

Good afternoon, everybody. Welcome to today's MIT Starr Forum: Power and Progress: Our 1,000 year Struggle Over Technology and Prosperity. I'm Evan Lieberman, Professor of Political Science and Director of the Center for International Studies, which is hosting this event, and I'd like to thank all of you for joining us today.

So the title of the talk is taken from the latest co-authored book by my colleague Daron Acemoglu, an Institute Professor here at MIT. And the questions raised in this terrific book could not be more timely. What is the value of new technology for the human condition? Should we faithfully trust that labor-saving ideas will make us all better off? How might social and political institutions mediate the effects of these new ideas?

As in his other terrific books, in this one Acemoglu once again reveals a great sensitivity to the importance of social and political power in structuring how we live and exist on this planet, and in a context in which generative AI and other technologies are diffusing with such incredible speed, we really need to listen to what he has to say. So before we get started, I'd like to remind everyone that the book, *Power and Progress* is for sale in the lobby. And I know that a lot of you have already done that. The line was great, so the sales have been good.

Please pick up a copy, and after the event please feel free to bring it up to the stage for a book signing. In addition, as per our custom, we're going to conclude with a question and answer from the audience, and we'll please ask you to line up behind the microphones to ask your one question. And I have an emphasis on one and on question-- which is not always heeded, but I will try as moderator to enforce that.

So without further ado, let me introduce both today's speaker and our discussant. So Daron Acemoglu really doesn't need introduction, but I'm here to do that, so I will. He's the Institute Professor at MIT, an elected fellow of the National Academy of Sciences, the American Philosophical Society, the British Academy of Sciences, and a member of the Group of 30.

He's the author of six books, including the *New York Times* bestseller, *Why Nations Fail: Power Prosperity and Poverty*, which was joint with James Robinson. *The Narrow Corridor: States, Societies, and the Fate of Liberty*, with also James Robinson. And this book, which he's going to talk about today.

His scholarly research covers a very wide range of areas, including political economy, economic development, economic growth, technological change, inequality, labor economics, and the economics of networks. Our discussant today, who I am sure did not rely on ChatGPT for her comments-- right, they are completely original, OK, good-- is Fotini Christia, the MIT Ford International Professor of the Social Sciences, and my colleague in the Department of Political Science.

She's also the faculty affiliate of the Center for International Studies, where she's Director-- well, effectively, of the Center for International Studies, and she's also the Director of the Sociotechnical Systems Research Center at IDSS, and she's Chair of the Doctoral Program in Social and Engineering Systems at the Schwarzman College of Computing. So I'll look forward, as I'm sure all of you will, to her comments. But, first, please join me in welcoming Daron Acemoglu to the podium.

[APPLAUSE]

**DARON
ACEMOGLU:**

Thank you very much, Evan. And thanks, Michelle, for organizing. This it's a true pleasure to be here. And I would like to point out that my partner in crime here, Simon Johnson, is somewhere here. Oh, here he is. Yes. So we have the double bill here.

So it's my pleasure to be here to share some of the ideas from our new book. And here is the book, and the title. And I think today, after so much hype and discussion about advances in AI, especially generative AI tools such as ChatGPT, I don't think we need to give a big introduction that there are tremendous and consequential changes in technology. But part of the reason why we have written this book is because we think there are some critical questions that need to be asked whenever there are new technologies, which is often in human history, and those are about the control of technology. Who controls technology, and how that actually shapes who will benefit.

In fact, for transformative tools such as generative AI, these are particularly important because there are so many different directions in which these technologies can be developed, and it is quite possible that they could bring broad based benefits or that they might actually enrich and empower a very narrow elite.

In fact, throughout history, there are examples of very consequential decisions being made following the visions of powerful agents. Today, those powerful agents are the optimistic technology leaders in places such as Silicon Valley. In the past, they may have been different people. For example, this gentleman, Ferdinand de Lesseps, was both a leading technologist of his day, and perhaps the techno-optimist of his day, because of his big belief that the world had to be opened up with big investments in public infrastructure.

And when other people thought that this couldn't be done, or couldn't be done in the way that would really enable ships to freely flow from the Suez Canal, he single handedly pursued that dream. He convinced others to come on board, he called for big technological advances when others thought that his schemes wouldn't work. And he was very successful, becoming one of the most famous figures of the second half of the 19th century.

But his belief in technology and his own understanding of where technology would go then made him blunder with huge consequences in the Panama Canal, where he completely ignored the science and the conditions on the ground, and his schemes completely collapsed, leading to the deaths of more than 22,000 people.

Now, these sorts of larger-than-life characters having a huge impact on where technology goes, again, is not something that we are unfamiliar with today. But even those things you might think are not so important because you might imagine that whatever these leaders of technology decide, there may be very powerful forces that ultimately, almost automatically, are going to bring broad based benefits.

And, in fact, a very critical part of that has to be through the labor market. Most of us earn our living through the labor market, and if any technology is going to create any type of broad based benefits, it must somehow lift people in the labor market. Here, the economics profession is very optimistic, in some sense. It is so optimistic, in fact, that a very important proposition in economics doesn't even have a name because it is so much part of our canon.

So Simon and I had to invent a name for it, and we called it the "productivity bandwagon," and it goes something like this. Technology improves, meaning our capabilities for doing things, for example, in the production process, gets better. As a result, productivity-- meaning how much we can produce-- increases. And this creates a series of economic forces-- in particular, firms wanting to go out and hire more labor because they have become more productive-- and via that channel and through the workings of the labor market, wages increase, and, as a result, workers benefit.

That type of story is the way that most economists think about why it is that over the last century, for example, as productivity went up, wages and employment went hand in hand, and brought pretty widely shared benefits. But if you go back in history, you will see there are many other examples in which things work out in a much more complicated way.

Here, we provide pictures of two of those, one from the medieval era, the other one from the early industrial period. The one on the left is a medieval technological breakthrough-- really revolutionary in terms of its conception and what it did to the production process-- windmills that massively improved capabilities in many sectors. But when you look at the data, you don't see the windmills creating this sort of productivity bandwagon that lifted the living standards of the workers or the peasants.

At the end, instead, what you see is that a very narrow group of people-- the landowners and the clergy who controlled land and the production process-- were the beneficiaries while the working conditions of farm workers did not improve much. On the right, you have even a starker technological transformation that, again, has a more complex effect. Eli Whitney's cotton gin, which, at one fell swoop, turned the US South from a complete economic backwater to the largest producer and exporter of cotton and fueled the critical phase of the Industrial Revolution and really created huge fortunes in the US South.

But the workers who actually did the cotton production-- the enslaved Black people-- did not see any benefits. In fact, their coercion increased, intensified. They were moved to the Deep South, where working conditions were harsher and longer. And, again, no sign of a simple productivity bandwagon.

But you might actually think that those two are handpicked examples and the bigger process that really defines our age-- the one that started sometime in the middle of the 18th century with the application of scientific knowledge and artisanal knowledge to the production process, the Industrial Revolution-- is very different. After all, when we talk about the dangers of AI or other automation and digital technologies for inequality for wages, one response that Simon and I get is, well, are you saying that this time is different?

And, in fact, this is the reason why the subtitle of our book is "Our 1,000 Year Struggle over Technology and Prosperity." No, in fact, we're not saying that this time is different. This time is very similar to what went on in the past. There has always been this tension about who controls technology and whether actually the gains from technology are going to be widely shared.

So, for that, let's turn to a re-evaluation of what happened in the British Industrial Revolution. And the story turns out to be much more complex than the simple story that the Industrial Revolution increased our productivity and we are all the beneficiaries of it. Yes, indeed, we are immeasurably more prosperous, healthier, and more comfortable today than people were 300 years ago, but, again, there was nothing automatic about it. In fact, the path to that improvement was much more circuitous.

The early phases of the British Industrial Revolution were characterized by something that's going to have resonance for today as well, automation. Meaning the application of technology for simplifying and reassigning tasks previously performed by humans to machines. This was visible in the textile industry, especially in the weaving processes. And weaving, which used to be done by people with hand looms or in their houses, became something that migrated to factories.

Did that improve productivity? Yes, the evidence is very clear that productivity increased as a result of that. But, in fact, the benefits of that, again, were not widely shared. The evidence is not completely clear, but most of it suggests that real earnings of workers in Britain during this period did not increase, while their working hours probably increased by about 20%. So their hourly real wages declined.

But even worse, the working conditions of people were much harsher-- much less autonomy, much less independence for workers in factories than they had as independent weavers. And working conditions are only one part of the story, of course. People also care about where they live, their health conditions, the pollution around them. And in all of these, the situation was decidedly not so good for the British working people.

So the factories were emblematically much harsher places than they were used to. This is captured by-- symbolically, at least-- by the picture we have on the left, which is the idea of a very famous economist Jeremy Bentham-- the panopticon, famous from Michel Foucault's writings or the movies of *Guardians of the Galaxy*-- but the way that Jeremy Bentham thought of this was a highly efficiency improving technology because it would enable employers to monitor workers better, or teachers to monitor students, or guards to monitor inmates better. And what could be wrong with that?

But, of course, that was actually what employers were very interested in doing in modern factories to the expense of workers, who were forced to work very long hours under very harsh conditions. At the same time, the cities in which workers concentrated became complete cesspools, much less healthy, much less comfortable living places. And life expectancy during this phase at birth may have fallen as low as 30 years old, a terrible number at a time when conditions economically were actually improving for factory owners, for example.

Now, this is, of course, not the end of the story, and there is some truth to people who say that, look, we have benefited so much from the Industrial Revolution. Indeed, we have. But the process for which needs to be understood, how it is that we benefited from it after this early phase.

And this early phase was not a short one. It may have lasted by about 100 years. The beginning of the Industrial Revolution is not clear, but you may date it to around 1750, and, by the 1840s, conditions were still very harsh, not just for workers in textile factories, but in every sector of the British economy, including coal mining, another one of the dynamic sectors of the economy where children as young as six were working 18 hours under hugely harsh conditions deep in mines.

But, of course, you might think, this is all history. Today is different. Yes and no. Today is different because today we are in an economy in which, for a while, we got used to a very different type of sharing of the gains. And these two charts that I have up here summarize both the ways in which today is different, but also the ways in which there might be some parallels to those older, not so good times.

What I'm plotting here is for men and women separately and for five education groups-- starting from workers with less than high school in dark orange all the way up to workers with a postgraduate degree in dark blue-- how their real wages have evolved over the last 60 years.

And, in 1963, everything is normalized to zero, so you can follow the cumulative change of the wage profile for each one of these 10 demographic groups. What you see in the 1960s and early 70s is actually a continuation of a trend that is also visible from other data sources in the 1950s, which is one of shared prosperity.

Real wages for all 10 demographic groups that I'm showing you here are growing in tandem, more or less on top of each other, from the fact that all 10 of these curves are growing very sharply. In fact, they're growing very, very rapidly about 2.5% every year in real term, which is a really remarkable rate of growth.

At that rate of growth, starting from poverty, in two decades you can reach sort of a much more comfortable level of living. But that period of rapid wage growth and shared growth comes to an end sometime around 1980 or the late 1970s. From then on, you see a much different picture. These curves are fanning out, indicating a greater level of inequality, a much greater level of inequality.

But even more jarringly, you see that the real wages of low education groups, especially men, but also for women, are actually declining. So the green line, which is for high school graduates, the orange line, which is for high school dropouts, are actually sharply declining from their values in 1980.

Not only the gains are not being shared, but some groups are losing out. So this picture therefore poses two sets of questions. How was it that the US economy-- in fact, it turns out, much of the industrialized world-- reached some sort of compact in which there was this fairly shared and rapid growth in the late 40s, 50s, 60s and 70s, and why did it come to an end?

So let's try to understand both of these. And let me skip this part, which shows that inequality is also increasing in other countries as well. Let me skip that, and instead just get to the bottom of the theory that Simon and I tried to develop in the book, which is that the productivity bandwagon is not a force of nature that applies under all circumstances automatically and with great force, but it is something that's conditional on the nature of technology and on how production is organized and how the gains are shared.

In particular, what changed in the second half of the 19th century in Britain and then continued in the United States and much of the industrialized world in the first part of the 20th century, and came into even greater fruition in the three decades that followed World War II, was this process of shared prosperity built on two big pillars. One is new tasks, and the other one is worker power. Of course, both of these terms need to be defined, and they may be a little bit simplifying.

But the new tasks are critical. Automation that I mentioned, which is the substitution of machinery and, today, algorithms for the labor of humans, has always been with us, or at least has been with us in great force since the beginning of the Industrial Revolution. Automation is a major force for increasing productivity, but it does not create the type of shared prosperity by itself.

Because, after all, what automation is about is to take tasks away from workers and have more machinery do it. So it reduces the importance of labor. It also reduces the labor share in output or in national income. So if we're going to have shared prosperity, automation needs to be coupled with something else. And that something else, critically, turns out to be new tasks. New activities in which human labor is critical that reinstates workers centrally into the production process.

So throughout the second half of the 19th century and the early 20th century, we see these as very important determinants of what's going on in the labor market. Emblematically, for example, captured by the picture up there, which is from Henry Ford's motor factories. Henry Ford was a leader in applying new technologies, including decentralized electricity and assembly line type technologies early on, and that was absolutely revolutionary, completely changing the car production, making cars affordable for the masses.

And a very important part of that is the use of machinery to do tasks that were previously done by labor, automation. But if you look at the Ford factories, that wasn't the only thing that they were doing. At the same time as they were introducing automation technologies, they were also creating new tasks for workers.

So it is no surprise that, in this picture, what you see is the advanced machinery together with the workers. Workers are now performing more technical tasks. They're operating this machinery, they're engaged in design, inspection, and other quality control activities. And if you actually look at the factories of the early 20th century, you'll see that production workers are joined with non-production workers, clerical workers, that are very much engaged in planning and other aspects of the production process.

It was this double process that was so important for the beginning of shared prosperity. Automation, which increases productivity, but also new tasks that give another boost to productivity and also create reasons for workers to share in those gains. But even new tasks by themselves are not enough, because if workers are making a major contribution to productivity but they don't have the power to take a share of that, they may go down the same path as the Black enslaved workers with the Eli Whitney cotton gin.

It may not be in the interests of the firms to share those gains with them, and they can get away with it because they have all the power. So, actually, a balanced distribution of power in workplaces and in society is also a critical part of it. That's why the second picture on this slide is one of the emblematic strikes in the auto industry that was still a leader in establishing the labor movement during this period, the United Auto Workers strike at General Motors in 1937.

So it was this double process that was so important, but not just for the early car industry, but, in the 19th century-- the second half of the 19th century that I've been referring to in Britain-- what was so distinctive about it-- when you think about what went on, the direction of technology changed. Completely new set of technologies in railways, in steel, and chemicals that were much more important for improving the productivity of labor and introducing new tasks for labor-- this was embedded in an society that was democratizing from the times of the early 19th century, where even the middle classes didn't have the vote, now universal male suffrage, and then universal suffrage came to Britain. And, also, the labor movement, which was completely banned and heavily prosecuted up to the last quarter of the 19th century, became a staple of British workplaces and was a very important part of improved working conditions and wages.

Now, this was about why shared prosperity's foundations were laid starting sometime in the middle of the 19th century, and then continuing to the 20th. But then why did it come apart sometime in the 1980s? And I think the same two processes now playing out in reverse are the key actors in this.

And to explain that, I'm showing here a modern car factory, which looks somewhat different, or quite different, from the Ford one. You again see the advanced machinery. Now the advanced machinery takes the form of robotic arms. But conspicuous in its absence are the workers. The workers are no longer playing a central role. The automation is rapid, but the new tasks haven't accompanied it.

So too much focus on automation but not enough on creating new tasks is the technological part of it. But accompanying that has also been an institutional change. And that institutional change-- sorry, before that, let me actually show you this figure to substantiate the claim that I made. This is a figure from work that I have done with Pascual Restrepo.

What it shows is something that's akin to the first graph that I depicted, the evolution of the real wages of the 10 demographic groups. Now a little bit more detailed, demographic groups distinguished by age, gender, education, and ethnicity-- each one of these circles refers to one of them. On the vertical axis, I'm showing you the cumulative change from 1980 to 2016.

So that period in which some groups were experiencing wage growth, other groups were experiencing wage decline. You can see the same thing from here, from the fact that many of these circles are below the zero. Those are the demographic groups that are experiencing wage declines. And on the horizontal axis, I'm depicting the extent of task displacement that a demographic group has experienced during this period. Namely, what fraction of the tasks that this demographic group used to perform across industries and across occupations in 1980 have since been automated.

So you can see that, for some groups-- mostly those like us who have postgraduate degrees or very high levels of specialized skills-- those numbers are very close to zero. We have not really suffered much automation of the tasks that we used to perform that are much more creative, much more problem solving, and high level. But if you look at those for high school education or less than high school education demographic groups, shown, for example, by purple and green, you'll see that up to 25%, 30% of the tasks that they used to perform have since been automated, and those are exactly the groups that have suffered the wage declines.

In fact, this regression line explains about 60% to 70% of the variation in equality between groups in the United States. So this is the automation part. But automation has had a very big effect and, in fact, its path was very much shaped by institutional changes in the US labor market. And those institutional changes have been in the direction of declining worker power.

And to understand declining worker power during this period-- in the same way that if you wanted to get to the details of understanding the increase in worker power in the 19th century, you would need to think about both how ideology has shifted and how organizations have shifted.

And both of those have gone against workers during the post-1980 era. One is-- perhaps I'm giving too much credit to our fellow economist Milton Friedman-- but the rise of new corporate visions which elevated managers doggedly working for shareholders and ignoring everything else. This was the beginning of the shareholders value revolution, or Milton Friedman's statement that the only social responsibility of business is to maximize its profits.

And that was coupled with the erosion of worker powers, for example, emblematically during the defeat by Ronald Reagan of the professional air traffic controllers strike. So these two changes together both shaped the way that managers wanted to approach how to run their business. For example, monitoring workers more tightly or automating in order to cut labor costs was completely welcomed because it would increase the returns to shareholders. But also there was no resistance to them from organized labor because organized labor was getting weak during this period.

Now, this is all before AI. Can the age of AI change it? Yes. The age of AI can change it because, at some level, if you look at the details, the promise of generative AI-- some of it is hype, but some of it is reality-- is that it can actually be a tool in the hands of workers. But if you look at the reality of it, you also see major roadblocks towards that kind of change.

What makes today such an important point in this type of discussion is that there are transformative and very consequential choices ahead of us. And, again, Simon and I think that one way of framing this is in terms of different visions. One vision about where AI and digital technologies in general are going to go-- emblematically summarized by that picture at the top, or the Turing test-- is towards autonomous machine intelligence. Meaning machines become more and more autonomous, more and more intelligent, and they start doing more and more of the tasks that humans used to do.

It won't take much imagination to see that if that's the emphasis, we're going to have a lot more automation. But, in fact, if we are right that automation doesn't create the foundations of shared prosperity, that spells trouble. But that's not the only way in which digital technologies were conceptualized, and it's not the only direction in which AI can go.

Long ago, many computer scientists understood a very different way of using technologies, which Simon and I call "machine usefulness" to contrast with machine intelligence. The objective is not to make machines intelligent in and of themselves, but more and more useful to humans.

Engineers such as MIT'S Norbert Wiener, JCR Licklider, who was briefly at MIT as well, or Douglas Engelbart tried to articulate both the philosophical foundations and the technological realities of this vision, and out of this came many of the technologies that we rely on. For example, when you use your smart phone, the menus, or the computer mouse-- which was revealed by Douglas Engelbart in a very famous event called the Mother of all Demos-- or hypertext, all of these came from an effort to make machines more usable and more useful to humans. And, in fact, AI could pursue that path.

Now, the problem, in fact, it's not just one of distribution. If you overemphasize automation, it's not that you're going to get huge productivity gains and they're just going to be unequally distributed. In fact, there is every danger that overemphasizing automation is not going to get you much productivity benefit either. And this is the concept that Pascual Restrepo and I and Simon and I tried to capture with the label "so so automation."

What you're trying to do is you're trying to get machines to do things that humans are pretty good at. So when you do that, you don't get a huge productivity boost because humans were doing it fine, but you get big distributional benefit or distributional costs because you're sidelining humans. Firms may become a little bit more profitable, but a lot of workers lose out. And self-checkout kiosks or excessively automated customer service, all of those are examples of so-so automation where the productivity benefits turn out to be not so much as people were hoping.

Now, of course, generative AI and ChatGPT could change all of that. So we asked ChatGPT itself whether generative AI could reverse these things. On this one, I think ChatGPT was quite on target. Perhaps, but probably not. It's not a magic solution. If generative AI is used to replace workers instead of support them, it could have negative consequences.

Now, we don't know whether it knew the answer that we wanted to hear or it read some of our papers, but we agree with this answer. Labor market consequences and inequality are not the only things we have to worry about. One of the other trends since the 1980s but, again, accelerating with generative AI, is about who controls information.

After all, even in the Industrial Revolution, it wasn't just automation, it was also how the modern factory system changed the method of control and who was in charge and what they could dictate. So one of the things that we are seeing with more and more digital technologies is surveillance. Surveillance in workplaces, surveillance in political views. Now, it takes different forms. In China, you may be more worried about it because it's in the hands of the government-- the social credit system, or facial recognition cameras everywhere where protests could one day break out.

In the United States, it's companies. It's Google, Facebook, Amazon that have all of this information. But, at the end of the day, in the book Simon and I argue that both of those are pernicious. It doesn't matter who has your control. As long as that information can be used without any constraints, it's going to be anti-democratic and it's going to be inequality inducing.

This has so far been a lot about the developed world. I have given examples from the US, a little bit about Europe, and, in fact, that's a national focus for Simon and I because we want to trace the counters of new technologies and how they are used. But let's not forget that new technologies that are developed in the United States and in China are going to be used throughout the world and, in fact, the International division of labor is already being reshaped by automation technologies.

One of the patterns you see around the world is that a lot of routine activities that were automated in the United States are also being automated around the world, or at least the amount of activities or labor that is assigned to these production functions is declining. There is every danger that AI, if it goes down the automation path, could be a highly an equalizing technology around the world. And, again, some of the surveillance implications are global as well.

The recent work by some of my colleagues here-- Martin Barrera and David Yang and others-- shows that Chinese companies are already exporting anti-democratic monitoring technologies to more than 60 non-democratic countries around the world.

So this is all potentially depressing because it says there are big dangers of [INAUDIBLE] inequality and democracy. But, from the beginning, I tried to frame this as saying, well, these are transformative choices because there is no necessity that AI is going to go one way or another. There is a high degree of malleability to all technologies, and that's doubly true for AI.

And if we make the wrong choices, they could have damaging consequences. If we make the right choices, they could be much better for society, it could be much better for workers, it could be even better for democracy. So the question is what we can do. Simon and I, in the book, suggest that we need a three step process for thinking about change.

The first one is changing the narrative. Our modest hope is that this book is a small contribution to changing the narrative, moving away from blind technological optimism-- everything is going to work out. This time was no different-- to try to understand how things have worked in the past, when things go right, when things go wrong. And part of changing that narrative is to also recognize that things are more likely to go right when there are more voices rather than technology being in the hands of some powerful actors, be it Ferdinand de Lesseps or Sam Altman and Elon Musk.

So the first is a change in narrative. But changing the narrative is not worth that much unless there are institutional and other developments that actually turn that into actual action and policy changes. So that is what we mean by countervailing powers. So part of the reason why things were different in the 1950s and the 1960s was because technological choices-- and how those choices were-- how those gains were being shared, was embedded in an institutional framework where government regulation was important and where there was civil society and labor movement constraints on what companies could do.

Some of that needs to be recreated. It has to be in the form of a new labor movement, perhaps other forms of bottom up organization for civil society, and also government regulation, especially in the field of AI. And the central idea here that Simon and I emphasize is that redirecting technological change has to be a major part of both the efforts of the labor movement and of government regulation.

Technology has many potential directions in which it can go, and there is no guarantee that the completely unfettered market process is going to choose the socially beneficial direction. There were many people, such as this gentleman Ted Nelson-- not just Douglas Engelbart and others-- who thought that the personal computer and other digital technologies would be fully liberating choices, both for workers and for citizens.

In the end, that's not the path that we ended up on. But that was a choice. It wasn't because they were completely wrong in thinking that technology could be a decentralizing force, it could be a tool in the hands of the workers, not just of corporations. And, in fact, Ted Nelson very much anticipated this when he was writing. On the one hand, he was optimistic, but on the other hand, he was very much emphatic that large corporations such as IBM would try to control the technology, and that would push it in a very different direction.

But all of this raises another question, which-- and I will end on this. There is some degree of optimism in saying that we can actually redirect technological change, and that's a social choice. Because the counterargument is, no, technology is a fully organic process. Every time you interfere with it, it's going to end up in your face.

Well, of course it is an organic process. But the fact that it's an organic process doesn't mean that it cannot be steered within bounds. And one example where you see that is in the energy sector, which is not, of course, something to be proud of. We are very much behind in combating climate change. But, today, if you look at our ability to generate clean energy, it's miles apart from where it was in the mid 2000s or even in 2010.

For example, various different types of solar and wind technologies are, today, cost competitive with fossil fuels, whereas they were about 10 times as expensive as fossil fuels as recently as 2010. How did that happen? It happened because there were some subsidies and some regulations about clean energy, and it was also a civil society movement. More pressure from consumers who wanted cleaner products, more pressure from civil society for companies to clean up their act. And even a modicum of that type of pressure led to a complete redirection of technology. And the reason why we end the book with this example is precisely because Simon and I think this can be done in the realm of production technologies with even more consequential consequences than energy perhaps. Thank you.

[APPLAUSE]

**FOTINI
CHRISTIA:**

Hi, everyone. First of all, I want to thank CIS for the opportunity to engage with Daron on this terrific book co-authored with Simon. It's great to have you in the room. It's a true honor to be here. As you saw from this terrific presentation, this is a book about the forces of technological change, but also all the challenges that this ushers in. And, I mean, it's an intellectual tour de force taking us over 1,000 years and bringing it to the very present, focusing on the challenges of digital innovation and the new era of AI.

So, unsurprisingly-- and now you know it's not just ChatGPT-- the first thought that came to mind was how relevant these topics were even in ancient times. So it's not just the 1,000 years, but these were real preoccupations also for the ancient Greeks. And in their myths and in their philosophical writings, they were really challenged by this idea of the wonders of innovation and technological change and what they meant for creating societal hierarchies and structures, but also a lot of cautionary tales of the responsibility that comes with these great wonders and these great capabilities.

So, if you may indulge me, I'll touch on three myths that came to mind, and I hope I'll briefly travel us back to the Greek mythology books that we all engaged in childhood. So the first is the myth of Daedalus. He was a master craftsman and inventor that was brought to the island of Crete by King Minos in an effort to house the Minotaur, the half-bull half-man that was a pretty dangerous being. And, of course, the innovation there was the labyrinth that enabled the housing of the Minotaur, but it was also-- it created a big inequality, and that was in the form of the human sacrifices of the Athenian youth for whom this became a death trap, the actual labyrinth.

The second myth is the related story of Icarus, the son of Daedalus. After this great craftsman fell out of favor of the King, he was closed up, locked up in a tower with his son. And, of course, he wanted to engineer their great escape, and he created wings out of feathers and wax. And while they were flying out and were actually making their trip back, Icarus was not listening to his father and decided to fly very close to the sun, which actually led to his kind of plummeting in the Aegean Sea and dying.

So what was interesting-- I mean, this is seen often as a myth about hubris, and very much hubris around invention and innovation. And, I think, thinking to contemporary times and kind of youth and hubris, I thought of cases like Theranos and FTX that I think are quite associated with some of these startups.

And the third and last Greek myth may be the closest to what we're discussing here, is Prometheus and his decision to steal fire from the Gods and share it with the humankind. And, in that sense, Prometheus was kind of the ultimate equalizer, maybe a union organizer of his time of the Titans, deciding to kind of challenge elite dominated decisions over technology. And he took this great leap, but he was also punished, very, very much so. I mean, he was chained-- eternal punishment, being chained on a rock and having an eagle eat your liver every day. And this would go back, and every day he would get rejuvenated and eaten again. So I wonder if there is an interesting lesson there, kind of the courage and self-sacrifice that may be required in order to rein in some of these really big interests.

And apart from myths, I mean, Greek philosophers were also very preoccupied with these themes. Plato, in *Republic*, talks about a controlled technological environment where the main goal is the pursuit of the good life, implying that some of these technologies are actually distracting from this goal. And then Aristotle, in turn, acknowledging, obviously, the great importance in society of having-- I mean craftsmanship and technology-- but distinguishing between technologies that sustain life versus those that actually enhance living, which I think is an interesting distinction.

I mean, so this is clearly a relationship that has been kind of a great topic of importance throughout life, and I was hoping to pose three sets of questions. I hope Evan will not block me. But this is kind of a bit of a warm up for our discussion. I'll pose all three, and then you can take them one by one.

First, the book tries to suggest that there is a certain heightened energy and urgency about digital innovation and AI that maybe did not exist with other past cases. But I wonder if this may have been what it felt like at the time for all the others, what it felt like at that particular point in time. What is different, actually, between the Gilded Age-- Rockefellers, Vanderbilts-- and the Bill Gates, Jeff Bezos, Elon Musk of our times? Why is the invention of railroads and steel at that time so different from digital innovation and AI now?

It seems like in every past case in history-- and I don't know if that makes me a great optimist-- but humankind kind of, in the long run, benefited from this technological change, managed to improve overall standard of living, averted disaster, moved on, and then went into a new cycle of innovation. So why and how is our case now different? And I particularly wonder about this in the context of the work around climate crisis and the environment, which you highlight very articulately in the book, and which seems to give us a certain kind of roadmap and recipe of how to think about this and how to move forward.

My related second question is whether it is technology, really, that creates the inequality, or lack of proper governance over technology. And, specifically, this is a question about understanding how technology relates to societal, economic, political structures, and basically kind of the role of the broader ecosystem of the cities and the state and other institutions in addressing these inequalities. So I wonder, for example, what does a union-- an effective union-- look like in the 21st century?

And I'm particularly interested in-- especially now, I think President Biden just joined the picket line in Michigan where the United Auto Workers Union is on strike against the three big auto producers here. So what does it mean to have an effective union in the 21st century, and what does it mean for collective action among citizens to be effective in this case?

So there is a new book by Frederick de Boer that I was just reading about on how elites ate the social justice movement-- this is the title of the book-- and he talks about the failures of some of the most recent movements, like the MeToo movement, the Occupy Wall Street movement, the Black Lives Matter movement. He basically claims that they may have succeeded in symbolic change among a lot of the academic elites and some of the more bourgeois citizens here, but they didn't really manage to make changes for the average citizen.

And he says that a lot of it has to do with the fact that they didn't get any legislation actually passed on these issues, and he also attributes some of the failure in the way these movements were actually structured, which is very different than some of the movements that we know from the 20th century, like the Civil Rights movement, for instance.

Beyond citizen movements and unions, I was wondering about the responsibility of state more actively in terms of regulation and the creation of social safety nets. I think there are some people out there who may even argue that the state has already been co-opted and what is regulating is in favor of the interests of these big elites and corporations and not really actually regulating in ways that are protecting people and citizens in terms of these inequalities. Also, what is the responsibility of international institutions?

We know Europeans are very keen on regulating on these issues on the EU level. So can we think about this differently in that context? Or academic institutions, like here. I mean, you made it clear that a lot of us that actually get money from these big tech companies need to be very clear about where our independence may potentially end or be compromised, and I wonder, what is the role of other initiatives, like ethnic and social responsibilities of computing initiatives, or other efforts within academia, to try to think about how to move beyond these inequalities.

And as the last question-- this is a little bit more about the democratizing of technologies and this idea that they have also flattened inequalities to a certain degree. So, for instance, there are some instances in the Global South where they have been faster adopters of technologies than we have, and how the openness of digital education, information has actually enabled them to make big strides. They've also been very keen on some of the AI health care tools, for instance, that they've taken on in some of their national health systems.

And even among educators, among creators, among artists one can see how technology has been an enabler. And can we say that maybe we can accept that there will be a certain level of inequality that is required to keep wanting to have change and move on towards progress, and how should we look at this in terms of the Global South and the Global North?

DARON How many hours do we have?

ACEMOGLU:

FOTINI Sorry, I mean, literally just to open up some thoughts of conversation. And we could even just turn to the audience right away.

CHRISTIA:

DARON ACEMOGLU: No, no these are-- Fotini, thank you for these very erudite and extremely far ranging comments and fantastic questions. But I will have to be brief. Let me first say that Simon and I used to have Prometheus as the beginning part of the book and try to tell the story of responsibility. Actually Prometheus himself is sort of a conflicted character, hubristic and responsible at the same time. But then we decided A, it was going to be a little bit tortured, and B, we would probably be caught by somebody who knows Greek mythology much better than we do.

FOTINI CHRISTIA: And then Oppenheimer, I feel, had already taken Prometheus. So there you go.

DARON ACEMOGLU: Well, we also had Oppenheimer and Szilard. We had the contrast between Leo Szilard-- if we had known that the movie was coming out, we would have kept that.

FOTINI CHRISTIA: Prometheus.

DARON ACEMOGLU: But absolutely. You're right. And that's part of the reason, is we think these issues are as old as humanity, that there were concerns as early-- and even probably earlier than-- ancient Greek civilizations about control of technology. So they are very much with us.

Coming to your questions. It is possible to read history by saying, look, we've had some hard times, but then we bounced back. But Simon and I don't think that's the right story, partly because A, there are many examples in which we did not bounce back. Or even when we even when we bounced back, it was far from automatic.

I think you cannot tell the story of the medieval Europe as one of bounce back. There was no sort of process towards sharing of those gains with farmers or with the farm laborers until the whole system collapsed. And when you look at the two examples that you hinted at-- the British Industrial Revolution and the rapid industrialization with new industries in the United States-- in both cases, there was a very, very radical change in institutions, and that was very far from a foregone conclusion.

The progressive movement, I think, even its participants were surprised that it could actually form that coalition and succeed. And the cards were very much stacked against democratic reform, the labor movement, and all of that in Britain. One example we mentioned briefly in the book because it's so telling, that, in the 1840s, the Chartists in the UK collected 3 million signatures. I mean, can you believe 3 million signatures in a time when you don't have any of the modern communication technologies? With the demands that are so tame in terms of universal suffrage and some basic sort of rights.

And they were so careful not to be labeled Socialists. And the response of Parliament was completely turn down all of these demands and jail all of the Chartists. So there was no automatic institutional process that could say, oh, yes, we're going to bounce back from this. So that is the sense in which there is both caution and optimism in what Simon and I are telling.

The caution is, there's nothing automatic here, but optimism is, yes, it is possible to do it, but it's not something that's going to happen by itself. And when we come to the question of is this technology or is it institutions-- in fact, this is a big debate in economics. Inequality-- what's the role of institutions, what's the role of technology? And, at some level, it's a false dichotomy.

It's a false dichotomy because I think, first of all, the two interact in very complex ways. But most importantly-- and this is why we put so much emphasis on what the objectives of the corporations were, what the vision of the tech industry is-- it is institutions, regulations, and social norms that shape the direction of technology. But we are also insistent-- and this is partly the research that I've been doing for many, many years-- that technology does have a really important role in understanding inequality.

So there is sort of a left wing narrative, for example, that everything is just about the union movement and decline in minimum wages. But the facts just don't line up with that, and that's the reason why I showed that chart. That chart is rather striking because it shows just the automation part of technology, which Simon and I think is the most important one, really is centrally important in explaining what's been going on in terms of inequality in the United States.

In other countries, it's a little bit more of a complex picture because institutions are very different and they interact with technology. But the direction of technology, especially the automation focus, is absolutely central. And that's why redirecting technological change has to be quite important. Yes, indeed, railways and steel and how they were used was critical. Again, that's the redirection of technology. But, again, going back to our overall interpretation, even that is not a foregone conclusion.

Companies are going to have many different choices about how to use technologies. And the reason why we are so insistent on digital technologies is because they have been used in a very particular way. The conception of people like Norbert Wiener, JCR Licklider, Douglas Engelbart, Ted Nelson, is precisely-- our emphasis is precisely because there was a different path of how to use these digital technologies that was more targeted at improving the productivity of workers and, therefore, as a result, would have been more beneficial for workers.

But at the end, digital technology is being used for automation, being used for surveillance. I think those are the choices that we have made, and those are very consequential choices. Yes, indeed, there are parallels from the past, but I think digital technologies really amplify these things. And if we come to AI I think the parallels-- again, this is why the book is sub-entitled "1,000 Year Struggle"-- but there are also some unique features of AI.

One is the speed of change. To the contemporaries, the introduction of textile mills would have appeared completely revolutionary. But even then it was a slow process over many years. For example, when you see the weaving industry, modern factories are coming up, but there are still thousands of handloom weavers. It takes several decades.

With AI, the speed of change is very fast, and it's also very pervasive. The potential-- again, some of it is hype-- but the potential is there that AI tools are going to be applied across many industries. So that, I think, raises the stakes, both about redirecting technological change and the regulation. So that's part of the reason why we think this is an epochal time in terms of making the right choices.

And it is also the reason why, when we come to your question about the labor movement or the general regulation, we think the focus has to be on technology, or, at least, a very important part of the focus.

Today, we are going through a spring of strikes. Some of it is understandable. There is a pro-union, pro-worker president. Wages did not increase while unemployment is low. And it is an interesting time. But with the exception of the WGA, the Hollywood strike, I think a very, very important absence is the discussions about technology and how technology is going to change things.

UAW, their attitudes to technology is just say no. I mean, if you don't come up with a solution about how we can transition to electric vehicles but still make that better for workers, I don't think the Union labor movement is going to be successful. Or let me put it another way.

How many leaders of the labor movement have really invested in understanding AI and thinking about how it is that AI can be used in workplaces while that's good for the labor movement? I think that's the part of it. In that sense, the WGA was very trailblazing because they made AI as one of the key topics, and they have the power, because of their visibility because of their special skills, to actually have a say.

And, in fact, probably this agreement is a success for the Writers Guild of America. But even in that case, I think the WGA needs to be congratulated on making AI such a focal issue, but they did not come up with a positive vision of how to use AI. I think, ultimately, the labor movement-- and you might say this is why we may need a new and refashioned labor movement-- needs to come up with new ideas about how to use these new technologies that's good for workers and ultimately acceptable for bosses.

So the WGA, in my opinion, should have articulated a vision of using AI that's good for its members-- that's good for the productivity and the quality of the entertainment industry-- that was different than what the Hollywood studios wanted. And I think that would have been much more powerful. I think that's what we need from the labor movement. And the broader things-- the Boer book, yes, I think Boer is 100% right.

In fact, James Robinson and I wrote an article published by, of all people, in an edited book by Michael Lewis on the Occupy Wall Street saying at the time that it wouldn't amount to much precisely because it did not have a plan for institutionalizing. So the progressive movement's success is that they wanted to change politics, not just protest.

And I think that's what's missing. And the labor movement or a broader democratic movement targeted at changing regulation, changing labor market institutions, changing technology, I think, needs to be embedded in a bigger institutional setting. And then the final thing I'll say about the international dimension.

I think the international dimension is key. After all, one of the things that's quite obvious from this discussion is that the choices that are being made in the United States and in China are going to have sweeping implications for every country around the world. But where is their voice? I think the biggest issue is that if you want to have a direction of technology that's responsive to global needs, including the needs of more than 4 billion people who don't live in Europe, China, and the United States, then you actually need their voices to be heard.

So we need international organizations or new vehicles for the emerging world's voice to be heard because otherwise there's a real danger that AI is going to go in an inappropriate direction for their needs, and, worse, it's going to become a more and more powerful surveillance technology that can be unleashed on the populations of these emerging democracies. And we're already seeing that. Thank you for those fantastic questions, Fotini.

[APPLAUSE]

EVAN LIEBERMAN: Daron and Fotini, that was terrific. But now we have almost a half hour to take some questions from the audience. So you'll see there are two microphones lined up on either stairwell. So if you could just line up in those, and we will-- if we have a queue on both, we will not discriminate against either side but rotate. And, again, if you could just ask a question, and just one-- and I see we have someone there, so you can start.

AUDIENCE: So my name is Shahzad. I'm a grad student at Tufts. So you showed us some data on wage inequality and how it was brought about by task displacement and by educational attainment levels.

But, at the same time, we all know that technological advancement brings about huge differences, huge improvements in life quality, right? So, for example, now we can connect with your loved ones 1,000 miles away in an instant. Now we have AC, right? So I'm wondering, would we have seen a different picture than the one presented in the book if we looked at the net effects of technology instead of focusing just on income effects?

**DARON
ACEMOGLU:** Absolutely. I think the effects of technology are very rich and complex. But I think the general principle that we want to push in the book is that for all the technologies, including those that affect health and very different dimensions of life, how we use them and how we develop them is of first order influence. So, first of all, I think some of the technologies that we are talking about are actually making many of the goods that we enjoy cheaper. That's already taken into account in the numbers that I have given you, imperfectly, albeit. But the Bureau of Labor Statistics makes an effort to take into account what's going on with the cost of living.

So despite the fact that we have benefited from cheap goods, the real wages of some groups are declining. I think health is a super interesting area. Tremendous, tremendous advances in health technology from which, again, we have all benefited-- antibiotics completely revolutionized everything. Beta blockers.

But, actually, the picture is even complex even there. If you look at the US data, it's a striking thing-- it's not COVID, so let's stop it before COVID. From the mid-2010s to 2019, life expectancy in the United States is falling. And is falling because low income Americans are actually experiencing worse health conditions, a higher mortality.

So, again, despite the fact that we have these tremendous sort of advances-- nanotechnology, targeted medicine, and some people are spending billions of dollars to reach 200 years of age-- people are dying sooner and sooner in the United States. So, again, I think there are choices that we have to make in order to improve how those technologies actually impact the way in which people live and benefit from them.

AUDIENCE: Do you see any potential for the development of worker ownership in America and elsewhere?

**DARON
ACEMOGLU:** That's a good question. That's probably much more radical than the things that Simon and I are advocating. Absolutely, there are places around the world where you see more collectives, and there are a few in the United States. But I am not sure that that's going to be-- that's got the potential to become the model organization in the United States. And even if it were, we would need to have more of a discussion about the pros and cons of that.

But I think the halfway house that Simon and I are suggesting is workplaces in which labor is treated better, it has more of a voice, for example, through labor organizations, there are limits to how harsh the working conditions are, how intrusive monitoring is. I think that does not require as radical a step. And especially if you couple that with efforts to increase worker productivity rather than just sideline labor, I think we're already quite a big difference from the current environment of how labor is being treated and where the workplaces are going. Thank you.

AUDIENCE: Hi, I'm Brian. I'm a graduate student at the Kennedy School. So early on, you mentioned that we need new tasks to replace the tasks that might be displaced by automation. And my question is, what are the characteristics that you envision for this new tasks? The reason I ask is because I can see it going two ways. Either it becomes more advanced tasks, which requires a reskilling of the economy, or it's more blue collar non-routine tasks, which would be a decline in working conditions. So I want to hear your thoughts about this.

**DARON
ACEMOGLU:** Thank you very much. Yes, I wasn't sufficiently detailed on that. Thanks for bringing it up. If you look at the 1950s and the 1960s, you see both kinds of new tasks springing up. For example, when you can-- my colleague David Autor and his co-authors have done that, and Pascual Restrepo and I have done some similar analysis as well. So you see various measures of new tasks, and you see that both in manufacturing, especially a lot of new blue color tasks, and you see a lot of non-manufacturing tasks in offices that are springing up.

I mean, if you look at a-- a very crude test. If you think of all of the people you know and the kind of tasks that they're engaged in, actually you would quickly convince yourself that many of them are new relative to, say, 80 years ago. Many of the occupations that people who will graduate from MIT-- management, consulting, design, computer programming, many of the engineering tasks did not exist. And even for things that existed, like being a lawyer or a professor as an occupation, the content of the tasks that these occupations actually perform are very different.

But one issue that this sort of reasoning immediately reveals is that, over the last several decades, we do see new tasks, but not enough of them. But also, the new tasks are mostly for more college educated workers. We see many fewer of the blue collard new tasks, say, compared to the 50s or the 60s, and that raises the question of, even if new tasks come, can we then use them for creating good jobs for workers who don't have college, which is, of course, a very important part of shared prosperity.

AUDIENCE: Thank you.

AUDIENCE: Hi. My name is Callin. I'm from the Technology and Policy Program here at MIT. I really like-- I think it's very powerful to look at this over history and in the context of AI.

But I want to try and pull out maybe some of the differences with AI that spring out to me. And one is, with that plot of automation against wage growth and decline, I think something that struck me is that one of the things that's new about the AI discussion is the idea that a lot of the jobs that are being automated and are now the college-educated jobs, the ones that have historically been growing in wages over time. I was wondering if you had any thoughts on whether that's a fallacy, first of all, but whether you had any thoughts on how that changes the new tasks dimension, but also the worker power dimension of it.

Is there a dimension of this where, because it's a new group that is being subject to automation, the political economy of the whole situation and the power dynamics might shift considerably considering it's now the more privileged groups?

**DARON
ACEMOGLU:** Thank you very much. Those are an excellent set of issues. They're complex ones. Let me give a two layered answer. First of all, you're absolutely right. There is this claim that AI will be equalizing because it will replace or automate more higher paid jobs.

But AI or generative AI is not the first technology that promises to do that. In the 2000s, there was a lot of talk of digital technologies getting rid of middle managers and so on. And what you see when you look-- this requires much more detailed analysis-- when you look at the data with digital technologies, first of all, you see that the displacement has really concentrated on lower education tasks.

And the same is true of pre-generative AI, say, the AI that was spreading in the 2018, '19, '20. But, moreover, even when there is some displacement of higher education workers, that's not the end of the story because those workers can then go and take the tasks away from lower education groups. So in the labor markets, you actually see downward pressure on the low education wages, even when automation targets. Some of the more middle skill tasks.

So I think my assessment-- preliminary assessment-- is that the claims that generative AI is going to be an equalizing force are overblown. And even if it tends to automate higher skill tasks, that may not actually reduce inequality. That being said, in the book, Simon and I argue-- and Simon, David Autor, and I have a new sort of policy paper that provides more details on this-- there are many possibilities for using generative AI as a tool for reducing inequality, but it would go through, just like the previous question was emphasizing, through new tasks.

Creating, for example, finding ways of deploying generative AI in a way that's going to be useful for people who are more engaged in manual occupations-- electricians, line workers-- so that they can be trained faster, they can use that for problem solving or dealing with more real time issues that are more complex. For example, as you get the electrification of the grid, you're going to need much more complex tasks being performed by electricians. So there is a path for generative AI, but, decidedly, that's not the one we are on right now.

AUDIENCE: Hello. My question is actually not about AI, but more about the energy transition, where we'll also see a lot of transformation of jobs. So currently there are many countries around the world where the majority of their exports and the foreign currencies that they receive are tied to fossil fuel related products or commodities. And what I see-- based on what I have what I see is that these countries, most of them will not transform in time before we will need much less of these commodities. And my question is, what will happen to those countries and the workers that are inside those countries-- countries like, for example, Russia, Azerbaijan, and Venezuela-- that are very much related dependent on exports of fossil fuels?

**DARON
ACEMOGLU:** Yeah, those are very good questions. And I think that transition is not easy. In fact, that transition is not easy in the United States. The evidence is that when there were environmental regulations that were introduced in the United States that helped reduce pollution, for example, the Clean Air Act, the communities that were focused, for example, in coal production, they were negatively affected. And the workers did not easily relocate to other jobs either because there weren't the programs or because their skills were very specialized. So that's part of the unfortunate thing, and part of the reason why the energy transition is so hard.

I think the energy transition is going to have distributional effects, and some of those are going to be negative. And I think part of the reason why this has become such a politically charged issue in the United States is because the pro-climate groups are perhaps not articulating a way in which they coal miners are going to be dealt with, how you can do this without harming the coal miners or workers who are currently working in dirty industries. And I think that's going to be part of the challenge moving forward.

AUDIENCE: Thank you.

AUDIENCE: Thank you. My question is about how to empower people to look for better jobs or advocate for better positions. You mentioned, for example, unions, but I guess maybe the people already need to have some kind of basic safety so that they feel comfortable to work in the union, like health insurance if they lose their job or things like that. So do you think there's a way to get that in, like why--

**DARON
ACEMOGLU:** I think that's a great question, and unfortunately I don't know the answer. So it's very easy to agree with what Fotini said. The three very big movements-- Occupy Wall Street, MeToo, and Black Lives Matter-- did not coalesce into a bigger reform agenda. It's easy to bemoan that, but it's not easy to articulate how they could have done that in a much more-- I think that is going to be a very critical question, and I'm not sure that I have the answer.

But I think one experience from history is that the labor movement has been the linchpin of something like that because it is much more broad based. It is about bread and butter issues. So it creates a better grounding for civil society and other types of political movements. And, in some sense, perhaps the unfortunate thing of many of these movements that we're talking about, they have been completely divorced from the labor movement.

AUDIENCE: Hello. Hi, Professor. Thank you so much for the sharing your thoughts so far. I'm a final year PhD student here at MIT. And, just personally, I came here with the intention to use technology for good. And as I plow through the PhD, the world just looks gloomier, and, to name a few political tensions, COVID and energy crisis, et cetera. So going forward, I guess for my generation for the decades to come, in your view, what is the greatest technology fear that you have, and what is the one thing you're most optimistic about? Thank you.

**DARON
ACEMOGLU:** Well, that's a very hard question. I mean, I think your attitude is exactly right. I think there is an ethical responsibility. And the kind of questions that you've just posed-- that technologists need to get engaged with. I think for a long time-- it's a caricature-- but the attitude epitomized by Facebook's motto "move fast and break things" was, I think, defining for the technology sector.

We'll just disrupt things, and then the social consequences don't matter. Everything's going to work out. I think we need to move away from that. Once you move away from that, I think there are so many different things that can be done. But the central emphasis that I've tried to put here is that it's not just like technologies to cure cancer. It's not just technologies to detect pandemics in advance. Even when it comes to technologies that are the bread and butter of the production process, there are huge distributional consequences.

And I think the questions about, what are their social implications, that needs to be factored in. And, to me, triggered by all of these things, I think very transformative technologies have to include those that are going to find ways of making workers more productive, including workers with diverse skills, not just those with PhDs and masters, but workers who have high school degrees or vocational skills. How can we make workers with every level of skills more productive in the production process? I think that's going to be a critical part of socially beneficial technologies in the future.

AUDIENCE: Thank you, Professor Acemoglu, for an excellent presentation. My question will be actually about the inequality among the countries, rather than inside an economy. More about macroeconomic inequality. You gave the example of how the cotton gin actually created slave societies inside of the United States, but the cotton gin and the Industrial Revolution also created a huge wave of colonization around the globe. I would like to ask how this new wave of technology can create a new world order, and how can this new world order change the inequality among the countries, and what will be your suggestions for preventive measures to prevent any event that we have actually witnessed during and after the Industrial Revolution?

**DARON
ACEMOGLU:** Those are very, very good questions as well. Let me just answer it in the context of the issues that we've talked about. The fear about AI that I tried to briefly articulate in the one slide I put up on the developing world is that if AI goes down the path of more and more automation, it would actually be-- it could actually be something that expands inequality between countries.

Because if you look at the emerging world in the last 50, 60 years, almost all of the rapid growth experiences from South Korea, Taiwan, Malaysia, and later China, they all leveraged their labor force, their human resources, which is absolutely normal because that's the comparative advantage. Rich countries are very, very intensive in capital goods, they have some very educated workforces. But countries like South Korea, China, used to have decent workforces that are not paid very high wages. That creates a competitive advantage.

So if you push technology more and more towards automation, you're really precluding this type of growth behavior, which would ultimately slow down global convergence.

AUDIENCE: Thank you.

**EVAN
LIEBERMAN:** Just a few minutes left, so if everyone could be succinct in their question.

AUDIENCE: OK. My question is, briefly, why can't we approach this AI and optimization problem as a strictly redistribution problem, and why should we try to redirect technology at all to create new tasks that we might not need in the future? And I understand the climate technology analogy where we were able to redirect technology to an extent, but it might also be argued that, albeit at a horrible environmental cost, we would have much more energy at our disposal through fossil fuels if we didn't have to care about those environmental externalities. So can we use more aggressive redistribution to handle those externalities in the case of AI? And I know that you're critical of UBI in the book, but maybe--

**EVAN
LIEBERMAN:** We have a bunch of questions on the table--

**DARON
ACEMOGLU:** OK. Let me summarize the point, which is, what's wrong if I makes only 10 people around the world productive, and they become so productive and they produce everything else, and we redistribute the rest to the rest? I think the problem with that is, I think, it's both infeasible politically and actually undesirable, even if it were feasible. The reason why it's politically infeasible is because if you really have a very small group of people who do all the contributions, they're going to become politically very, very, very, very, very powerful.

You're already not going to convince people like Sam Altman and Elon Musk and Mark Zuckerberg to agree to redistribution if they become 50 times more productive and everybody else becomes more dispensable. The political economy of it would be worse. But even if it were feasible, I think that would be a very dystopian world in which social status would be highly unequal. There is a 90%, 95% of the people who are not real contributors, and just a group of geniuses who do everything that's worthwhile. I think that would be a very, very unequal world. Perhaps not in terms of consumption, but in terms of a lot of the other things that people care about.

AUDIENCE: Hello. Thank you very much for your presentation. I'd like to make a question about how to redirect technology to enhance democracy, considering that over the last years that we've seen that technology has increased polarization and made the public debate very problematic. So I'd like to listen your thoughts on how the technology could be used to make better decisions and more constructive debates instead of increasing polarization?

DARON ACEMOGLU: So that's a fantastic question, and the honest answer is, I don't know. But let me say two more sentences. First of all, I think when the internet and then later social media first came out, many people thought that these were going to be democratizing technologies. They turned out to be wrong, but they were not completely delusional.

When you look at the structure of social media and the internet, it creates pathways for people to interact on a level playing field without as much hierarchy, to form more deliberative bodies. And there are models, both in the United States and in other countries, that are actually attempting to do that. One that we discuss in the book is, for example, what Taiwan has done under the leadership of Audrey Tang that has actually led to tangible results.

Whether it can be scaled up, whether it can actually survive the other uses of technology, for example, for censorship and surveillance, whether it has a future once we have reached this level of polarization, those are all excellent questions. But I think technically it's feasible. And then the question is, can we socially go in that direction as well? Thank you.

AUDIENCE: Hi. My name is Nectarios, and I basically also wanted to ask you about the role of democratizing in AI, and to pick up one of Fotini's questions. Do you think we need kind of inequalities to foster progress?

DARON ACEMOGLU: Yeah, I mean, I think that's a great question that Fotini asked, and sorry I did not respond to that one. Yeah, I mean, I think that's very standard in economics that I think most economists would say, yes, it's going to be impossible and probably quite costly if you try to create a perfectly equal society. So that's why the focus that we have here has been not on eliminating inequality but keeping it to a reasonable level.

I think when you have CEOs earning 1,000 times as much as their line workers-- or even worse, the pattern that I showed, some people are really becoming much richer and other people are actually becoming poorer and are becoming unhealthy and are becoming more and more marginalized-- I think those are the kinds of alarming things. And I think those are the ones that are reversible.

No, I don't think the dream of creating a Communist society of everybody is equal-- that's neither feasible nor what we're talking about.

EVAN LIEBERMAN: We're just about out of time. I'm going to ask all of you to quickly ask your question--

DARON And I'll select one of them.

ACEMOGLU:

EVAN Any of them that you want, or maybe none. But we'll ask you to ask your question, then we'll get the last two questions.

AUDIENCE: Yeah, I'll keep it short. Sorry. We've been seeing a trend that I think people are finding college more undesirable now, it's getting more unpopular. I was wondering if that poses any challenges to task creation or coming up with new tasks and automation?

EVAN You ask your question.

LIEBERMAN:

AUDIENCE: Hi. My question is, with authoritarian governments utilizing new technologies to increasingly surveil their citizens as well as subversive powers within democratic governments increasingly doing so as well, what steps would you suggest that citizens in the public take to try and push back against this from happening, and do you think they'll be successful?

EVAN The last question.

LIEBERMAN:

AUDIENCE: Hi. I'm afraid of tech and polarization and the way it impacts societies, so do you think our polarized communities will hinder the progress of the future to come?

DARON I couldn't hear the last--

ACEMOGLU:

AUDIENCE: Do you think our polarized communities will hinder the progress of the future to come?

DARON That's a very interesting-- so let me just say one word on each one of them. I think that's a very interesting question. I think polarization has actually been much more micro, and inequality is much more micro in the United States with segregation. I think that raises a lot of interesting issues. We haven't really dwelled on them. It's something I've worked on in the past, but let me not say more on that.

ACEMOGLU:

On what is it that we can do in terms of pushing technology in a more democratic direction and resist sort of surveillance, again, I think that's a fantastic question. Some of the ideas like, for example, creating more ownership of data, collective ownership of data and better regulations for privacy protection, I think, are part of it.

But, also, I think if there is demand from civil society and for organizations and individuals for tools that are better at protecting their privacy against surveillance, against censorship, I think that could be a trigger for the development of these kinds of technologies. And then finally on college. I think that's an interesting-- there have been a couple of news articles on that, that college is becoming less central.

It's hard to know where we are. Actually, when you look at data over the last five years, still the status gap, income gap between college and non-college is as strong as before. So I'm not sure whether this sort of college is not so important is a reaction to a lot of people who went to low quality colleges and they're not getting the returns, or it's a passing fancy. We'll see.

EVAN
LIEBERMAN: Great. Well, I really want to thank Daron Acemoglu. That was terrific. And also to Fotini Christia for all of your comments, and all of you. Those were great questions. I'm delighted that MIT Center for International Studies could host this event.

Please join our email list, which is outside, if you want to receive updates on upcoming events. There's a sign-up sheet in the lobby. And for those of you who would like your book signed, I understand that Daron is willing to do so, so please make your way up to the lobby. Thank you again for joining us, and please show your applause.

[APPLAUSE]