THE BROOKINGS INSTITUTION

WEBINAR

THE PROMISES AND RISKS OF ARTIFICIAL INTELLIGENCE: A CONVERSATION WITH MIT INSTITUTE PROFESSOR, DARON ACEMOGLU

Washington, D.C.

Monday, December 13, 2021

PARTICIPANTS:

ANTON KORINEK
David M. Rubinstein Fellow, Center on Regulations and Markets
The Brookings Institution

DARON ACEMOGLU
Institute Professor, Department of Economics
Massachusetts Institute of Technology

* * * * *

P R O C E E D I N G S

MR. KORINEK: Hello, I am Anton Korinek. I'm the David M. Rubinstein fellow at the Center at Regulation and Markets at Brookings, as well as a professor of economics at the University of Virginia. It is my pleasure and honor to welcome all of you, and especially you, Daron, to today's fireside chat on the promises and risks of artificial intelligence.

Let me also say thank you to the team that has made today's event possible, Sanjay Patnaik, head of the Center on Regulation and Markets that is hosting the event, and Megan Warring (phonetic), who has put a tremendous amount of work into the preparations.

We have seen tremendous progress in the field of artificial intelligence over the past decade. In fact, many view AI as the defining technology of the 21st century. Now, every great technological revolution comes with promises and risks. Some predict that AI will usher in an age of unprecedented growth and prosperity that smart machines will allow humans to expand economic output and living standards at a pace we have never seen before. Others point to risks such as the displacement of human workers leading to greater inequality or the risk that AI will reduce competition, privacy, and consumer choice. And, in fact, it's not an either/or. These promises and risks may materialize together.

Our guest today is Daron Acemoglu and he will help us shed light on these questions and on what policy can do to build on the promises and mitigate the risks of AI. Daron is the Elizabeth and James Killian professor of economics and also an institute professor at MIT. He was awarded the John Bates Clark Medal in 2005. Among economists, that's a greater honor than a Nobel Prize because it is only awarded biannually to a single person in the field.

By some measures of citation counts, Daron is actually the most influential economist in the world. I was honored that Daron recently wrote a paper on the harms of AI for a handbook that I'm editing for the Oxford Handbook of AI Governance. In this paper, which is also available on the Events page from which you are viewing this conversation, he provides a broad overview of potential risks or harms generated by AI.

I thought it may be worthwhile to start our conversation today by looking first at the promises of AI, then at the risks. And in the course of doing so come to a synthesis of the two.

Incidentally, if you're watching this event live you can tweet any questions that you have under #FutureofAI and I will do my best to incorporate them in our conversation.

Daron, let's start with the promises of AI. May I ask you to explain the mechanism through which technological progress in general leads to higher living standards? For example, how have all the technological advances that we have seen since the onset of the industrial revolution contributed to increasing our prosperity?

MR. ACEMOGLU: Well, first of all, thank you, Anton, for inviting me to this fireside chat. It's fantastic to be joining you and all the members of the audience. It is an important topic and I'm delighted to be part of the conversation.

So let me give a two-part answer to your question, because I want to first clarify how economists think —- or used to think —- about technology and then how AI fits into it.

So I think if there is any area of agreement among economists, then technological progress being a force for good is a pretty good candidate for it. Economists normally think of technology, I think very much along the right lines, as something expanding our capabilities. So if using a given amount of input, labor, material, we were able to produce ten goods and technology enables us to do now twelve, that is obviously a big progress. And I think that really is the essence of the simplest ways in which economists approach the problem. Technological progress directly translates into productivity increases and those productivity increases are good for society and then they will trickle down to factors or production, such as labor.

You can think of this as a neutral technological progress, meaning that it is expanding our capabilities without changing any sort of deep balances in society. And it is more or less the staple form in which economists write down the effects of technological progress in their textbooks and their models. But of course many people have realized that this is a simplification.

Luddites understood very well that technological progress wasn't something neutral that was going to benefit everybody with some simple process of trickle down. They were organized artisans in the textile industry at the beginning of the 19th century and they were very much threatened by the introduction of spinning machines that would take away their fairly well paying jobs. They put up quite a bit of resistance, but at the end the government suppressed the Luddite revolt and their ability to break

the frames and the machines was limited and in history now Luddite is sort of a symbol of people who don't understand the strength of the march of technology.

But actually that's a very uncharitable view of the Luddites because I think they had a much more nuanced understanding of what was going on and in some ways they were right. It's not like they were completely off the mark.

Other prominent people who actually understood how technology could be very non neutral was the great economic John Maynard Keynes. In 1928 — or '29, sorry — he gave a lecture at Cambridge where he pointed out, or forecasted, very patiently that technological change would continue to increase our measures of output, GDP per capita, but more or less 2 percent a year. But he thought that this would be very non neutral and it would create a term he coined "technological unemployment" and this would be quite a sweeping change which would require huge social adjustments. Despite the fact that Keynes is still many people's favorite economist, the post-war economics profession concluded Keynes was wrong partly because the data did not bear it out.

So the question is was Keynes wrong in general or whether there were other changes in technology or how we use the technology that turned out to create a more neutral and more beneficial improvement. So I think that is a critical discussion that goes beyond AI, but AI makes it very relevant.

Now, if you put AI into this context, of course AI is expanding our technological capabilities. So if we use it well it can bring huge improvements. But just like what Keynes thought and Luddites understood, AI is highly non neutral in ways that we're going to go into. It benefits some people, it disrupts other people's jobs, it disrupts balances, it could create misery and it could create huge power imbalances in society.

But before we get there, let me conclude by saying I am certainly not anti AI and I see three very distinct, very important broad areas in which AI can help, which actually are very much areas in which other technologies have helped as well. First of all, AI is about more data and better data. And many problems that we face, such as cancer, pandemic tracking, recommendation systems are all about data. So we can improve many of these aspects of our economy.

Second, I think as we move into the heart of the 21st century there's going to be an increasing understanding that we need to reduce our footprint on the environment, not just for carbon but

a whole host of materials and other things that we are using. AI is pushing us toward more services that if designed well can replace a lot of the things that are based on materials and consumption of resources. So AI could be a very important part of a post resource economy.

And then, third, I think AI so far has shown its greatest — I would say it's only — success in helping decision making and human decisions are seriously flawed in certain dimensions. It's true when we are consumers, it's true when we are parents, it's true in our entertainment choices, and it's definitely true in our business decisions. And AI can help businesses make better decisions, better hiring, better strategies in terms of which products to market and how to price products. It can help consumers for certain decisions, like in the recommendation systems I mentioned. So in all of these ways there are a lot of promises. But I will also argue in the rest of our hour together that none of these promises are being realized, and instead of promises we have a lot of perils.

MR. KORINEK: In your paper you emphasized that there is a risk that some recent efforts by the AI industry don't actually contribute to growth in the way that you just described, but they may instead undermine the livelihood of workers without much benefit in terms of productivity gains. We call them "so-so technologies".

Can you explain that a bit more?

MR. ACEMOGLU: Thank you, Anton. Absolutely. That's a topic very close to my heart, I know also yours.

And I'll just pick up where I left off. So I emphasized that economists most optimistic read of the benefits of technology won't work out so nicely if technology is not neutral and it disrupts certain existing balances. I think one very important aspect of this is what technology does to labor. What Keynes was worried about is that technology would increase productivity, but at the same time it will be labor saving, meaning that it would reduce the labor requirements and hence employment.

So this is a type of technology we call automation today. So the spinning and weaving machines at the beginning of the industrial revolution, such as the spinning machines the Luddites were so opposed against, are clear examples of automation. They were taking tasks that humans used to perform and allocating them to capital. Mechanism of agriculture with tractors and harvesters was another example of automation. In the course of about 80 years, labor requirements plummeted and from

50 percent to less than 2 percent of the population of the United States became to be employed in agriculture. There are many examples of automation technologies in manufacturing, but lately key examples of manufacturing are digital based robots and AI based algorithms. In offices, in sales, clerical function, in back offices we see all of these examples.

Now, some amount of automation is a fact of life and could actually be a very good thing. In much of the 20th century rapid productivity growth when hand in hand with automation, but there were two reasons why automation worked out to be reasonably okay for workers and for the economy. One is that it was high productivity automation. So you automated certain tasks, but that improved productivity and that in turn enabled firms to expand production in non-automated tasks, sucking some of that labor back. And second, it went hand in hand with other technological changes that created new tasks that were labor intensive, so reinstated labor due to production costs.

Now, unfortunately, despite all of the promises on behalf of AI, that's not what we are seeing today. AI is not generating very high productivity automation and it is not being used for generating new tasks that reinstate labor. So we get a double bad whammy. Now, that will surprise people perhaps for several reasons. First of all because we are told by AI enthusiasts how amazing AI is. So you would expect huge productivity. And second, we are told by McKinsey and all these consulting groups and AI companies that they are creating jobs as well as replacing jobs. So unfortunately neither of these seem to be true.

So let's try to unpack. Why is it that AI is not generating huge productivity improvements? Well, actually the answer is very simple. Because AI is driven by a particular ideology that thinks machines and algorithms are superior to humans. But actually in reality humans are pretty good. So AI is doing things that humans are very good at, facial recognition, language recognition, pattern recognition, problem solving, identification of anomalies, personal advice. All of these are things that humans excel at. Of course we make mistakes. We are fallible. But we have 200,000 or more years of evolution on our side with amazing brains that help us do that. And AI, despite all the promises, is actually not working so well.

So take an example, customer service. That is a hugely important part of our economy and it's actually a perfect job for humans. Other humans like interacting with humans, so talking to an

individual on the line is very important. Most of the people who actually call for customer service have a non-routine problem. So you can't just say, okay, your system is down, I'll restart your system. So you have to recognize what the problem is, understand it from imperfect description, and then find a solution to it. Humans are very good at that.

And now let's see what AI based systems are doing. They have a rule book. Despite the fact that it's supposed to be intelligent, it's actually pretty dumb. Every time you call they go through the same sequence whatever you tell them. It wastes a lot of times, it often doesn't resolve the system, it frustrates the customer, but it saves on labor costs. So this is the so-so technology. It's replacing humans, it's doing things that humans used to be very good at, it's doing it badly, but it's marginally profitable for companies and our shareholder value obsessed cost cutting companies jump on it. And perhaps many of them are actually jumping on it even when it's not cost cutting because they have just drank the Kool-Aid and think that's the technology of the future, we have to do it.

So rather than use AI for doing something more useful, like increase the productivity of their workers, this is what they're going, they're doubling down on this. And that's terrible. We're not getting the good kind of AI, we're not getting the productivity growth, we're not getting the reinstatement, and we're getting just the disruptive effects that Keynes was concerned about.

MR. KORINEK: You and I have both written about this need to steer technological away from these so-so technologies. What policy measures would you recommend to an audience of policy makers to steer technological progress in that way? How do we put this into practice?

MR. ACEMOGLU: Well, look, that's a very, very important but a very difficult question. I am a huge believer that innovation, one of the most important activities, is one that brings a huge number of social effects externalities, distributional consequences, and therefore it's not something you should just say blanket, oh, the market can take care of it. But regulating technology is very difficult. In the past many attempts to regulate technology ended up blocking technological progress, and that's certainly not what any economists would advocate nor would I ever advocate that. But it is also critical to redirect technological change. If that sounds abstract to you, let's take a concrete example -– climate change. If you let the energy sector technological changes to the market, today we would be in even a worse situation than we are right now because the market would invest in what's most profitable, and what's

most profitable was for many, many, many decades -– at least since the Ford and others decided that electric cars weren't the right idea and they had to go through the internal combustion engine and oil became a cheap source of energy and coal became a very practical source of energy -– has been fossil fuels. And there has been a huge number of patents, a huge number of innovations in fossil fuels, and that's what we would have continued. But it's also clear to anybody I think who has any knowledge of this topic, that fossil fuel creates a huge number of externalities, a huge number of negative effects -– pollution, but even more importantly global warming. That's a worldwide problem.

So how are we going to solve that problem? If you think that the solution will be for us to reduce our energy consumption, that's nowhere. That's not going to happen because reductions in the level of energy consumption that would be required to reduce emissions to something that could stay even in the neighborhood of the very non ambitious target of the Paris Agreement would be a huge reduction in industrial output. No country is going to accept it. But fortunately we have an alternative, which is transitioning to clean energy, green technology, renewables. For example, hydro, solar, wind. But if you look at the data from the 1980s when incontrovertible evidence that climate change caused by fossil fuels was a problem, all of these renewables were two orders of magnitude more expensive, more than an order of magnitude more expensive than fossil fuel based energy. No profit maximizing company by itself would do it, but society had to do it. So that's redirection of technological change. And how did we do it? We did it through three levers.

First, in some countries, including parts of the United States, we started pricing carbon. That is an indirect way of redirecting technological change. You change the input preferences of firms and that trickles up to changes in innovation. The second is consumer pressure. Consumers, sometimes employees, started putting pressure on companies to clean their carbon footprint. And that again is an (inaudible) effect on innovation. You have to clean up and you recognize the only way to do that is by technological investments, or it makes technological investments in green technology more attractive.

But there's a direct way of doing it, and that's been actually the most important one –– subsidies to clean technology. So the U.S. government, European governments, Chinese government, pour billions of dollars subsidizing green technology. And thanks to them, today actually if you removed –– if you don't remove subsidies to fossil fuel, which are actually quite of the order of $5 billion a year,

around the world, actually solar and wind are cost competitive with fossil fuel on almost everything. If you remove those, they're actually cheaper. It's the untold story of success and it's a redirection of technological change success.

So we can too and should do the same thing when it comes to AI as well. We should redirect technological change in AI in a more human friendly, in a human beneficial direction. The problem is it took a while for us to work out how to do this in the case of climate change, even though I think it was a simpler problem because at least we had a metric, you know, carbon content. So we worked out a measurement system for that and that's the basis of carbon prices and classifications for clean technology.

So we have to do the same in the case of AI, but I think before we get there, which is a difficult problem, there are lots of other things we can do. And the most important one is that we actually already have many artificial subsidies to the bad use of AI or bad use of technology. So, for example, if you look at the U.S. tax code, we tax labor, we subsidize capital, so much so that at the margin there's a 20 percent difference between producing with capital versus producing with labor. So that means that a firm that decides to fire workers and hire capital instead will actually get a subsidy — effective subsidy of 20 percent. So the first thing is to actually get rid of this asymmetric taxation of capital.

Second thing, which I think perhaps we'll talk about in much greater detail, I think we are not producing AI technologies that are profit maximizing. We are producing AI technologies that a handful of companies have decided are the technologies to bank on. They have their own big war chest of cash and venture capital heft behind them and they're doubling down on it. So it's really a handful of companies in the U.S. and a handful of companies in China who are determining our future. Our technology — so economists are right, technology is the most important part of our future economy — so our future is being shaped by perhaps 200-300 people in a dozen companies. That is just an unacceptable state of affairs to anybody who thinks about it.

So that's the other thing that we have to deal with. We have to try to democratize the decisions of how decisions about the future of AI are being made. More companies, more perspectives, more researchers, more diverse viewpoints have to be part of that conversation.

MR. KORINEK: Thank you, Daron.

And I think you also have a paper that came out in the Brookings papers on economic activities two years ago where you specifically look at the effects of tax policy on influencing the direction of our technologies, right?

MR. ACEMOGLU: Correct. Exactly. And that 20 percent number I just gave you comes from that paper.

MR. KORINEK: Now let me turn a little bit to the longer run. As the progress in AI continues, some predict that machines will be able to perform all human labor more cheaply than humans, let's say if wages cannot fall below subsistence levels. For example, at the survey at the NBER Economics of AI Conference in September, the median prediction for this form of human level AI was in the 2070s, about 50 years from now.

So let me ask you two questions related to that. What are your thoughts on the long-term future of work and in the spirit of being ready for potentially catastrophic scenarios, what can economic policy in the present do to be prepared?

MR. ACEMOGLU: Well, thank you, Anton.

First of all, I wasn't in that I conference this year, so my data is not part of the median of 2070. And if I had to say, I would certainly put that number past 2200. So I think it is true, AI is making progress in performing certain tasks better and it is leading to replacement of labor in those tasks. But those are very narrow tasks and they're very few. Their effects on the economy are still large because if you replace workers in 2 percent of the tasks in the economy, that's still a huge change. But we are talking 2-3 percent for AI, not 20 percent at the moment.

Can that accelerate and become 30 percent, 40 percent, 50 percent in the next decade? No. Because contrary to what many experts tell us, I think the current architecture of AI and where all of the money goes is actually completely unsuited to what is called the artificial general intelligence. So we are capable of using AI, huge amounts of data, and a lot of processing power to perform some narrow tasks, but many of the things that people do in their jobs, and humans are very good at it, again I will add, is much more complex than that. So we can take away the dispensing of cash and account information from tellers, but there's a lot of personal aspects and problem solving that tellers do that's going to be much harder to take away from them. We can take away homework help and some monitoring activities

from teachers, but I shudder with fear if AI will ever replace teachers, because I think that will be the end of educating our youth.

So for that reason, I think unless we crack artificial general intelligence, there is no risk that AI is going to replace humans in more than some high but still less than 50 percent of tasks that humans do at the moment.

Can we have a shot at cracking artificial general intelligence? Absolutely not on the current path. The current path absolutely has no —- not even a fighting chance of getting to artificial general intelligence. Look at all of the success stories that are pedaled in the press, Alfa Zero, AlfaGO, but they are very far from anything to do with artificial intelligence. They are programs tasked with a very, very, very narrow thing and very closed system —- chess and go have rules. Most human activities operate in environments that don't have rules. Rules are fluid and changing and involve synthesis and analogies and extrapolation of very limited information in order to make tentative decisions that are dynamically updated. Those are not the kinds of cross domain updates and extrapolations that the AI programs are doing. And there is I think no hope of the current approach. And there are some groups that are working artificial general intelligence, but I think they are lost and they're not really making much progress.

So, given that, I think, sure, in 100 years' time perhaps things might change, but the chances are not great that we're going to crack this problem any time in our lifetimes, our children's lifetimes, or grandchildren's lifetimes, or their children's lifetimes. But that does not obviate the need for the second part of your question. We still need to be prepared. There will be jobs that are going to disappear. So if your son, let's say, wanted to become a welder, loved that job, you should probably still tell him that it's not a great job because robots are doing all the welding jobs and there are not going to be that many welding jobs left. If all you wanted to do was accounting, well, again, that's probably one of the areas where AI is going to make a lot of progress because that's very simple problem in the grand scheme of things. You know, use the tax code, et cetera. So we have to recognize what are the tasks in which AI can substitute for humans and what are the tasks in which they cannot and what are the tasks in which it's going to complement humans, especially if we make the right investments.

So I think we have to be prepared, prepare the next generation, prepare ourselves, taking

that information of what's complementary, what's substitutable to AI. And we have to be prepared by trying to shepherd AI energy into the right domain. So those are really critical decisions.

MR. KORINEK: Thank you for your perspectives on that.

Let's stay in the realm of work but focus on working conditions as opposed to incomes. So this includes, for example, questions like worker monitoring. What do you view as the main potential benefits of AI in this space and what are the main potential harms?

MR. ACEMOGLU: That's a great question.

I think first of all it's a great question because it's a perspective that's sometimes often missing from economics and also from public discussion. You know, a job is not just an income stream, it's a meaning social activity. We define ourselves by the thing that we do because we spend the majority of our waking time doing those activities. Our social networks, our role in society, all of those are defined by our jobs. A job is a very, very important social construct. And in that context you have to recognize what are the things that people value about their jobs. Well, the social networks, the self-worth, the autonomy, the sort of sense that we are doing something useful. I think all of those are very, very important. And I think sociologists and anthropologists, for example, have recognized how disruptive social and economic changes are when they rob people of these meaningful, symbolic aspects. So there's a huge literature, for example, that shows how Native Americans got completely pushed into almost group level depression in many cases because they're most symbolic and meaningful social activities, hunts, the way of celebrating things, et cetera, were robbed from them by European settlers and Americans. So you have to think about what are those aspects of jobs that we have to think about. And autonomy I think is a very important part of it and being worth something because we are contributing something unique is part of it.

So in that respect, well designed AI and automation can help. If you ask workers who are in well managed automating manufacturing firms, who have survived and who have not left their jobs, they actually love the robots because the robots have increased productivity and they have taken some of the hardest and most dangerous jobs away. And if the company is well run, and not many of them are, but well run they also manage to save the workers from the robots who can actually chop your head off if you get too close to them. So managed well, some of the most routine, boring, uninteresting jobs can be

taken away by algorithms and machines and that would be good.

But, in fact, if you look at how AI is being used I addition to automation, and a very common use of AI is for monitoring work. And this monitoring often takes the form of not helping the workers, but completely robbing them of autonomy and making them work excessive hours. So Amazon on paper is a great employer, right. It pays a much higher wage than Dunkin Donuts or Wendy's or Walmart. But many employees don't like Amazon. Why? Because Amazon is at the forefront of using data and AI based processes to monitor workers, collect huge amounts of information, and actually make them work very hard. So it gets very efficient, but not by making them more productive, but making their work harder.

So there is a couple of publications that collected data — I don't think they are completely representative — but collected data on injury rates at Amazon and they are very, very high by industry standards. And they spike up during periods of peak work, like holiday season.

So that is just the tip of the iceberg about how little autonomy is left because of this monitoring. And I think if this monitoring continues unabated and unmonitored itself — forgive the pun — it's actually going to change the nature and social meaning of work in very important ways. And it's also going to have major effects on income distribution because monitored workers cannot ask for higher wages.

MR. KORINEK: So based on these observations what would be the one — or maybe two main policy measures that you would recommend?

MR. ACEMOGLU: Well, I think — again, this is such a new area that we have to be a little bit cautious in policy. You know, how do you distinguish between the good use of data and the bad use of data? But if I tell you I'm your employer and I'm collecting data so that I help you to become more productive, but actually I use that data in order to monitor you. So it's a very, very gray area. But there are at least some minimal things that we just absolutely must abide by, which currently Amazon or other employers are not.

We should never collect information and use information about individuals without their consent. Right now, nobody knows in the Amazon workforce what Amazon collects information about and how they're using that for promotion, for firing, for setting the work requirements for you. We should

also have some basic standards, like safety and autonomy that individuals should not be turned into more flexible robots. Again, where are the boundaries of that? Of course, if I am your employer and you're my driver, I can ask you to drive from point A to point B without taking a two hour lunch and then visiting point C on the way. But is it reasonable for me to say you should drive the fastest possible way, the fastest possible rate as if there is no traffic when there is actually traffic so that you actually have to go over 80 miles an hour in order to make that time requirement? So those are all gray areas. I think they have to be part of the public debate and public policy, but I don't have custom made sort of policy recommendations about where to draw the line on some of these issues.

MR. KORINEK: Thank you.

Now, you have also spoken about the ability of AI systems to collect and process vast amounts of data about its users, its consumers. What are the potential benefits and down sides to all this data collection and data processing?

MR. ACEMOGLU: Well, I think this actually gets to the heart in my opinion of what I call disrupt existing balances. You know, my view has always been that human societies, when they function even to a moderately acceptable level, they are based on a number of balances. The employer's power has to be balanced against the power of employees, the state's power against that of society, different political parties, within family, between groups in social interactions. And I think the biggest problem of the data economy that is disrupting these balances, it is empowering companies against their customers and it's empowering states against their citizens. And those two are lightning rods for me. They are very dangerous developments.

We are there perhaps, or will be there soon, but many companies will know more about you than you know yourself. That fear of private domain, where I make my decisions, I have autonomy, I have something approaching free will, that will all become meaningless. But worse, the same thing for the state. The state will know so much about its citizens, how can you do dissent, how can you do protests, how can you generate new ideas when that's the case? When your boss knows so much about you, can you be innovative?

So I think there are a lot of intersecting questions there.

Now, those are the bigger questions and I think we are going in a very unhealthy

direction, precisely because AI is doing the bidding, our current AI research is doing the bidding of a few very large companies who are amazingly expansionist in their vision. So these are not companies like some of the old big companies like GM that want to sell you a lot of products. They are much more like standard oil that want to control everything.

You know, Standard Oil was such a threat to American livelihood and democracy, not because it was making a lot of money — it was — but it also wanted to control the entire economy. It wanted to control the whole transport system, even when it was run by other companies. It wanted to completely be in charge of the whole industrial structure because it was controlling the life blood of industry. And it was actually exploiting that. So I think current big tech mindset is worse than Standard Oil. So therefore it's really feeding into this using information for empowering companies against customs.

And then the economics of it — so that's like the broader social thing — economics so it doesn't work out either. So privacy. You know, how can I have privacy when you are collecting so much data from all the people that I know, not just from me. So privacy debate is like framed as oh, well, consumers can opt out of things and they can protect their data if they don't really want to. Well, it doesn't actually work that way. Data is the life blood of many AI applications. AI cannot be thought without abundant data. And some of the good uses of AI really depend on that data. So if you wanted to cure cancer, again, no change of that using AI in the next several decades in my opinion, but we can make progress. If you want to make progress in that you need to have a lot of data about what different types of systems and what different types of behaviors are associated with different onsets and developments of cancer. So that means you take a lot of data and pool them.

But the problem is once you have that capability for collecting a lot of data and pooling them, then I can collect a lot of data about you from your friends, from things that are correlated with you without your consent. And that really throws the whole system into disarray.

MR. KORINEK: Let me perhaps ask are there any specific regulatory solutions that you propose with respect to data, perhaps in the realm of privacy?

MR. ACEMOGLU: Yeah, I think we have to strengthen the privacy requirements. But we also have to sort of revive consumer protection. What do we mean by consumer protection in the age of

data and AI? You know, Ralph Nader, for example, before he became a third-party candidate disrupting the political system in the U.S., was a hugely important political figure because he had a defining role in how we perceive of consumer protection, what happens when you get faulty products, what happens when safety is not prioritized in cars. So that was not a very well defined problem before Ralph Nader started doing the public advocacy and revealing data about what was going on.

So I think we need to do the same here. We need to understand how it is that these companies are using the data and which are the ones that are acceptable and which are the ones that are not acceptable. I personally don't find it problematic that Amazon gives me recommendations. But there might be other things that I find quite problematic about what Google knows about.

So I think we have to decide which ones individuals generally find problematic but don't have the agency to deal with it. And that's where public policy and the regulatory system has to come in.

On the other side, and in contradistinction to it, we also have to find ways of making data more available to other competing companies. So you see that these two objectives are contradictory. One says protect the consumer, especially their privacy and what you can do with data, but the other one says now, if Anton collects data about Daron, then Anton competitors should also have that as well. So these two are both very important because if we don't do one of them, the consumer is ill-treated, but if we don't do the second, this time Anton gets to have artificial monopoly because he's the one who got there first and started collecting too much data about me. So we have to solve both of these problems and you cannot solve without the other. Why? Because if you don't have a good privacy infrastructure you cannot make that data available. So you have to solve both of them at the same time.

MR. KORINEK: Thank you, Daron.

Now, in your paper the final category of potential harms from AI were harms to the democracy. And I should emphasize to our audience that Daron has not only contributed groundbreaking works to economics, but he is also the author of two books at the intersection of economics and political science, together with James Robinson — "The Narrow Corridor" from 2019 is a book in which they argue that liberty requires a very delicate and precarious balance struck between this data and society.

Now, in your research, Daron, you're making the case that AI as practiced by many of today's leading social media platforms is contributing to echo chambers and other adverse effects that

undermine our democracy. I cannot help thinking of the Facebook files, for example, in this context. And additionally it also exacerbates the big brother effects that allow autocratic governments to suppress dissent from their citizens.

Can you explain in a little bit more depth what this is about? And in the spirit of balance, what could AI system actually do to enhance our democratic societies?

MR. ACEMOGLU: Well, first of all, thank you again.

And on this one I think I am in good company. There are a lot of other people who have over the last three or four years woken up to this danger and have written — some at high level, some with more data — and I'll mention a few of them. But just let me start from the last part of your question. It's actually ironic. The first wave of thought on the effects of digital technologies, internet, on democracy was completely the opposite. In the early 2000s everybody and their dog thought that we were entering a new era of openness, democratic journalism, democratic dissent being much easier because everybody could write in Wiki and we could use the internet to monitor corrupt politicians. And of course I think today nobody clings to that view anymore. But it's not because that view as fundamentally wrong, it is because that view would have required a very different direction of technological change and a very different balance of power.

Once the internet was recognized, once data was recognized as an important resource by authoritarian governments and companies, they had much more resources to control the technology. So a good example was the Iranian Green Revolution. So when students, youth, and some workers rose up against Ahmadinejad's incompetent dictatorial regime, at first they were using modern social media. Even though Iranian state is not at the forefront of technology, they had enough resources that once they pooled their resources they could turn the tables on them. Now they use internet and all of the existing social media to catch all of the protestors and throw them in jail. And it's a cautionary story, which is again much better told perhaps through the wide lens of China of what I call big brother effects, that now what George Orwell feared in 1984 seems like child's play.

So the amount of information that you can collect about somebody who uses social media, if you're a government or intent on doing so, is just enormous.

So that really strengthens all sorts of governments, authoritarian and non-authoritarian,

against their citizens. It's quite clear what the implications are if you're China or Iran, but I don't think the implications are going to be any more innocent if you are USA and NSA. So that sort of balance is completely out of the window when the government can collect so much information.

And there are a lot of different problems that this creates. Dissent becomes much harder, opting out becomes much harder, because now opting out becomes a signal. So if everybody is using a browser that government can track very well and they see that you cannot be tracked, that must be using Tor. If you're using Tor you must be a dissenter. So there are a lot of second round effects that all help the powerful governments.

Now, that may or may not be the more important problem, but the one that actually has afflicted the U.S. more recently is Facebook. So Facebook I think is not unique to Facebook but is emblematic of what the particular business model of many of these tech companies is doing. First of all, Facebook, like Google and many others, has a business model that's ad financed or ad monetized. So that means that they have to maximize engagement. And in order to maximize engagement they have to feed you things that are going to get you hooked, that often tend to be more sensation stories that are going to create either echo chambers or artificial sort of reactions. And they can do that with micro-targeting. So then you see the two parts of the business model. One is maximize your hooked nature to the platform, how much time you're going to spend impulsively on the platform. The second is complete disregard for your privacy so that I can collect as much information about you and I can micro-target things that are going to make you hooked or pull you into an echo chamber and get the biggest reaction from you. And that's very specific to Anton. If it was for Daron, it's something, for Megan, it's something else.

So that we have seen if it's unregulated can be quite disastrous. There was a huge amount of extremism fed through all of these social platforms. Some of them have cleaned up their act a little bit. Like Reddit was one of the worst and they started cleaning up. YouTube actually was horrible. YouTube's feed early on, once they started introducing AI based recommendation system, was keeping a lot of young people hooked and then taking them more and more to extremist websites. So you would listen to somebody complaining about woke culture — they didn't call it woke culture eight years ago, but something like that — so then they would feed you a video by Jordan Peterson and then you listen to

Jordan Peterson and then they would feed you an alt right one and soon you were listening to Neo Nazis.

And see it's a public policy problem. Lots of young people being indoctrinated into alt right or Neo Nazis is not an individual choice that we can say oh, that's a consumer choice between the company and the individual.

So those are the issues that we really have to sort of worry about and again requires a much deeper diver into public policy. What do we expect from social media? Do we set a much lower standard for them than traditional media, which is what we're doing right now? In fact, we subsidize what they do. Or do we actually understand all of their broader social effects.

And look, again there's nothing about the nature of AI per se. The early ideas that AI could empower the centers, journalists, et cetera, that could have been true, but it required a different path of technology. So what can we do with AI, we can go back to those aspirations, we can try to develop AI much more like Tor or much more like differential privacy or encryption that protects individual autonomy, protects individual privacy, protects dissent. We're not doing that. We're not going to do that when Google, Facebook, or Amazon makes the choices. So that's where we need diversity of approaches.

MR. KORINEK: Thank you, Daron.

So we have been through a whole range of areas related to AI and unfortunately we are in last six minutes. So let me ask you what broader lessons should we draw for the need for AI regulation? And perhaps also to what extent would competition help or hurt us with the problems that you are describing?

MR. ACEMOGLU: Well, let me actually start from the second part of your question, because I think if you are an economist, if you look at the current landscape, it looks like a monopoly, smells like a monopoly, quacks like a monopoly. You have these huge companies and they are dominant in their sectors. And now we haven't taken action against them because of a stupid criterion that was propagated by people like Robert Bork and others, which is that you have to prove almost beyond reasonable doubt that a monopoly creates harm to consumers. So that's why we have approved a lot of mergers, or U.S. courts became very company friendly.

So therefore if you're an economist, your first instinct might be — and I applaud that —

well, let's get rid of this Borkian tradition, let's tackle these monopolies with anti-trust. And there are some economists who are finally waking up to that and there are some legal scholars who are finally waking up to that and now Lina Khan is in charge of competition policy and she is a big advocate of using antitrust. I applaud that, but I think it's completely inadequate because we're not really dealing with a monopoly problem. Monopoly is part of the problem, but we're not dealing with a monopoly problem.

You know, the antitrust tradition in the U.S. goes back a long time, but one defining person was Brandeis and he was writing at the age of trusts where really the problem was a monopoly problem. Standard Oil was using its control of the U.S. economy in order to charge higher prices. And, moreover, with the regulatory state that existed at the time, there was nothing else other than antitrust, using the Clayton Act and the Sherman Act, that they could have used.

But today, we are not dealing with a monopoly problem. We are dealing with something much more pernicious. I've pointed out what this is doing to our democracy, to the future of work, to the citizens' information, to privacy, to misinformation. None of those you're going to fix if you get these companies to charge lower prices. In fact, they're not even charging prices anyway, they're just making money from ads. So you need a much more expansive way of dealing with it. So I applaud the anti-big tech sentiment. I think that's much needed. But we have to channel that the right way. And I think the right way to channel it is to really think of a broader regulatory framework.

So let me try to amplify that point for 30 more seconds. Say, for example, if you have the narrower vision of what to do with regulation triumph and we break up Facebook into WhatsApp, Instagram, and Facebook proper. Perhaps it helps, but most likely it won't because all three of these companies will still have huge effects and they will have exactly the same business model. And some of them may become even more extreme, less self-restraint in order to grab market share from others.

So what we need is a complete overhaul of their business model and a complete model in the direction of technological progress. None of this you're going to automatically achieve by competition policy. So that's why we need a broader debated. And as part of that broader debate we need to really develop much better ways of regulating what these companies do, what their power is, wire information, wire control, and we have to have a much more robust approach to redirecting technological change.

MR. KORINEK: Thank you, Daron.

Now let me ask one final question. In light of the problems that we have just been discussing, one temptation is to feel overwhelmed by the magnitude of all our challenges and just throw in the towel, but that seems hardly like the best course of action. So we'd like to ask you if you have one actionable piece of advice for the members of our audience, many of which are active in policy circles. What would it be?

MR. ACEMOGLU: Well, actually I'm a huge believer in the importance of discourse. When the tone of the discourse changes everything changes. Let me give you a small example. Three years ago it seemed like it would be crazy to expect any new public policies from the government in terms of fighting poverty and supporting children and financial difficulty. Today, U.S. has actually lower child poverty rates and it is possible that we'll pass legislation that will actually help poor families with children tremendously by a variety of levers.

How did this become possible? Because suddenly the discourse changed. So the future of technology is all about discourse. So the first thing we have to do is change the discourse. The current discourse is AI is wonderful, AI is going to give you everything you want, you have to let AI decisions in the hands of the people who know best and everybody else, all these activists and all these citizens, you don't understand IA, you don't understand its promises, just stay away from it. That's the discourse we have accepted, that's the discourse the New York Times has fed you, that's the discourse that the Democratic Party as well as the GOP are completely on board with. And I think that discourse has to change.

Sure, you and I and the activists are not going to know how to best design an algorithm, but we sure need to be consulted if that algorithm is going to catapult our lives and is going to collect huge amounts of data about us and is going to impoverish us. So the consequences of these new technologies are things that we all have to have a say on, and that's the change in discourse that we first must go through first. And that's where I think we can all make a difference.

MR. KORINEK: Thank you, Daron. Thank you for participating in today's event and thank you for doing research on all these difficult questions that we have discussed today.

Meeting the challenges that we are facing in the age of AI will require hard work from

everyone, from bright minds in academia in our policy community, and most importantly that includes you listening to our conversation.

I hope that you have found both inspiration and actionable insights in today's conversation and that you will join us again for future events in this series.

And for today, let's end with a round of virtual applause for Daron.

Thank you.

MR. ACEMOGLU: Thank you, Anton. It was great to be part of this conversation.

* * * * *

CERTIFICATE OF NOTARY PUBLIC

I, Carleton J. Anderson, III do hereby certify that the forgoing electronic file when originally transmitted was reduced to text at my direction; that said transcript is a true record of the proceedings therein referenced; that I am neither counsel for, related to, nor employed by any of the parties to the action in which these proceedings were taken; and, furthermore, that I am neither a relative or employee of any attorney or counsel employed by the parties hereto, nor financially or otherwise interested in the outcome of this action.

Carleton J. Anderson, III

(Signature and Seal on File)

Notary Public in and for the Commonwealth of Virginia

Commission No. 351998

Expires: November 30, 2024