

Daron Acemoglu

On Henry Ford, making AI worker-friendly, and how democracy improves economic growth

Daron Acemoglu is one of MIT's nine university-wide Institute Professors, the university's highest faculty rank. One of his predecessors, Robert Solow, developed a pathbreaking mathematical model of economic growth in the 1950s. Today, Acemoglu says hurray for economic growth — but is also concerned that choices made by policymakers and companies are channeling the gains from that growth away from workers. And as he sees things, the powerful AI technologies that have come to the fore in the past several years, embedded in products such as ChatGPT, should be regulated with the economic interests of workers in mind.

Acemoglu's research on the role of technology in economic growth, the decline in labor's share of income, and other topics has made him, according to *Research Papers in Economics*, the third most-cited economist in the world. He is also the author or co-author of six books, including the 2012 bestseller *Why Nations Fail*, in which he and James Robinson argued that differences in affluence between countries are mainly driven not by differences in natural resources or climate but by their economic and political institutions. His latest book, published in May, is *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*. Among his numerous professional awards is the American Economic Association's 2005 John Bates Clark Medal, recognizing the American economist under the age of 40 who is judged to have made the most significant contribution to economic thought and knowledge.

He is married to an MIT computer scientist, Asu Özdağlar, a frequent co-author of his who is an expert in optimization theory, game theory, and social network theory. He credits her with helping him understand the new generation of AI technology as well as giving him an insider's view of her field — as he puts it, “how the computer science discipline works inside and outside academia.”

David A. Price interviewed Acemoglu by phone in April.



EF: How did you become interested in economics?

Acemoglu: I became interested in economics when I was in high school and coming of age in Turkey, which was under a military dictatorship at the time. The country was having a lot of economic problems, including widespread poverty. So I started becoming drawn to these issues and also wondering about the linkages among dictatorships, democracy, economic growth, and the things that came to occupy my research decades later.

What I thought was economics at the time turned out to be not quite exactly economics, and what I thought I was going to do changed quite significantly when I went to college and graduate school. But I liked what I saw — trying to approach the questions of social science and, really, of humanity's existence using empirical, mathematical, and conceptual tools.

EF: In your new book, *Power and Progress*, you and Simon Johnson argue that the degree to which workers share in the gains from new technologies depends on the legal rules and informal expectations governing management. The rules governing management have obviously changed over time — have the expectations also changed? Does our society expect something different from CEOs than it used to?

Acemoglu: A large chunk of my work focuses on economic growth. My early work within that followed what many economists do, which is to look at growth — the growth of GDP — with the expectation that as the economy grows, that's going to be beneficial for all segments of society, including workers. *Power and Progress* is a culmination of

my research over the last 15 years, which has made me less certain about that. It has certainly worked out fine during certain episodes; for example, in the decade that followed World War II, U.S. GDP grew rapidly and so did wages. Inequality remained stable. So this was a classic period of shared prosperity. But during different episodes in history and even today, there's a variety of evidence suggesting that this doesn't always work so seamlessly.

Power and Progress tries to approach these questions, putting emphasis on three things. One is the nature of the technology of the time; the second is the institutions that shape the bargaining power of workers; and the third is expectations and norms.

And on the last one, yes, I think expectations and norms are particularly important, especially in the modern world. We are in the midst of a big transformation in which new technologies, new organizations, new ways of living are spreading around the world, and who is going to get the benefits of these is very much up for grabs. The expectations and views about what is acceptable are going to play a very important role.

One sneak peek at that comes from the research I've done recently with Alex He, who is at the University of Maryland School of Business, and Daniel le Maire, who is at the University of Copenhagen. What we find is that business-school-trained managers in Denmark and the United States significantly cut wages relative to their competitors; they don't give enough of the gains from productivity improvements to workers. Our evidence suggests that that's because they subscribe to the view that it's more efficient and more just to look after the interests of shareholders and try to make the corporation leaner. That shows one small part of a much bigger whole, which is that the visions

or expectations or ideas of powerful actors are playing and will continue to play a major role.

EF: Looking at all the forces that you just outlined, what do you think they mean for the way AI is likely to be adopted in the workplace and its likely effects on workers?

Acemoglu: All of these forces, I think, are important during every period, but they become particularly critical during transformative eras when orga-

"There isn't much AI-related hiring activity as late as 2013 or 2014. But around 2015 to 2016, you see almost an inflection point where many establishments in many different industries start looking for AI-related workers. So that would be our best estimate of when AI technology really started spreading in the U.S. economy."

nizational technologies are changing. We're in the midst of one of those because of the rapid spread of AI — including large language models — and other digital technologies. It's also happening because the business community in the industrialized world has subscribed to a vision of relying more and more on these tools instead of on humans.

The issue of these norms and expectations is going to become particularly important because the direction of technology is open. We could use these tools for eliminating workers, sidelining them and thus not using their unique skills. Or we could find ways of creating new tasks and new opportunities for workers.

EF: You were able to look at this question empirically in some recently published research of yours that measured the effects of AI on jobs based on online job postings. What did you find out?

Acemoglu: There was very little evidence of how quickly AI is being adopted and what its effects are. We're not doing a great job in general in the economic profession of measuring technologies. That's doubly true for AI.

So the idea that David Autor, Jonathon Hazell, Pascual Restrepo, and I had was to look at the near universe of online vacancies in the United States, which comes with detailed information about what types of tasks workers are being sought for and what kinds of skills they have to bring. We're looking at the level of an establishment like a store rather than a firm; for instance, a Burger King store rather than the whole Burger King firm. With this data, we were able to pinpoint which kinds of establishments are adopting AI and hiring AI-related skills and what else these establishments are doing in terms of their hiring.

We get a very interesting picture. First of all, there isn't much AI-related hiring activity as late as 2013 or 2014. But around 2015 to 2016, you see almost an inflection point where many establishments in many different industries start looking for AI-related workers. So that would be our best estimate of when AI technology really started spreading in the U.S. economy.

Second, we look at what are the kinds of establishments that are doing this hiring. And the answer seems to be that they're the ones that have fairly simple tasks that can be replaced by AI technology. So you don't see the creative tasks or firms that require very complex functions going to AI; it's more simple IT security, simple clerical jobs, and so on that are going big-time AI.

This, then, sort of confirms our suspicion on the basis of the prior waves of automation technologies that these technologies, especially when they spread rapidly, are destroying some jobs and not always increasing wages or demand for workers.

I should add that this research was several years before large language models, so it isn't informative about ChatGPT or GPT-4, which may have different effects. And in fact, I believe they will do different things in some dimensions.

EF: Based on all this, how would you recommend that workers adapt to the changing demands of the labor market that AI may bring?

Acemoglu: Well, that's a very natural question. And obviously, it's a good question from the point of view of a worker. But let me push back and say it's not the only question or even the right question. Because that question, when you ask it from the society's point of view, buys into the narrative that AI is an already-happening avalanche with a given direction and the only thing we can do is adapt to it.

Of course, we have to adapt to all new technologies. But I think that the question you ask has to be coupled with or even preceded by a different question: What type of AI do we want? What are the technologies of the future that would be most beneficial to society, particularly workers? I cannot imagine any technology that would be harmful to workers for a long period of time and yet would be beneficial for society.

And therefore, my view is that right now we are going in the wrong direction in the AI community. We are going in the wrong direction in the tech community, because there is no regard paid to what these technologies are doing to workers' jobs, democracy, mental health, all sorts of issues. So we really need to ask, can we redirect these technologies?

But coming back to your question, of course workers need to adapt as well. And I think workers who have skills or choose to specialize in things that one way or another are going to be done by machines are not going to do well. So I think social skills, social communication, teamwork, adaptability, and

Daron Acemoglu

■ PRESENT POSITION

Institute Professor, Massachusetts Institute of Technology

■ SELECTED ADDITIONAL AFFILIATIONS

Member, National Academy of Sciences; Fellow, American Academy of Arts and Sciences; Research Associate, National Bureau of Economic Research

■ EDUCATION

Ph.D. (1992), London School of Economics; B.A. (1989), University of York

creativity are going to be rewarded by the labor market. The way that machines augment humans, humans should also augment machines.

But make no mistake, it's not just those skills. Today, and I believe in the next 10 years, the United States economy is going to need a huge number of carpenters, electricians, plumbers, lots of people who do very valuable, very meaningful skill-requiring, expertise-requiring combinations of manual and cognitive work. It's a mistake for us to think everything is going to be digital. And it could be very beneficial for us if we tried to make new machines, including AI, in such a way that they complement electricians, plumbers, carpenters. I think that complementarity is really critical.

EF: Arguments for regulating AI along economic lines seem uncommon now. More usually, one sees arguments about AI and alignment, about AI and long-term threats.

Acemoglu: Those arguments really confuse the debate. I'm not worried about artificial general intelligence coming and taking over humanity.

EF: Why do you think economic policy arguments about AI aren't more salient?

Acemoglu: There are many reasons. I think one of them is Hollywood and

science fiction. I love science fiction, don't get me wrong, but it has conditioned us to think about the scenario in which the machines become humanlike and compete against humans.

But second, even more importantly — and this is, to me, a foundational mistake in the AI community, going back to Turing's work and to the [1956] Dartmouth Conference on AI — it was a mistake framing the objective as machines being intelligent, developing humanlike capabilities, doing better than humans. I think we should have framed the question from the beginning as a machine that's useful. We don't want machine intelligence in itself; we want machines that are useful to us having some high-level capabilities and functions.

Today, still, the way you get status in AI research is by achieving humanlike capabilities. On top of that prestige, the biggest sources of funding right now for engineering, computer science, and AI are companies like Google and Microsoft. Put the two effects together and you have an amazing bias.

And then the third is the economics profession. You know, economists are right: We owe today's prosperity to technology. We would not be 30 times as prosperous as our great-great-grandparents who lived 250 years ago if it wasn't for the huge breakthroughs of industrialization, of communication, of improvements in pharmaceuticals, all of these things. Yet that does not imply that technological change is always good for workers or always good for society. So we really need to develop a perspective of how can we harness technology for the better. But if you subscribe to the view that technology is always and everywhere good, it's like a sin to ask questions about regulation of technology within the economics profession. And if you put that together with the ideological disposition of the AI community, I think you get the current picture.

EF: Objecting to the effects of new technology on labor is sometimes

casually linked with the Luddites. As you know, the Luddites were a group of 19th-century English textile workers who responded to automation by destroying textile machinery. Setting aside their methods, what were the Luddites right about and what were they wrong about?

Acemoglu: There's a debate about Luddites, and I think the public almost always sees one aspect of the Luddites — that of the rabble-rousers who went around creating trouble. There was that; it's not deniable. But Luddites were part of a broader nascent working-class movement that was trying to articulate ideas about worker rights, worker participation in decision-making, and how work could be organized in a way that was beneficial for workers. So the Luddites had some ugly parts and some forward-looking elements as well.

But specifically in the context of the weaving machines, which is what animated the Luddites, they were right that those weavers were the losers out of technological progress. Their high wages got destroyed. They were shifted into much worse working conditions for longer hours, for lower pay in factories, or lost their jobs.

What they did not do is that they did not articulate a coherent view about how we could harness and leverage technological change in a way that would be beneficial for workers as well. But that's probably asking too much from them.

EF: In contrast, in your new book, you describe the adoption of electrical machines by factories in the late 19th century and early 20th century as highly beneficial economically to workers. Why did workers share in these gains?

Acemoglu: Why is it that electrical machinery was so beneficial and

the textile machinery of the late 18th century wasn't? That brings me to the key concept of the framework that I developed in academic work with Pascual Restrepo: new tasks. If you want to think about workers benefiting, you have to think about what new tasks they can perform. And the key thing about electrical machinery — and the Ford factory in the early 20th century is a great exemplar of this — is that it generated a whole series of new tasks.

"I am one of the signatories of the declaration that we should pause, until we understand where we're heading, the further training and the development of large language models for a period of six months or longer. Because I think right now this rollout is speeding ahead of our understanding in the regulatory framework and it's locking in a particular direction."

With the introduction of electrical machinery, production became more complex. So you needed workers to attend to the machinery and then you needed a lot of supporting occupations: maintenance, design, repair, and a whole slew of engineering tasks as well as many other white-collar occupations. So what really was beneficial both from the point of view of the workers and from the point of view of productivity wasn't the fact that those factories were substituting electrical power for some other kind of power. They were completely reorganizing work in a way that made it more complex and thus created more gainful activities for workers.

Not everything was rosy. It was hard work. Compared to today, workers were worn out. They found it very difficult to keep up with the pace. It was still much noisier than the kind of factories that we would see later.

And Henry Ford himself, especially later in his career, became zealous for anti-union activity. So it's not like saying Ford was a visionary in every

dimension. But Ford exemplified a new type of industrialization, which created new tasks and thus opportunities for workers.

EF: What policy choices might be made to see to it that AI is adopted in a worker-friendly way?

Acemoglu: Well, that's a very difficult question. I think there is no silver bullet. But let me give you several answers.

First of all, I do not argue that we should stop AI or even stop large language models, but I do strongly believe that we should regulate how they are developed, how they are rolled out, and how they are used.

For example, in the case of large language models, I am one of the signatories of the declaration that we should pause, until we understand where we're heading, the further training and the development of large language models for a period of six months or longer. Because I think right now this rollout is speeding ahead of our understanding in the regulatory framework and it's locking in a particular direction.

Likewise, in the case of the technologies that, for example, Facebook uses to promote certain kinds of content and monetize it on the basis of digital ads, there's again a question of how the technologies are going to be used. Are we going to allow Facebook to play on the insecurities of teenagers to get them more hooked or is there a better regulatory framework? I don't have all of the answers to these questions, but those are questions that we need to ask.

In my thinking, there are a couple of levers that are highly important in doing this. First of all, we need to deal with data control. All of these tech models that we're talking about with their pernicious effects are predicated on free collection and control of data by large companies. I think that's

something we need to revisit.

We also need to worry about whether the business model of the tech industry, especially in the communication area, is leading to highly negative effects; the core of that business model is digital ads. So I am in favor of policies giving alternative business models — for instance, based on nonprofits, such as Wikipedia, or on subscriptions, such as Netflix — room to emerge and become more of a viable alternative across the tech industry.

But I think most importantly, we need to change incentives of the research community and the leading firms such as Microsoft, Google, OpenAI, and so on, such that they divert their attention away from automation and control toward things that are human-friendly, meaning increased worker productivity, empowered citizens.

EF: You've argued that by various measures, the position of the median U.S. worker has become worse since the late 1970s. You suggest that automation and globalization have been synergistic in driving this trend. Please explain.

Acemoglu: Yes. I think the data show that median wages in the United States increased from 1980 or the mid-1970s onward, but only slightly. So if you exaggerate a little bit, you can say median wages are stagnant. Average wages in the United States have increased; that's partly because workers have become more skilled in the United States, so the educational achievement of workers is quite a bit higher than it was in 1980.

And if you look at some large demographic groups, such as a man with a high school degree or a man without a high school degree, you see that their real wages have actually declined significantly since the late 1970s or

early 1980s. All of that is to highlight that this has not been a good labor market for workers on the whole.

There are some groups that have benefited. If you are very highly skilled, if you have a postgraduate degree — especially in engineering or one of the other disciplines that has benefited from the tech boom — you have done very well. If you are a skilled surgeon, you have done very well. People with specialized skills have done quite well. But the majority of the workforce hasn't benefited much, and some people have lost out.

Why is that? I think technology is a major part of it, and globalization is a major part of it. Both technology and globalization have hit workers who used to be in manufacturing or who used to be in nonmanufacturing but earned a decent living. So you see this pattern of workers in the middle of the income distribution being particularly badly affected by these forces.

What we argue in our book is that the effects of globalization and technology were not inevitable. In both cases, the specific choices that we made were quite important. And particularly in the case of technology, it's about where we started — whether we are going to use these technologies for automation or for creating new tasks. And once you increase productivity, how are you going to share this productivity?

And in the case of globalization, there was the rapid flow of cheap Chinese imports and offshoring, which happened as a result of excessive attention on cost-cutting. These choices have not worked out well for working people.

EF: Shifting topics a little bit, in your 2019 article "Democracy Does Cause Growth," you and your co-authors found that democratic institutions are associated with economic growth.

Why is that? What are the mechanisms behind this?

Acemoglu: I'm proud of that paper because, for some reason, there was an emerging consensus within political science and economics that democracy was not a good system for dealing with economic problems. I think it was fueled in part by China — people seeing China's tremendous leap and saying look how well autocracy works, and at the same time also witnessing gridlock and economic problems in various democracies.

We were very suspicious of this, which is the reason we started this project. As soon as we started, we realized if you organize the data in the most neutral way, it is amazingly apparent that democracies actually grow quite a bit faster. And one way of doing that is just to look at the same country before and after becoming democratized.

Before, when they are under a dictatorship, they have a lot of economic problems. And then after democracy, it takes a while for stability to set in, but after a while a rapid growth experience exists.

Why is that? Well, one of the things that democracies do is they increase taxes; democracies raise more revenue and spend more money.

And where do they spend it? Well, some would say waste, of course; it's the nature of bureaucracy. But a lot of it goes to health, education, and public infrastructure. That's part of the answer. But also contrary to the conventional wisdom, we find that democracies are better at doing reform. They are much better than dictatorships at dealing with monopolies. They are better than dictatorships at increasing the capabilities of the workers, especially low-income people in the community. So there are a number of dimensions to the link behind democracy and growth. **EF**