

Forward Thinking on technology and political economy with Daron Acemoglu

The influential economist connects the dots between artificial intelligence, productivity, wages, and inequality, and how to counterbalance the impacts of automation.



In this episode of the McKinsey Global Institute's Forward Thinking podcast, hosts Michael Chui and Anna Bernasek speak with Daron Acemoglu. Daron is a professor of economics at the Massachusetts Institute of Technology (MIT), a coauthor (with James A. Robinson) of [Why Nations Fail: The Origins of Power, Prosperity, and Poverty](#), and the author of many influential academic papers. His research covers a range of topics, including political economy, economic development, economic growth, technological change, inequality, labor economics, and economics of networks.

The conversation focuses on how we make growth happen in a world of technological change that is largely disrupting labor markets. As Daron puts it, "I believe that the political economy of growth ... and technology, automation, AI—those [things] are intimately connected."

An edited transcript of this episode follows. Subscribe to the series on [Apple Podcasts](#), [Google Podcasts](#), [Spotify](#), [Stitcher](#), or wherever you get your podcasts.

Michael Chui (co-host): Hi, and welcome to *Forward Thinking*. I'm Michael Chui.

Anna Bernasek (co-host): And I'm Anna Bernasek.

Michael Chui: We've got a fun episode today because we're talking to one of America's favorite economists, at least according to other economists.

Anna Bernasek: Hold on, hold on. That sounds really geeky to me, Michael. I hope you're going to translate for the rest of us.

Michael Chui: There won't be any translation required. Our listeners will be able to get this immediately. Daron Acemoglu is a professor of economics at MIT, and he's one of the most widely cited economists working today.

Anna Bernasek: OK, that's right. I know his expertise is in growth, technological change, and the consequences of that on things like inequality. Plus, he's really into understanding [the impact of AI](#) and automation, which I know is why you're excited, Michael, because that's one of your areas, right?

Michael Chui: We have done a lot of work in this area. In fact, MGI has published a report on [the future of work after COVID-19](#), and we found that over 100 million more people could have to change jobs over the coming years. But Daron's work on the impacts of labor is only one of the areas for which he is known. He has also done some very influential writing on political economy.

Anna Bernasek: Those two topics are really of the moment. Let's go to the interview and hear more.

Michael Chui: Daron, thanks for joining us. I appreciate it. Why don't we get started? I'd love to start with just a little bit about your own personal journey. Where did you grow up, how did you go from where you were to now becoming an Institute Professor at the Massachusetts Institute of Technology?

Daron Acemoglu (guest): I was born in Istanbul, Turkey, and I grew up in Istanbul until the age of 19, when I moved to the UK to study. And my interest in economics and politics took shape when I was in Turkey. And in fact, many of the questions that I later came to study, such as democracy, the relationship between political institutions, and economic growth, and dictatorship repression, those were things that I started wondering [about] when I was a teenager, especially growing up in a country that had recently suffered a military coup. [Its] economy [was] in the doldrums, inequality was rising, so those were the questions that I thought I would go and study in economics.

Little did I know that that's not what economists study. But when I went to enroll at the University of York, I still liked what I saw, so I got into it. But then a number of years later, when I was towards the end of my PhD dissertation, the same passions came back. And now I felt I could take a little bit of time off from more conventional topics: unemployment, human capital, and so on.

But over time I think my two interests have merged. Moving forward, I believe that the political economy of growth, how we make it happen, what sort of institutions we have to have in order to undergird growth, who benefits from growth, how you regulate growth, and technology, automation, AI, the direction of technological change—those are intimately connected. Those are the two legs of my research that have, in recent years, come even more together than they were before.

Michael Chui: And certainly things have worked out well. Amongst other things, you have a John Bates Clark Medal, which has been a reasonable predictor of a Nobel Prize in some cases. Do you ever think about that? You've published hundreds of papers. If you go to Google Scholar, I think there are like 37,000 different results that come up. And then you were also ranked third behind Paul Krugman and Greg Mankiw in the list of favorite living economists under age 60 by other economists. Which of these was your favorite award?

Daron Acemoglu: My favorite award is probably the BBVA award I received a couple of years ago because I think it's a very serious award, and they have an amazing list of recipients in the past, and also [a] very, very good selection committee, and I'm not saying that just because they gave it to me.

But they are objectively some of the best economists, and they do it very seriously, and so I was very gratified and honored. But all of them are like my babies—I love them all.

Michael Chui: Why don't we start on one of the legs of your research stool that you mentioned, and that was my entree into your work. [At] the McKinsey Global Institute, we also have been studying AI, automation, the future of work, [productivity](#), and those sorts of things, and you and I have both been involved in the Partnership on AI for people and society. You do have an extensive body of work, of course, that many have reviewed, along with Pascal Restrepo, David Autor, and others. How did you get into it? Were you thinking about robots when you were a kid, for instance?

Daron Acemoglu: No, I wasn't. On that one I wasn't. Essentially, the way I got into it was that very early on in my career, about 25 years ago, I started working on what I sometimes call [the] directed technological change issue, which is that technology is not a sort of scientific process that has a preordained path, but its exact direction, which type of technologies we go after, who benefits from it, which gambles we take, which products we aspire to, all of those are choices.

The direction of technological change is something we determine socially and economically. And I worked quite a bit on that, both theoretically and in the [context of inequality](#), for example. What makes technology more skill biased? What makes technology more favorable to capital versus labor? And I think those models and insights I'm still very keen on.

But at some point I also started noticing the direction of technological change within the context of the most common approaches in economics, which is that technology just increases the productivity of some factors, and therefore its inequality consequences will come depending on

whether it increased the productivity of capital or labor, or skilled labor, or unskilled labor, agriculture versus industry.

You know, those are useful, but they don't do justice to the nature of technology and to its more sweeping distributional effects. And in particular, what was missing from this framework is exactly what you would think of as automation, which is the possibility that new technologies work not by increasing the productivity of certain factors of production, but they replace them in the production process.

Think of, for instance, the spinning or the weaving machines of the early Industrial Revolution in Britain replacing skilled artisans. The implications of this both for inequality and for how we should think of the determinants of technological progress turn out to be quite fundamental.

And then once I started thinking about these issues, then robots, AI, et cetera, immediately fell into place because of course it's really interesting to wax lyrically about the British Industrial Revolution, but that was in the past, whereas robots, numerically controlled machinery, AI, algorithms—they are our present and the future.

That's how my entry into the area of robots at the end of the 2000s, early 2010s, took place, and then a large part of the work I have done in this area can be thought [of] as both a corrective and as an effort to develop a new framework. The corrective part is, because as I have just tried to articulate, I think there is something unsatisfactory, and perhaps even misleading, about the ways that we have thought of technology as economists, and that was the corrective part.

It takes a framework. We need to come up with something that's workable, that has predictions, that shows you how things work, and has testable implications so that you can actually look at data, confront the data. And those are the efforts that I have been involved in over the last 12 years or so.

Michael Chui: Can you state simply what was wrong about the way that economists thought about technology?

Daron Acemoglu: There are three interrelated things about it. One is actually its descriptive realism. I think it matters quite a bit that what we describe has a clear counterpart in the data. So many economists still subscribe to the mythological point of view that [Milton] Friedman articulated almost 70 years ago now, where he said, "Realism of assumptions [doesn't] matter."

You have to just look at predictions. His favorite example was, I can model a billiard player by thinking he's calculating all the angles. That's not realistic. That's not what he's doing. But if I get good predictions about where he's going to hit the white ball, that's fine. I think there's something attractive about that methodology. But I think it's also very limiting.

And a good, rich model has to have close connection with the data. And if you look at the way that economists think about technology, it's this latent variable that makes you just more productive. But very few technologies actually do that. Electricity didn't make workers more productive. It made some functions in factories more feasible, and some few items more productive. A hammer doesn't make you more productive in everything. It makes you just more productive in one single, simple task—hammering a nail.

And many technologies don't even do that. The example of spinning and weaving machinery that I gave, or the factory system, or, more recently, databases, software, robots, numerically controlled machines, they are mostly about replacing workers in certain tasks that they used to perform.

But the second is that the problem of what technology does to wages is very much entangled with this notion.

If technology is just this latent variable that makes you more productive, it will help you. At the end of the day, you're becoming more productive. Of course we can come up with scenarios in which things are so topsy-turvy that you may become more productive, but at the end it doesn't help you all that much. But essentially it's a very robust prediction, that the classic economist way of thinking about it, where technology is an augmenting factor, tends to benefit workers.

It may benefit some workers more than others. It may well be that computers augment educated workers more than high school dropouts, so inequality can increase. But at the end you shouldn't see the high school dropouts lose out. Their real wages shouldn't decline. And the real wages of workers shouldn't decline.

But, in fact, one of the striking but very robust features of the last 40 years of economic development in the US and the UK has been that many groups, especially low-education or middle-education men, have actually seen their earnings fall, some groups by as much as 25 percent, in real terms, since 1980. Phenomenal. This isn't the American dream.

In the traditional economics approach, this is a nuisance that we often sweep under the carpet. We look at relative wages so that the decline in the absolute wage level of some groups is sort of hidden. But it is something that doesn't really fit into this technology as augmenting framework. But when technology, at least in part, is about automation, replacing, displacing workers from their tasks, then this happens quite often.

You can have productivity improvements—capital benefits, firms benefit, but workers, especially some types of workers, all workers overall can lose out in real terms.

And then third, once you go to this micro level, then the direction of technology, the future of technology looked at through the perspective of what type of technologies we're going to build on, that becomes much richer and much more interesting.

It's not just whether we're going to increase the productivity of skilled workers versus unskilled workers, both of which benefit all of them since they are complementary. It's more like, are we going to completely give up on unskilled workers? Are we going to try to replace them? Are we going to try to replace humans? Are we going to create new tasks for humans? How are you going to use the AI platform? All of these questions about the direction of technology become much more alive, and then also the productivity implications become much more interesting.

Michael Chui: Let's take the second of those. I'm very interested in this idea or this question. People have criticized technological determinism, that we know what will happen. But let me probe with you. The link between technology and inequality or depressed worker wages of certain workers, is that a causal relationship that you could say you could predict?

Daron Acemoglu: Yes, it is a causal relationship in the sense that I can say with some degree of certainty about what caused the huge increase in inequality in the US labor market over the last four decades, and technology is at the forefront of that. But, no, it's not a preordained, deterministic relationship whatsoever. Technology's played a very different role. Look at the macro facts.

It actually screams at you. It's pretty clear. Look at economic growth in the 1950s, '60s, early '70s. It was very rapid. But it was also very shared. If you look at the real wages of different demographic groups—men, women, high-education men, low-education men, college graduates, postgraduates—and plot them on a graph, they're all on top of each other.

Everybody is experiencing this real wage growth between 2 and 2.5 percent in real terms per annum. It's remarkable. And then you come to the 1980s and a completely different picture. First of all, median wages are essentially flat. Real wage growth in the US, especially for men, but overall, has stopped. Second, the top and the bottom are coming apart, [a] huge amount of inequality is opening, definitely not shared.

But third, it's not that the top is going ahead and the bottom is not keeping up with it—the bottom is actually falling. So those patterns, what caused them? Of course many things might have caused them. It may well be that low-education workers became less skilled. It may well be that it's globalization. It may well be that it's because the minimum wage has fallen in real terms by, oh, about 35 percent since the 1970s, or collective bargaining has come to an end.

Several of these things do play some role. But my own research claims, and hopefully shows, that technology is at the forefront. The reason why, in

the '50s and the '60s and early '70s, we had this shared prosperity is because we did have quite a bit of automation. It's not that there wasn't automation, but automation was counterbalanced by other technological changes, especially these new tasks.

And for every sector that had a significant [decline in labor share](#) because they were automating a lot, there were some new sectors that were coming up and demanding more labor, paying more to labor, and it was this counterbalancing nature of different types of technologies that create that balanced progress. And then you look at the 1980s, completely different picture. What you see is the automation picks up speed, but even more remarkable is all those countervailing types of technologies are absent.

Michael Chui: With that being the case, what is the hope for the future here? Do we need to unlock the creation of new industries?

Daron Acemoglu: I think we need to do a little bit more diagnosis before we can answer that question. The diagnosis is about exactly what is happening, of which we talked about one part, and I want to talk about one other part of it, and then second, why it's happening, and that will give us some clues.

We talked about automation. There is a lot of automation. But one thing we did not mention is that even if we automated a lot and did nothing else, there's one condition under which that would not be disastrous for workers. For example, it wouldn't translate into real wage declines. And it is when automation is actually highly productive.

Automation needs to be counterbalanced by these new tasks. That's the simplest way of thinking about the problem. But automation by itself creates two forces as well. One is the displacement. I'm getting rid of my blue-collar workers and many of the tasks, but if at the same time I increase my productivity so much, or reduce my cost so much, that I expand my production, I hire back those workers to do the other remaining tasks.

Essentially what we are seeing in the US, and some other economies as well, it's not only that automation is not being counterbalanced, but it's actually not sufficiently productive. And it's not generating the additional labor demand that would

come from the productivity effect. And it's not a surprise to people who have looked at national income accounts, because the last two and a half decades, almost three decades, have seen [some of the slowest productivity growth](#) in US history.

Michael Chui: Unquestionably the slowdown in productivity growth has been an issue, but others have also noted a decoupling between productivity growth and wage growth as well. Are you explaining that?

Daron Acemoglu: Yes, yes... Essentially there are two separate issues here. One is that productivity growth is slowing down and that, everything else equal, that's going to be really bad for labor. And second, automation by displacement effect always creates a wedge between productivity and wages.

Or a different way of thinking about that is that automation—again, this is my framework that I'm peddling, so it's not true in the standard framework; it works somewhat differently. But automation always reduces the labor share, always creates this wedge between productivity and wage growth. If productivity is very high, you can have that wedge, but wages might still increase. But then if you have productivity growth and a lot of automation, or productivity growth slowing down and a lot of automation, it's a double whammy. And that's what's really problematic.

Michael Chui: You said we need to diagnose before the solution, but what is the solution?

Daron Acemoglu: This was part of a diagnosis because now, let me come back to the two broad theories you can have, and this last fact will help us perhaps at least change our postures a little bit on those. One is we are going through a "singularity abundance" type of moment. We have worked out a way of producing more machines with machines, ultimately culminating to something like singularity or superintelligence, but it doesn't even need to go there.

Even if it just accelerates but it doesn't take us through to superintelligence, we can have that. We're going to have a lot of automation because we've worked out how to do machines with machines, and more and more efficient machines of that sort. And that's what's leading to this imbalance

between more pro-labor technologies and more automation. But if that was the broad explanation, you should also see a lot of productivity growth because it's working through new opportunities for doing much better with machines.

Michael Chui: What's an example of a pro-labor technology?

Daron Acemoglu: The new tasks that I was talking about. In the context of, for instance, mechanization of agriculture, where did all these workers who were displaced from the agricultural sector, which made up about 50 percent of employment in the US towards the second half of the 19th century—where did they go? Well, they didn't go back to agriculture to do other tasks.

And they didn't just go and do exactly the same things that their parents did. Instead, during the same time, you see this complete transformation of both the service and manufacturing sectors in the US. The manufacturing sector is electrifying, is adding all these nonproduction workers, clerical workers, back-office workers, engineers, maintenance workers, doing much better design, much better cost control.

And these are the new tasks that are soaking up all of that labor. And the retail sector, education sector, they're also expanding with completely new activities. Those are part of the more pro-labor things.

Michael Chui: And just let me push on this a little bit, because we've often heard this trope, and it comes in different forms, but in 20 years, 60 percent of the jobs that will exist will be ones you've never heard of. And we've tried to uncover that, and it seems that that rate of change in occupations, and even tasks and activities, is quite a bit slower than that. What is your research saying?

Daron Acemoglu: That's exactly what I'm talking about. We haven't done enough of that. We don't have the same detailed data, but at a more ocular, statistics level, at high level in the 1920s and '30s, it's much faster. We do have many new tasks. It's not like we don't have any, but they are not enough, exactly. That is the imbalance that I'm talking about.

Michael Chui: Are we just not inventive enough to come up with new tasks? Or do we need [universal basic income](#), because we've come up to such a pace?

Daron Acemoglu: That brings me to my second explanation. It's not that we have this imbalance because we have suddenly found much better ways of doing automation, doing machines, doing algorithms. Perhaps it is that we are inefficiently biasing our direction of technology. You can think that perhaps the possibilities offered to us haven't changed, but because of who's making the decisions, incentives, institutions, aspirations, we have gone more and more in an inefficient direction of doing too much automation and not enough of the other things.

Let me give you one example. I don't think it's the most important example, but let me give you one example to explain what I mean. Our tax system, if you look at today, if a company buys a machine to displace workers, it will pay about a 5 percent tax on the spending on that machine. If instead you hire the workers, you'll pay something like 25 percent to 30 percent tax.

The US tax system [provides an enormous subsidy to firms to buy the machines and replace the workers](#). Did this subsidy exist before? Yes. But much, much smaller. In the 1990s, the difference was 25 percent to 15 percent to 20 percent. So we are increasingly subsidizing firms to automate. If you have many incentives like this—and we can talk about what the other ones are, which are softer perhaps, but are no less important in my opinion—then you're going to go more in the direction of automation. You're going to leave lots of low-hanging fruit. And those low-hanging fruit might have been very important for that productivity improvement.

Michael Chui: You made reference to institutions and how they operate. This is probably going to segue into another area. You and your coauthor and coresearcher, James Robinson, wrote this book in 2012, *Why Nations Fail*. You talked about political economy as being something that is of deep interest to you. We could talk for hours on this, and we don't have that long, but a key distinction you make is between extractive and inclusive institutions. Can

you just expound a little bit for folks? People should read the book, but if they haven't read the book yet, can you give them a teaser or preview?

Daron Acemoglu: Let me give you one concrete example rather than give you a definition. The institutions are essentially the rules, political and economic rules—who owns property, how you contract on it, who can enter into what occupation. But let me give you one simple example. The South African economy in the middle of the 20th century under apartheid [was] an economy in which almost all of the firms and assets were controlled by a white minority that makes up about 5 percent to 6 percent of the population.

The 90 percent Black population was forced to live in the most unproductive part of the country in very, very crowded and bad circumstances. All of the good land was taken by the whites. And then they had no public goods, no public services, no education, so they had no chance but to supply their labor so cheap to mines or agricultural estates.

But even that was not enough. Actually the South African government passed something called the color bar. It says that Blacks cannot hold any other occupation than the most menial, the most unskilled. You cannot be a supervisor, a surveyor, a clerk, an accountant, a bricklayer, anything that is beyond the most menial, lowest-pay, lowest-skill occupation; Blacks were completely banned from it.

This is what we call an extractive political institution. Extractive economic institutions concentrate economic capabilities in the hands of a small elite. Extractive political institutions concentrate political power in the hands of an elite. And what we argue in the book is that extractive economic institutions bolstered by extractive political institutions are very common. They often lead to high inequality. They may sometimes allow some growth. But it's not sustainable growth, it's not innovation-based growth, and it's not ultimately growth that's going to bring productivity growth and broad-based prosperity. The alternative [is] inclusive institutions that are the opposite, that create these broad-based opportunities, anybody can enter into business or whatever profession they want, but also the means of achieving that: infrastructure, health, education,

status. I think those are going to be much more conducive to long-run productivity growth.

Because the experimentation, the innovation, the creative destruction that's necessary for true technological progress is going to be very difficult under an extractive system. Catch-up growth is fine, or if you have a clear competitive advantage such as oil or copper, sure, you can do that, but innovation-based growth, new products, new technologies, new ideas, that's going to be harder.

And why nations fail—even if troubled, the US is an example of an inclusive nation, it has active political participation by a large fraction of the population, it has an open system. You cannot ban anybody from becoming an entrepreneur or an engineer. You could in the US South at some point, but not anymore. But here's the problem, if my diagnosis in the first part of our conversation is correct, that we are creating more and more of this automated future, and AI will continue to trend, and especially for workers without specialized skills: programmers, the top lawyers, the top surgeons, OK, perhaps they can escape automation, but then we're going to create a world in which 60 percent, 70 percent, 80 percent of the population doesn't have much economic value.

What they do either can be done by machines or their price is depressed. They don't have much social status. They will increasingly get alienated from public life and political life. Therefore, in my assessment, fixing the direction of technology, fixing the direction of AI is important not just for prosperity, but it may actually be important for our democracy as well.

Michael Chui: What would it mean to fix the direction of AI?

Daron Acemoglu: To me, it has two legs. One is we have to free ourselves from the excessive obsession on automation. It is true in the area of AI. It's true in other areas, too. [In] our current business community, for a variety of reasons, some of it is cost cutting, some of it is because where the technology leaders in Silicon Valley have sort of set the agenda, some of it is because government policies are just too focused on automating everything.

Instead, we have to come back to a world in which we put as much effort in increasing human productivity, both in the tasks that they already produce, but also creating new tasks in entertainment, in healthcare. There are so many new things that we can do, especially with AI, but some of it with just our existing technologies, some of it with virtual reality or augmented reality. There are many, [many things ranging from judgment, social skills, flexibility, creativity, that humans are so much better at than machines.](#)

But we're not empowering them right now. That's the first leg. That second leg is that we also have to pull back from using AI as a method of control. And again, that's about how we use the AI technology.

Do we use it to empower individuals? To be better communicators, better masters of their own choices and data? Be able to sort of understand the veracity or the liability of different types of information? Or do we develop these tools in the hands of platforms so that the platforms themselves are doing all of that thinking and all of that direction for the individuals? I think that those two are very different futures as well.

Michael Chui: That's a great place to leave it. Daron Acemoglu, thank you for joining us.

Daron Acemoglu: Thank you, Michael. This was great talking to you.

Daron Acemoglu is an Institute Professor at MIT and an elected fellow of the National Academy of Sciences, the American Philosophical Society, the British Academy of Sciences, the Turkish Academy of Sciences, the American Academy of Arts and Sciences, the Econometric Society, the European Economic Association, and the Society of Labor Economists. He is also a member of the Group of Thirty. **Michael Chui** is a partner of the McKinsey Global Institute, where **Anna Bernasek** is a senior editor.

Designed by the McKinsey Global Institute.
Copyright © 2021 McKinsey & Company. All rights reserved.