly I signed it because I think we need a wake up call. And I wrote this book because I think we need a wake up call.

Tim Phillips [00:22:35]:

Pauses for six months struck a lot of people that I spoke to in a technology business as not remotely practical. No one's going to stop. They're too enthusiastic.

Daron Acemoglu [00:22:44]:

I know, but we need to develop the alternative perspective. I knew that they wouldn't stop it for six months, but just building that broad coalition and I don't think stopping it is the right solution. Are AI taxes, which you brought up, the right solution? I'm not sure about that, but let me tell you what one thing that I think would help right now. The UK, US and many other countries tax code already excessively subsidizes capital and digital capital at the expense of labor. We impose income taxes, payroll taxes and lots of other taxes on companies, directly or indirectly, via their workers. If they hire more employees, we subsidize them. When they adopt capital, they're often adopting that capital to replace the worker. So we're subsidizing automation. So why not level the playing field and equate the marginal taxes on capital and labor?

Tim Phillips [00:23:31]:

Would it work as well to break up big tech?

Daron Acemoglu [00:23:34]:

I do think so. I do not think that it's a holistic solution. I don't think it would be enough by itself. But yes, I think it would be a good idea because these are the largest corporations we have seen in humanity's history and they are very powerful both politically and socially, leave alone the economic power. They just set the agenda. They can acquire anybody they want, they can influence politicians, they set all of the policy priorities. So I think as a step in reducing their grip on power, breakup is something we should consider.

Tim Phillips [00:24:08]:

Finally, Daron, you talked a little bit at the beginning about what economists are like and how they view this as well. If you look at the economics textbooks, the two words that are in your title, Power and Progress, they don't really talk about power at all. When they talk about technology, it's in a kind of undifferentiated thing that gets you from not very good to much

better. Should we be thinking about how we teach economics?

Daron Acemoglu [00:24:35]:

Yes, I think we need to bring power much more into it. Politics much more into it, institutions much more into it. And when we talk about technology, yes, I think it both at the graduate and the undergraduate level, the view that technology is something that just makes labor more productive, that's a too reductionist view. So we have to think about what is it that technology is doing? Is it providing you better information? Is it automating? Is it creating new tasks? So there are things to be criticized in the economics approach and the economics curriculum. On the other hand, let me say that everybody is working on the consequences of AI and everybody is thinking about how to think about technology, from computer science to sociology to political science to philosophy. And I think economists are very much at the frontier. Some of the best work understanding the consequences of technology and discrimination and inequality in labor markets, in power is coming for economists. Economics has really established itself as a conceptually innovative, empirically minded profession. But there is a lot of baggage as well. So I think we should celebrate the things that economics is getting it right, but we should also try to modernize some of the things in which we are somewhat backward looking.

Tim Phillips [00:25:49]:

Ideas of power and of the uses and abuses of technology. They would have been quite familiar to 19th century.

Daron Acemoglu [00:25:54]:

Absolutely. I mean, political economy was about that, and through political auto economy. On the other hand, if you look at what young scholars are working on today, many of them are working on political economy. Now, all leading economics departments have political economy courses and fields. There's a lot of new work coming on politics, institutions, culture, struggles. I think that's to be cherished.

Tim Phillips [00:26:20]:

Have you had any feedback from tech CEOs yet?

Daron Acemoglu [00:26:22]:

Not yet, but I have spoken to them in the past and some of them are actually quite receptive. But of course, the view that you don't understand it, this is all for the better, that's very deep rooted. And another view, which is who are you and who are the regular public to talk about AI? We are the ones who understand the algorithms. And my answer is, of course, I don't know how to design better AI algorithms than you guys, nor do I know how to build nuclear weapons. But

that doesn't mean that I shouldn't have a say on whether nuclear technology should be used for peaceful means or weapons.

Tim Phillips [00:26:56]:

Daron, thank you very much.

Daron Acemoglu [00:26:57]:

Thank you, Tim. This was great.

Tim Phillips [00:27:08]:

The book is called Power and Progress our 1000 Year Struggle Over Technology and Prosperity. The authors Daron Acemoglu and Simon Johnson. It will knock your socks off.

Outro [00:27:25]:

This has been a VoxTalk from the Center for Economic Policy Research. We bring you the best new research in economics every week, so subscribe wherever you get your podcasts, we're there. And if you like what you're hearing, please leave us a review. Next time on VoxTalks when researchers collaborate, are mixed sex teams more or less successful.