Navigating the Future of Work

PERSPECTIVES ON AUTOMATION, AI, AND ECONOMIC PROSPERITY

Erik Brynjolfsson, Adam Thierer, and Daron Acemoglu

MARCH 2024

AMERICAN ENTERPRISE INSTITUTE

Navigating the Future of Work

PERSPECTIVES ON AUTOMATION, AI, AND ECONOMIC PROSPERITY

Erik Brynjolfsson, Adam Thierer, and Daron Acemoglu

Contents

Introduction

The Workforce Futures Initiative is a research collaboration among the American Enterprise Institute, the Brookings Institution, and the Project on Workforce at Harvard Kennedy School's Malcolm Wiener Center for Social Policy. The initiative aims to develop concise and actionable reviews of existing research for federal, state, and local policymakers. Since August 2021, the group has provided a forum for researchers and practitioners to discuss policy ideas, evaluate evidence, and identify priorities for new research on the future of work and the public workforce system.

In the first report, *Beyond the Turing Test: Harnessing AI to Create Widely Shared Prosperity*, Erik Brynjolfsson revises his view on AI, criticizing the Turing Test for equating human mimicry with intelligence and warning against economic consequences. He argues that true technological progress lies in augmenting—not replacing—human capabilities, historically increasing the value of labor. He criticizes the current trend of developing technology that substitutes for human labor, citing misaligned incentives among technologists, entrepreneurs, and policymakers. He advocates for innovation that complements human abilities, exemplified by companies like Cresta, which uses AI to assist, not replace, human operators. Brynjolfsson emphasizes the need for policy changes—such as equal taxation of capital and labor—to encourage such human-centered technology, arguing that the future of work depends on our choices regarding technology's role in the labor market.

In the second report, *We Can't Predict the Future of Work*, Adam Thierer explores the skepticism surrounding predictions about technology's impact on employment. Highlighting the tendency for overly pessimistic forecasts, he challenges the accuracy of such predictions with historical data. As examples of this overestimation, Thierer cites the recalibration of AI-related job loss estimates and the unexpected growth in certain job sectors. His report emphasizes the complexity of predicting future jobs and skills, advocating for flexible, adaptive workforce development rather than rigid government programs to navigate the evolving technological landscape.

In the final report, *Automation, AI, and Wages*, Daron Acemoglu examines the debate on automation and AI's impact on job creation and productivity. While some, such as *The Economist* and the McKinsey Global Institute, view AI as a driver of new jobs and growth, others express concerns about its potential to cause job loss and exacerbate inequality. Acemoglu argues that automation has not significantly increased productivity or jobs to offset losses. He highlights automation's limited success in creating good jobs and the growing inequality in labor markets, partly attributed to automation. He scrutinizes AI's role in the labor market, suggesting cautious adoption to avoid negative outcomes. His introduction also touches on the need for complementary investments and a balanced approach to leveraging automation and AI for society's benefit.

Beyond the Turing Test

HARNESSING AI TO CREATE WIDELY SHARED PROSPERITY

Erik Brynjolfsson

Alan Turing famously asked, "Can we create a machine that imitates humans so well that we can't tell which is which?" When I was a teenager, I remember thinking, "Oh, that's really good! If a machine is indistinguishable from a human to a group of testers, that must mean it's intelligent."

I have since completely changed my view. The Turing Test is a bad test of intelligence. It's about as reliable as assessing gravity's existence by asking if a magician can levitate someone to the astonishment of a live audience.

But more importantly, making machines that perfectly mimic humans would have some strikingly negative economic effects. First, if a machine closely imitates humans, then it's an economic substitute for labor, and that tends to drive down wages. In turn, that can create a trap—I call it "the Turing Trap"—in which many workers lose not only economic power but also the political power to reverse their predicament.[1]

Many think that, by definition, tech progress entails this sort of inexorable substitution of machines for humans. However, the historical reality is most tech progress has not substituted for humans but rather amplified and complemented our capabilities. One marker of this is that for over a century, an hour of human labor has generally increased in value (though not for everybody and not for all groups). For instance, manufacturing workers are paid about 10 times more for each hour of work today than they were paid in 1860.[2]

Why is an hour of labor more valuable now than it was in the past? Because today, we leverage our hands and brains with a lot of technology—hard technologies, such as bulldozers and computers, and soft technologies, such as business-process innovations. Technological progress that augments humans has increased wages.

Second, merely mimicking humans sets a ceiling on progress. If we are simply taking what's already being done and using a machine to replace what the human is doing, that puts an upper bound on how good you can get. For example, if a business automates the process of, say, making clay pots, then the clay pots can be made more cheaply and, as a result, you have a lot of inexpensive clay pots. However, the bigger value comes from creating an entirely *new* thing that never existed before, such as a supersonic jet, a nanoscale actuator, or a new way of solving protein folding to create medicines. We have iPhones because somebody invented something new. They didn't simply make a cheaper telegraph. Most of our increase in living standards comes from the invention of new goods and services, not from making the same things more cheaply.

The third important part of my Turing Trap argument is that three different groups—technologists, entrepreneurs and businesspeople, and policymakers—currently have misaligned incentives. Many technologists, though not all, focus on making machines that match humans in various tasks. It's an inspiring goal, passing the Turing Test. Some are working to make a robotic hand that's as dexterous as a human hand.[3] Others create technologies that play poker, chess, or other games that humans play.[4] Still others work on machines that can handle a telephone reservation or a medical consultation without human help.[5] These technologists are asking, "How can we replace humans doing existing tasks?" But in my view, they should more often ask, "What entirely new thing can we now do that we've never done before?" One reason they don't is the second question requires a lot more imagination.

I spend significant time with entrepreneurs and executives. I visit their organizations to watch them at work and I teach at a business school, where I study their decision-making. Once again, too often I see them focus on a task their business is already doing and think, "How can I replace the human worker with a machine?" as opposed to "How can we do something new?"

Finally, consider policymakers. The tax code, investment tax credits, and many other policy-guided decisions today heavily skew toward encouraging capital and discouraging labor. For instance, marginal tax rates on labor are currently much higher than tax rates on capital. Back in 1986, they were the same. But since then, they've changed in a way that discourages innovations that employ and reward labor and favors innovations that shift value to capital owners.

Therefore, for technologists, executives, and policymakers—and thus for our whole economy—innovation and investment do not create a level playing field. They skew toward creating technologies that substitute for humans rather than technologies that complement humans.

It doesn't have to be that way.

I work with several innovators and entrepreneurs who are doing something different. One company, Cresta, was started by Sebastian Thrun and Zayd Enam to help contact centers. But it's not a company that has a robot operator answer your call or a robot text generator respond to you. Instead, they keep humans not only in the loop but in charge. Customers talk to a human operator, and that person receives real-time tips by an artificial intelligence system. The system recommends topics that will be most useful to the caller, such as reminding the operator to mention a relevant product or a new price rebate or instructing them how to fix a particular problem. By augmenting humans this way, the operators have done fabulously. They can handle a much broader range of questions. There's higher customer satisfaction and higher throughput. Even the employees are less likely to quit.[6]

Using AI for augmentation turns out to be much more effective than trying to get the machine to handle the queries alone or having the humans work alone. The Cresta system combines the strengths of humans and machines. Lindsey Raymond, Danielle Li, and I have found the less-experienced workers benefit the most from this augmentation method, so it more equally distributes income as well.[7] This approach has been a win in terms of effectiveness, efficiency, and equity.

How can we encourage more companies to innovate toward complementing humans instead of substituting for them? One way is taxing capital and labor equally to create a more level playing field. A tax system that eliminates the existing incentives toward automation instead of augmentation would allow millions of managers and technologists to make their own local decisions without the government putting a thumb on the scale. Better yet, we have other tax systems, such as a value-added tax or X tax, that treat investment decisions much more evenly.

I'm not a technological determinist, and I don't think any particular outcome is inevitable in terms of how technology will affect work. The extent to which we augment human labor is a choice. We need to carefully consider what kind of world we want to live in. Do we want a world with widely shared prosperity? Do we want a world where everybody has some bargaining power? If we do, I believe we can create that. The mission of Stanford University's Digital Economy Lab is to do the research to understand what economic levers matter, what policies will make a difference, and how can we measure things more carefully so we can build a prosperous society.

We Can't Predict the Future of Work

Adam Thierer

In a 2002 speech on speculation, science fiction author Michael Crichton lambasted experts and the media for their "tendency to excess" and "crisisization of everything possible" when predicting the future.[8] Others have noted how sensationalism dominates forecasting because not only does bad news dominate media headlines[9] but "pessimism has always been big box office,"[10] with dystopian scenarios at the center of almost every story involving technology.[11]

Against this backdrop, pundits and politicians continue to make pessimistic predictions about the dangers of technology-induced unemployment. They do so even though the historical record tells a different—and quite positive—story about the relationship between innovation and jobs.[12] "Futurists don't know any more about the future than you or I," Crichton argued, and when reviewing their past predictions, "you'll see an endless parade of error" and a record that is "no better than chance."[13]

Indeed, a coin flip is typically a better predictor of future technology and employment trends. A 2012 report prepared for the Department of Defense evaluated over one thousand science and technology forecasts from academia, industry, government, and others.[14] The meta-survey revealed an average success rate of just 33 percent, with short-term forecasts (35 percent) faring only slightly better than long-term predictions (27 percent).[15]

Bad predictions are forgotten quickly, however, and replaced with other headline-grabbing pessimistic prognostications. Over the past decade, two major reports predicted massive job dislocations due to artificial intelligence. In 2013, Carl Benedikt Frey and Michael Osborne of the University of Oxford published a widely discussed study that surveyed hundreds of occupations and considered how likely they were to be automated.[16] They analyzed 702 professions and estimated 47 percent of US jobs were at high risk of being lost. Two years later, the McKinsey Global Institute published a report predicting as many as 45 percent of jobs (representing about $2 trillion in annual wages) "can be automated by adapting currently demonstrated technologies."[17] Seizing on these reports, headlines lamented, "Robots May Shatter the Global Economic Order Within a Decade."[18]

These reports were wildly off the mark. McKinsey recalibrated its model just two years later, admitting in 2017 that "very few occupations—less than 5 percent—are candidates for full automation."[19] Meanwhile, almost a decade after Frey and Osborne's study debuted, the US economy has added 16 million jobs. The profession they said would face the highest risk of technological disruption—insurance underwriters—instead has seen employment grow 16.4 percent since 2013.[20]

AI will cause job dislocations, of course, but no one can accurately predict which or how many jobs will be affected. Forecasting the future workforce is haunted by the same problem experts have always faced: We do not even possess a vocabulary to describe the jobs or skills of the future. When skimming old Bureau of Labor Statistics reports, such as the agency's mammoth 1969 *Tomorrow's Manpower Needs: National Manpower Projections and a Guide to Their Use as a Tool in Developing State and Area Manpower Projections*,[21] one finds no mention of any of the jobs that would eventually flow from the personal computing or internet revolutions. Even when old government reports or academic studies made passing mention of the future need for "computer skills," they offered no detail about what *specific* skills workers would require.

Employers, workers, and others instead had to master new skills and business models on the fly through constant iteration.[22] When mainframe computers dislocated an entire generation of human "calculators," who did hard math by hand for firms and government agencies, they got busy creating more and better computing devices. Once free to do more creative things, those calculators became the programmers who gave us the digital revolution. Some pundits now predict "the end of programming," with many of those workers losing their jobs to algorithms.[23] More likely, AI will once again free up workers to find still better things to do.[24]

A new book, *Working with AI: Real Stories of Human-Machine Collaboration*, provides almost 30 case studies showing how firms are currently integrating algorithmic technologies in the workplace and "practicing augmentation, not large-scale automation."[25] The common theme across these case studies is that "they involve highly complex collaboration,"[26] with humans and machines learning together through positive feedback loops.

Flexible workplace experimentation with new automation technologies will likely be the most crucial component of building a workforce that is better prepared for the future.[27] Unfortunately, that is also the hardest thing to devise in advance or facilitate through government programs, especially as firms are tapping an astonishing variety of new skill development models to adjust to AI-related automation.[28]

Many argue we still need to take steps to re-skill workers and prepare for the future, yet past government retraining efforts have fared poorly.[29] "Government job training programs (with the exception of apprenticeships) appear to be largely ineffective," concluded a 2019 report from the Trump administration's Council of Economic Advisers.[30] The biggest problem with these efforts and other proposals is they speak of a monolithic "workforce" when many workforces are constantly morphing. Meanwhile, pundits talk about "skills gaps" without even defining what they mean by it or bothering to prove such a gap exists.

The right mix of needed policies probably comes down to some combination of improved STEM education, better online learning and "micro-credentialing" programs (which are more focused than traditional four-year college degrees),[31] technical recertification efforts (especially more flexible retraining partnerships facilitated through community colleges), and creative vocational apprenticeship models. Many of these proposals have been floated before.[32] If policymakers want to prepare workers for the future by facilitating greater real-time learning by doing, then they must also act to remove barriers to flexible work and labor mobility,[33] including portable benefits solutions.[34]

Of course, difficult technological labor dislocations will still happen and will drive many policymakers to call for additional policies to help. Some transitional support mechanisms can help alleviate some of the pain associated with fast-moving technological change and keep the door open to ongoing innovation. A certain level of unemployment assistance will always be needed to cushion that blow. Beyond these steps, however, it remains hard to predict or plan for the uncertain future ahead.

Automation, AI, and Wages

Daron Acemoglu

Many are excited about the potential of automation and artificial intelligence to create new jobs and boost productivity. *The Economist*, for example, argues that "by lowering costs of production, [AI-based] automation can create more demand for goods and services, boosting jobs that are hard to automate. The economy may need fewer checkout attendants at supermarkets, but more massage therapists."[35]

McKinsey & Company's statement at the 2022 World Economic Forum was similarly optimistic: "Fourth Industrial Revolution technologies [are] driving productivity and growth across manufacturing and production at brownfield and greenfield sites. These technologies are creating more and different jobs that are transforming manufacturing and helping to build fulfilling, rewarding, and sustainable careers."[36]

When this optimism is challenged, many accuse their challengers of being Luddites. This is not entirely fair. Certainly, emergent technologies such as AI hold great potential, but that does not mean they won't be used for nefarious purposes or they will not create costly disruptions, especially if they are used for automating tasks previously performed by humans.

Thus far, automation has not shown sufficient productivity or job increases to offset the consequent job losses. Speaking to *The Economist*'s argument, the US Bureau of Labor Statistics Occupational Outlook Handbook estimates we will lose more than 10 times the number of cashiers as we will gain in massage therapists over the next decade.[37] Further, it is not obvious that displaced cashiers would necessarily make good massage therapists, at least not without additional training, credentialing, and so on. I suspect these work-eliminating applications of automation will continue to produce disruptive results, and the productivity gains are likely to be less than optimists expect, because we haven't made the adjustments and counterbalancing investments necessary for creating jobs—especially good jobs—for a broad range of skills.

If you examine US data, there is already plenty to worry about. Much of the data predates AI but not automation. Expanding inequality in US labor markets is a relatively recent phenomenon. As Figure 1 shows, before 1980, wages for all workers grew at similar rates, with low-education workers even seeing faster wage growth than high-education workers in some periods. However, after 1980, wages for high-education workers began to grow much faster than did wages for low-education workers. This trend is particularly evident for men. Even more worrying than the growing inequality these trends show is that, since 1980, the real earnings of low-education groups have stagnated or even declined.

This problem is not unique to the US; it is a global problem. Such a trend, I argue, is caused by automation.

Before I go further, I want to step back and decompose the argument about AI that goes like this: "Technology might cause some disruption, but in the end, it's going to bring huge productivity benefits." To do this, I will build on my work with Pascual Restrepo, in which we develop a task-based model to study the implications of different types of technologies, including automation, for productivity and wages.

In this framework, we suppose production requires completing a range of tasks. For example, to produce a car, the tasks involved include design, engineering, assembly, and so on. Firms need to decide how to perform these tasks most efficiently.

To simplify things, I focus on the relationship between capital and labor. In Figure 2, the y-axis is the cost of production, and the x-axis is the range of tasks. The orange line shows the cost of producing tasks

Figure 1. Evolution of Real Weekly Earnings Across Demographic Groups, 1963–2017

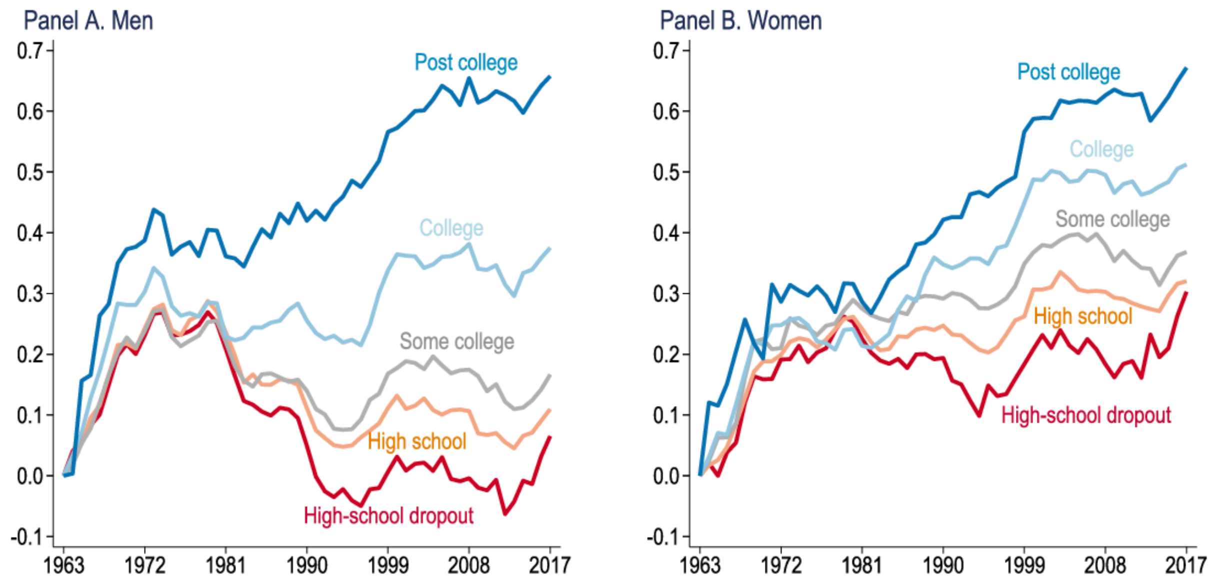

Source: David H. Autor, "Work of the Past, Work of the Future," *AEA Papers and Proceedings* 109 (May 2019): 1–32, https://www.aeaweb.org/articles?id=10.1257/pandp.20191110.

Figure 2. Allocation of Tasks to Factors

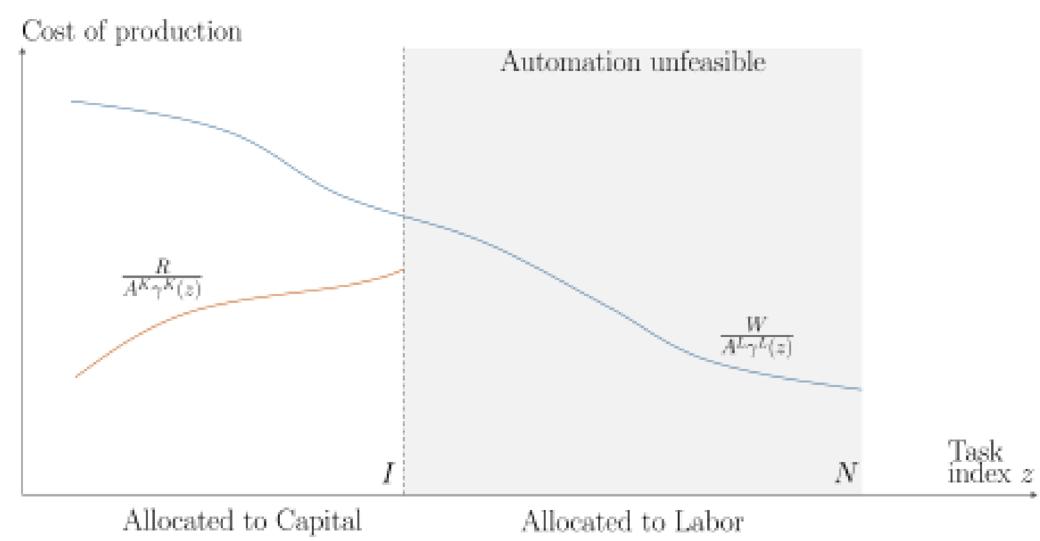

Source: Author.

using capital, and the blue line shows the cost of producing tasks using labor. Cost minimization means choosing the lowest envelope of these two curves. The "I" on the graph indicates we can't automate all tasks using capital because we lack the technological know-how to do so.

Now, let's try to understand the optimism economists tend to have about technology. The easiest way

Figure 3. Labor-Augmenting Technological Change

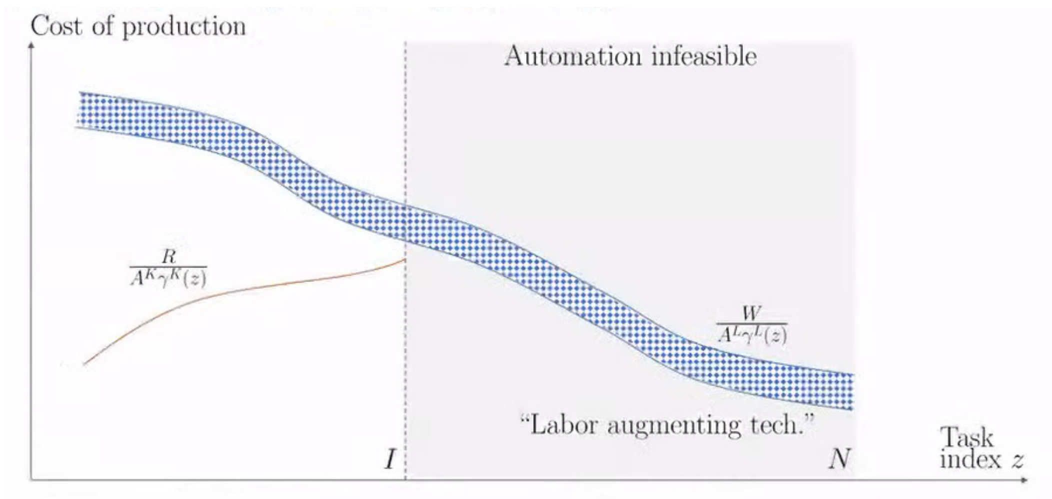

Source: Author.

to think about technology is that it makes the tools of production more productive. For example, technology can make labor more productive in all the tasks it performs, as shown in Figure 3.

Technology would shift the blue cost curve down, as shown in Figure 3. If this happens, we would get all the blue area as productivity benefits. This would be great because it would mean we as workers were becoming more productive, which would lead to higher wages, rising consumer welfare, and stronger economic growth. Also centrally, as labor becomes more productive, the allocation of tasks between labor and capital remains almost the same. Hence, labor is not displaced from the tasks it used to perform because of technological change.

If, on the other hand, we make capital more productive, the orange cost curve would shift down, and we would get all the orange area as productivity benefits, shown in Figure 4. There is, once again, no or little displacement of labor. So, if the world looks like this, then *The Economist* and McKinsey argument I referenced earlier—"short-term disruption, long-term production"—would be correct. However, this is not what automation is about. Automation involves the substitution of capital or machinery for tasks previously performed by labor. This would shift the blue cost curve from I to I′, as shown in Figure 5.

As you can see, this would lead to a much smaller productivity gain because the productivity benefits come from only the automated tasks—and only to the extent that capital is better at producing these tasks than labor is. The upshot is that while automation will lead to some productivity gains, these are unlikely to be as large as the gains that would come from making labor or capital more productive in all tasks.

Secondly, automation creates a huge displacement effect. There's no such thing as a free lunch: Workers who used to perform the tasks of making a car will be replaced. Now, that is bad for the transitioning workers, but, more generally, these workers will need to be reallocated to other jobs, which could actually push down wages. (They will certainly push down labor shares.) In subsequent work, Restrepo and I document these displacement and inequality effects, explained in Figure 6.

In Figure 6, we consider how automation has affected different types of labor. We looked at 500 different demographic groups—distinguished by gender, education, experience, and ethnicity—and estimated

Figure 4. Productivity Gains from Capital-Augmenting Technological Change

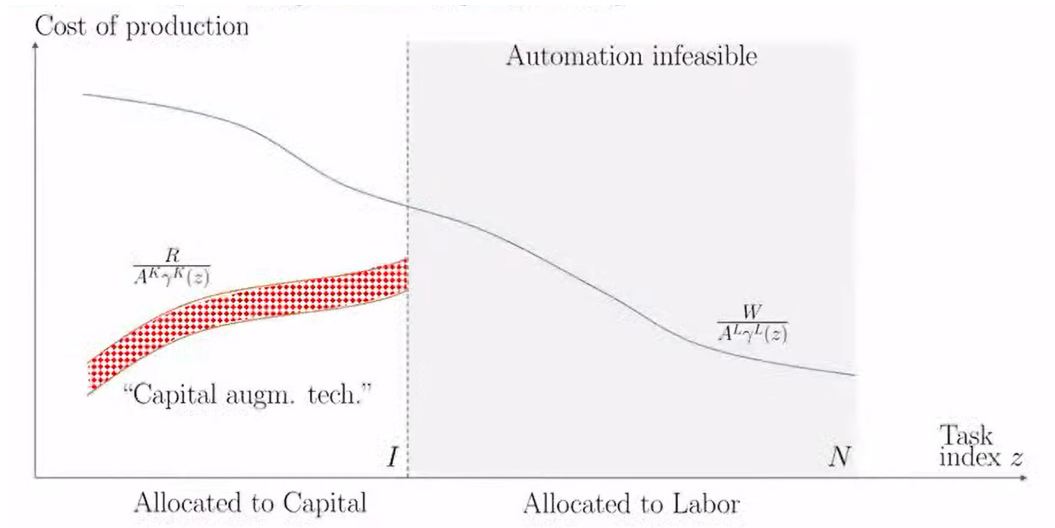

Source: Author.

Figure 5. Productivity Gains in Displacement from Automation

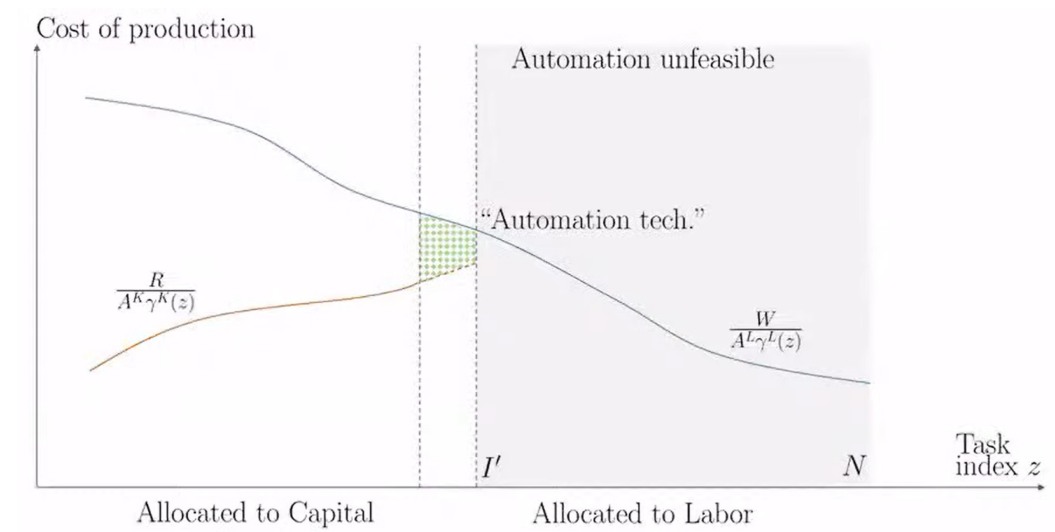

Source: Author.

what fraction of tasks that each group used to perform in 1980 has since been displaced by automation. We found the more a demographic group's work has been displaced by automation, the less its real wages have grown since 1980. In fact, automation explains about 70 percent of the between-group inequality in the US since 1980.

Now, these trends are mostly about automation before AI. What, then, do we know about AI's role in these automation trends? Talk of AI goes back to the late 2000s, but it wasn't until around 2016 that postings for AI-related jobs significantly increased, according to work I've done with David H. Autor, Restrepo, and Jonathon Hazell (Figure 7).[38]

Figure 6. Changes in Hourly Wages Across Demographics (1950–1980 and 1980–2016) Against Task Displacement (1980–2016)

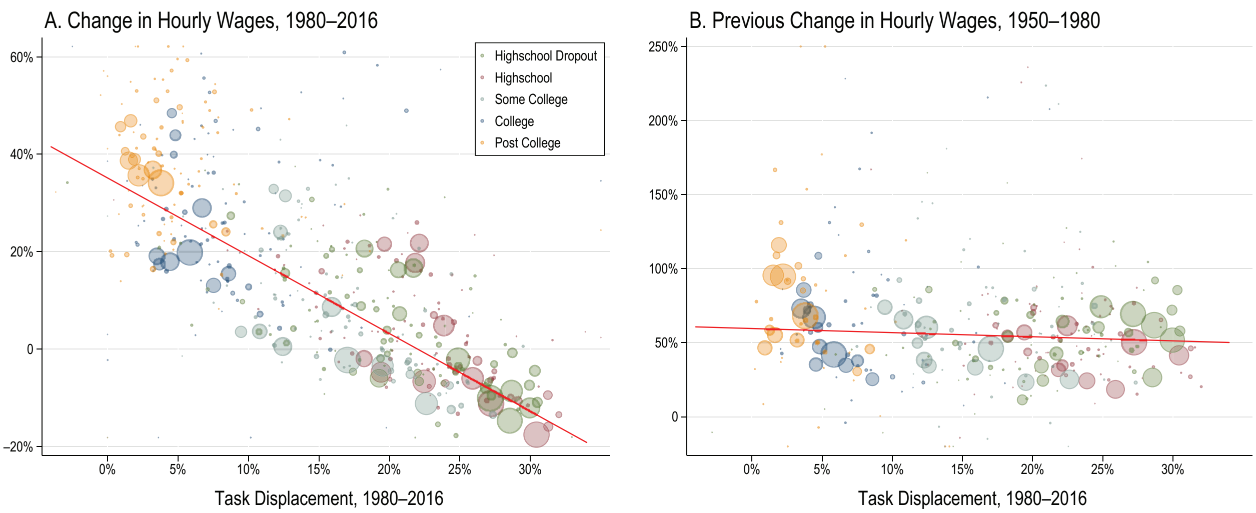

Source: Daron Acemoglu and Pascual Restrepo, "Tasks, Automation, and the Rise in US Wage Inequality" (working paper, National Bureau of Economic Research Working Paper, Cambridge, MA, June 2021), https://www.nber.org/papers/w28920.

Figure 7. Share of AI-Related Vacancies in Burning Glass and by Broad Industry

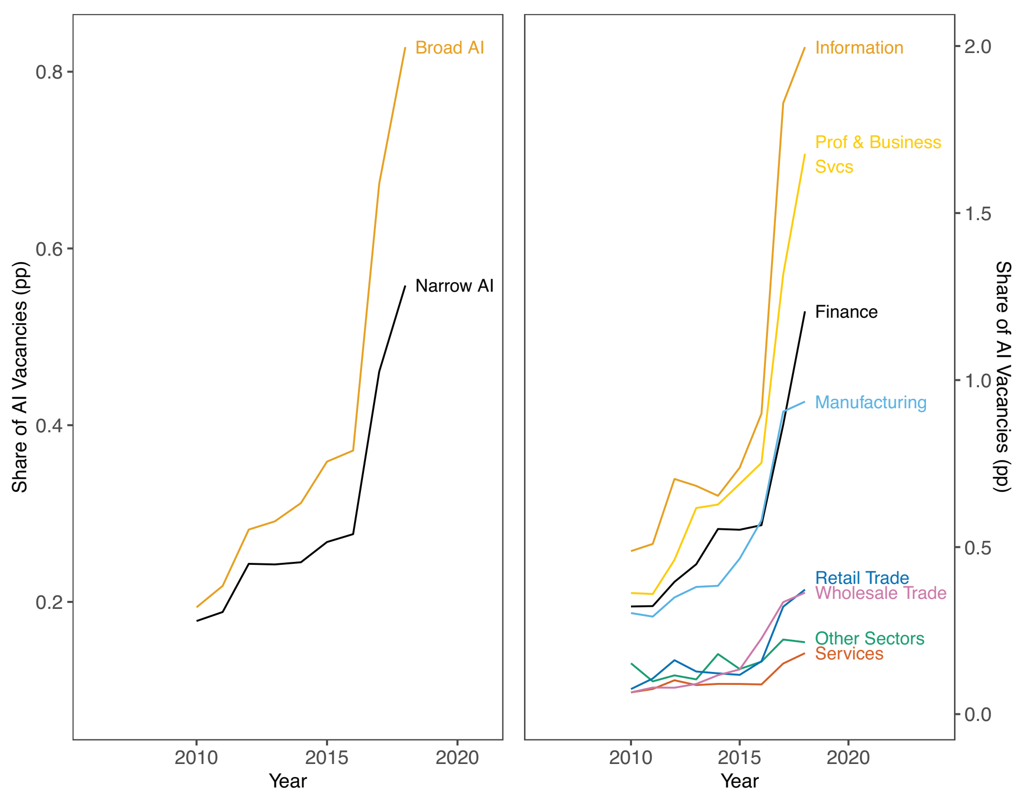

Source: Daron Acemoglu, "Automation, AI and Wages" (PowerPoint presentation, Workforce Futures Initiative, October 2022), https://www.aei.org/wp-content/uploads/2023/02/Daron-Acemoglu-Automation.pdf.

Figure 8. AI Vacancies by Level of AI Exposure

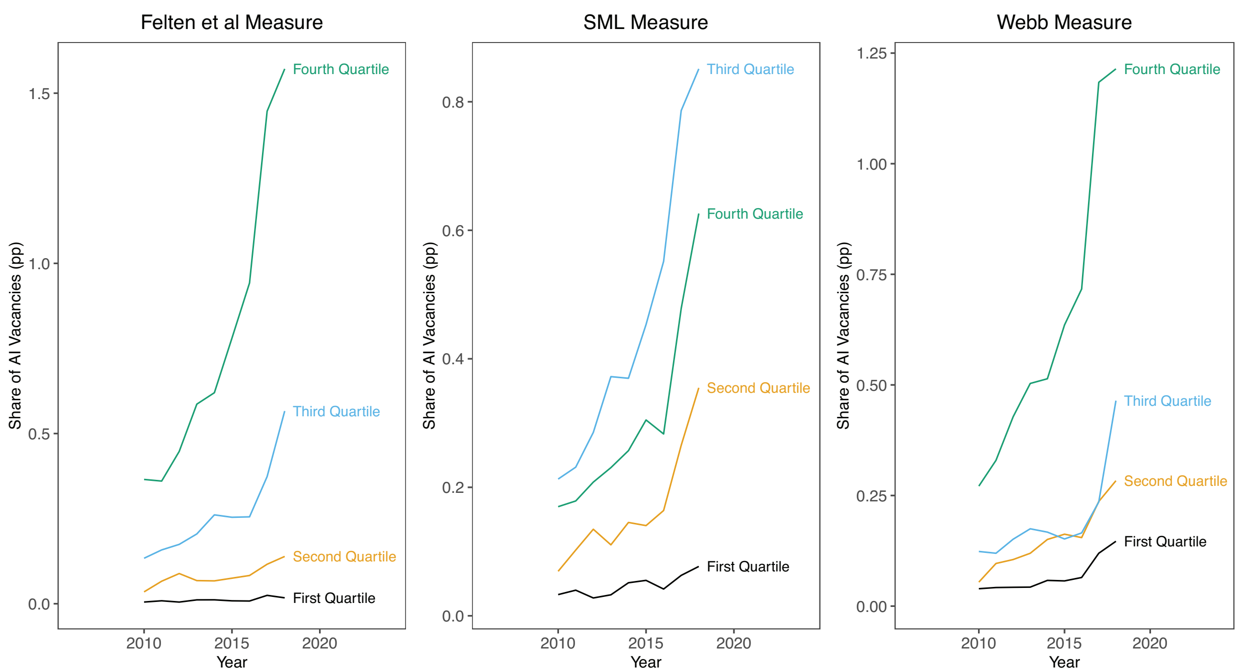

Source: Daron Acemoglu, "Automation, AI and Wages" (PowerPoint presentation, Workforce Futures Initiative, October 2022), https://www.aei.org/wp-content/uploads/2023/02/Daron-Acemoglu-Automation.pdf.

This suggests AI is starting to significantly influence the US labor market. But how will we use AI? This is the right time to ask these questions, and reevaluate the implications of task automation, because AI, as a general-purpose technology, can have many different uses. One important direction is what early computer pioneers, such as Norbert Wiener, Douglas Engelbart, and J. C. R. Licklider, envisioned—to augment human capabilities and productivity. Or AI can be used for automation, continuing what digital technologies in offices and robots on factory floors have done since 1980. There are important choices about how to use AI, but a first step is to recognize how it has been used so far.

We get a first sense of this by using AI-exposure indexes compiled by several teams of researchers—in particular, Edward W. Felten, Manav Raj, and Robert Seamans;[39] Erik Brynjolfsson, Tom Mitchell, and Daniel Rock;[40] and Michael Webb.[41] Autor, Restrepo, Hazell, and I used these data to measure which types of labor have more tasks that can use AI and thus are more likely to adopt AI early on. This exercise shows that, as expected, establishments with more tasks that can be performed by AI are at the forefront of AI adoption, as shown in Figure 8.

However, as shown in Figure 9, we also find the same establishments slow their hiring of non-AI workers (and much more so than they increase the number of AI specialists they hire).[42] So overall, the early rollout of AI seems to have taken the same automation path other digital technologies of the 1980s, 1990s, and 2000s followed.

This overall decline in hiring therefore confirms AI is replacing human workers in some occupations, even after accounting for AI-skilled job creation. There are many possible explanations for this. For instance, businesses may be prioritizing short-term cost cutting over long-term job growth. Another possible factor is that the US tax code heavily subsidizes capital, which ultimately makes it cheaper for businesses to invest in automation than to hire workers with relatively higher payroll taxes. Today, there is

Figure 9. Non-AI-Related Vacancies by Level of AI Exposure

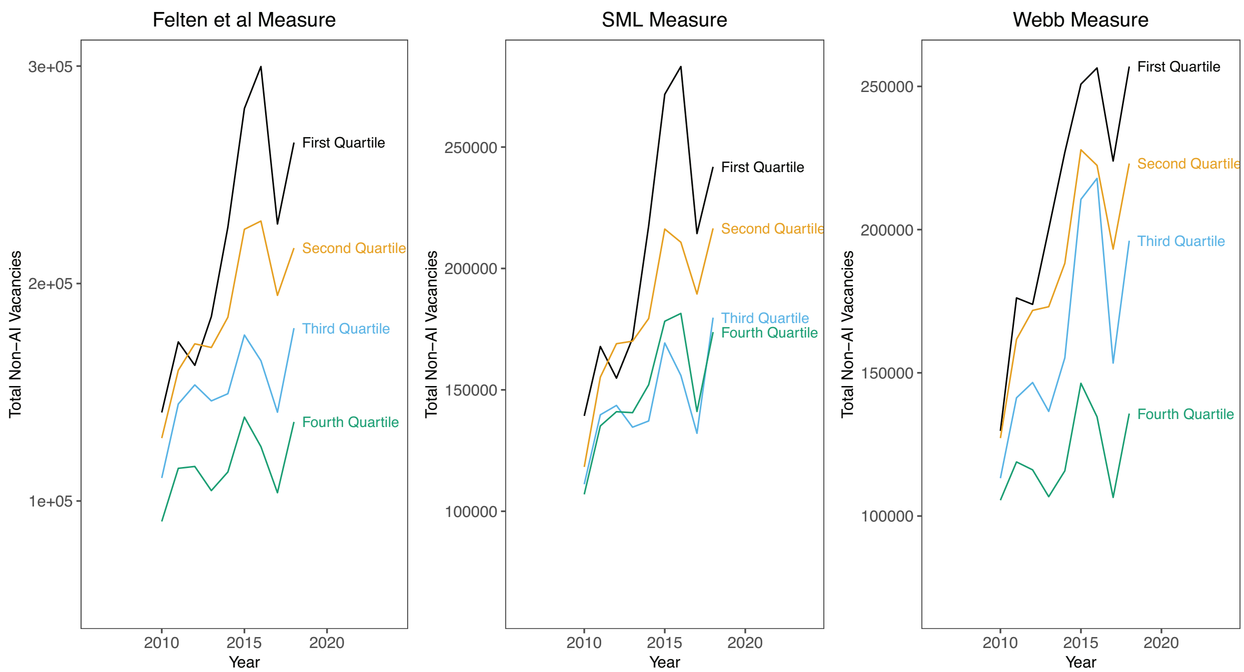

Source: Daron Acemoglu, "Automation, AI and Wages" (PowerPoint presentation, Workforce Futures Initiative, October 2022), https://www.aei.org/wp-content/uploads/2023/02/Daron-Acemoglu-Automation.pdf.

a 20 percent gap between the marginal tax rate for labor and capital, as shown in Figure 10.[43]

Note that this is only preliminary evidence, and more research is needed to understand AI's full impact on the labor market. However, the findings so far suggest we need to be careful about how we adopt AI technology. We need to ensure we do not automate jobs that are essential to our economy and society.

There is good automation and bad (or "so-so") automation. Good automation increases productivity and creates new jobs. Bad automation reduces employment growth, fails to deliver worthwhile benefits to productivity, and has negative distributional consequences. Over the past 40 years, we have seen more bad automation than good. I predict AI will continue this trend, although there is no technological necessity for it to do so.

The optimistic outlook holds that we shouldn't worry about AI because technology has a good track record of creating jobs. But in the past, automation has sometimes led to job losses and decreased wages. For example, in Great Britain, the first wave of automation during the Industrial Revolution led to about 90 years of real wage declines, increased working hours, and poor working conditions. The lives of workers improved only when Great Britain became a democracy, resulting in the legalization of trade unions and the abolition of child labor.[44] In the US, the mechanization of agriculture wasn't such a smooth process either.

Importantly, technology's benefits have not come automatically. Rather, they resulted from social and political struggles. I am worried we're moving to an increasingly unequal future because things are already unequal at present. We need to advance our efforts to prevent that, be careful about how we adopt automation technology, and use this technology in a way that benefits everyone, not just the wealthy. Here are some suggestions for how to move forward.

Figure 10. Effective Taxes on Labor and Different Types of Capital, 1981–2018

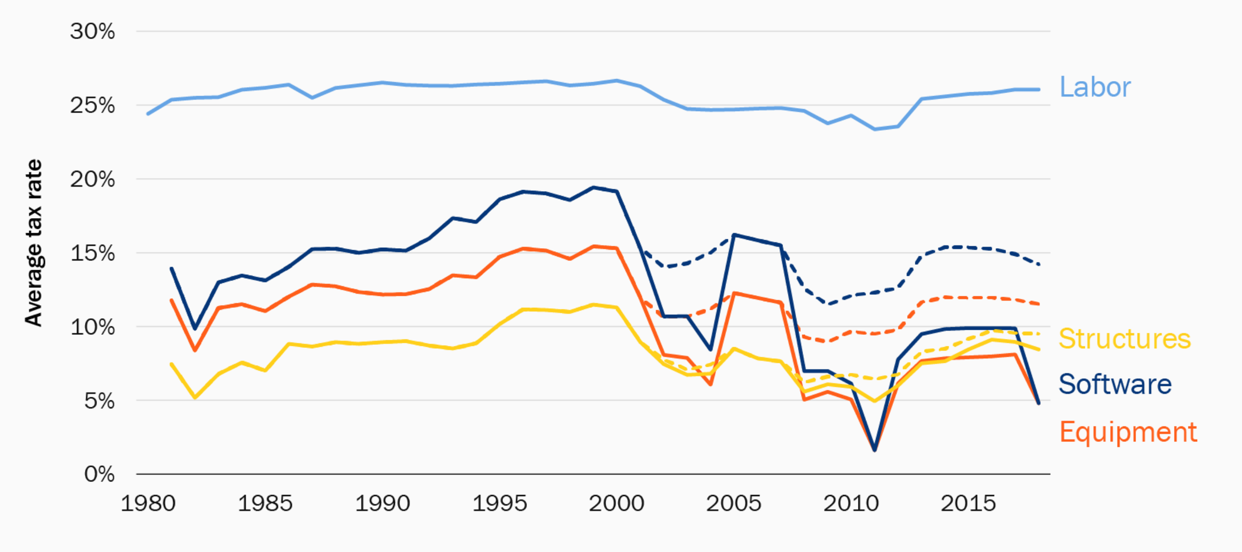

Note: Solid lines are observed effective taxes. Dashed lines are effective taxes if treatment of allowances had remained as they were in 2000.
Source: Daron Acemoglu, Andrea Manera, and Pascual Restrepo, "Does the US Tax Code Favor Automation?," *Brookings Papers on Economic Activity* (Spring 2020): 231–85, https://www.brookings.edu/articles/does-the-u-s-tax-code-favor-automation.

- Complementary investments play a crucial role and may be the most significant factor in achieving desired outcomes.
- While automation is an important focus in today's business world, driven by cost cutting and large-scale growth opportunities, it is not the sole path to success.
- Diversifying our technological investments can bring key advantages. During the mechanization of agriculture, the US economy continued to grow and the labor market did not collapse, because there was a tremendous amount of investment in manufacturing and clerical occupations, such as design and technical jobs. In fact, labor share increased during that period, in contradistinction to today's steep decline in the labor share.

Ours is also an institutional problem. Labor unions have a complex relationship with automation. In the United States, a labor union's presence impels employers to automate because it creates a conflictual relationship. In Germany, work councils and labor unions actually *encourage* firms to automate while investing in training and upgrading skills so blue-collar workers can become technical employees.[44] We need a much stronger worker voice. We also need a more competitive business environment so a handful of large companies can't create a monopoly of (and through) automation.

Lastly, and as previously mentioned, the tax code is an important part of the story. But so is a redirection of effort and focus in the AI research community. Without coordination, Brynjolfsson and I (in my case, jointly with Michael Jordan and Glen Weyl) wrote papers around the same time on "The Turing Test Gone Wrong."[46] I believe we both, from different perspectives, arrived at the same conclusion: The focus on human parity and human replacement—which can be traced back to Alan Turing's vision of technology as machines that indistinguishably imitate human behavior—is leading us astray. By drawing attention to ethical considerations in the research community, we can reshape this vision. Such a vision is vital if we want technological advancements that benefit individuals and society as a whole.

About the Authors

Daron Acemoglu is an economist and the Institute Professor at the Massachusetts Institute of Technology. A leading figure in the field, Dr. Acemoglu is known for his research on political economy, economic development, and labor economics, and he has coauthored several notable works such as the 2012 book *Why Nations Fail*, which focuses on the role of institutions in shaping nations' economic outcomes.

Erik Brynjolfsson is the Jerry Yang and Akiko Yamazaki Professor at the Stanford Institute for Human-Centered Artificial Intelligence, director of the Stanford Digital Economy Lab, and cofounder of Workhelix. His research and teaching focus on the economic implications of AI and digitization.

Adam Thierer is a senior fellow in the Technology and Innovation policy program at the R Street Institute. This report is derived from a forthcoming R Street policy study, "Can We Predict the Jobs & Skills Needed for the AI Era?"

Notes

1. Erik Brynjolfsson, "The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence," *Daedalus* (Spring 2022), https://digitaleconomy.stanford.edu/news/the-turing-trap-the-promise-peril-of-human-like-artificial-intelligence.

2. Specifically, in 1960, average manufacturing wages were $0.75 to $1.50 per day or about $0.10 per hour. (See US Census Bureau, "Chapter D: Labor," in *Bicentennial Edition: Historical Statistics of the United States, Colonial Times to 1970, Part 1*, 148–49, https://www.census.gov/library/publications/1975/compendia/hist_stats_colonial-1970.html.) By 2021, the average hourly manufacturing wage was $29.52, or over 290 times higher in nominal terms. However, the price level became roughly 27 times higher over this period. (See US Bureau of Labor Statistics, "Consumer Price Index for All Urban Consumers: All Items in U.S. City Average," FRED Economic Data, Federal Reserve Bank of St. Louis, https://fred.stlouisfed.org/series/CPIAUCSL.) Combining these, the net increase in hourly wages was a bit more than 10-fold between 1860 and 2021.

3. Shadow Robot, "Dexterous Hand Series," https://www.shadowrobot.com/dexterous-hand-series.

4. Keith Romer, "How A.I. Conquered Poker," *New York Times Magazine*, updated June 15, 2023, https://www.nytimes.com/2022/01/18/magazine/ai-technology-poker.html; and Chess.com, "Deep Blue (Chess Computer)," https://www.chess.com/terms/deep-blue-chess-computer.

5. James Vincent, "Google's AI Reservation Service Duplex Is Now Available in 49 States," The Verge, April 1, 2021, https://www.theverge.com/2021/4/1/22361729/google-duplex-ai-reservation-availability-49-us-states.

6. See Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond, "Generative AI at Work" (working paper, National Bureau of Economic Research, Cambridge, MA, April 2023).

7. See Brynjolfsson, Li, and Raymond, "Generative AI at Work"; and Andrew McAfee, Daniel Rock, and Erik Brynjolfsson, "How to Capitalize on Generative AI", *Harvard Business Review*, November–December 2023, https://hbr.org/2023/11/how-to-capitalize-on-generative-ai.

8. Michael Crichton, "Why Speculate?" (speech, International Leadership Forum, La Jolla, CA, April 26, 2002), http://web.archive.org/web/20070714204136/http://www.michaelcrichton.net/speech-whyspeculate.html.

9. Jeff Jarvis, "Optimism Doesn't Sell," Fear and Technopanic, May 25, 2015, https://medium.com/change-objects/optimism-doesn-t-sell-1f3723cff6d7.

10. Matt Ridley, *The Rational Optimist: How Prosperity Evolves* (New York: Harper Collins, 2010), 294.

11. Adam Thierer, "How Science Fiction Dystopianism Shapes the Debate over AI & Robotics," Discourse, July 26, 2022, https://www.discoursemagazine.com/culture-and-society/2022/07/26/how-science-fiction-dystopianism-shapes-the-debate-over-ai-robotics.

12. David H. Autor, "Why Are There Still So Many Jobs? The History and Future of Workplace Automation," *Journal of Economic Perspectives* 29, no. 3 (Summer 2015): 3–30, https://www.aeaweb.org/articles?id=10.1257/jep.29.3.3.

13. Crichton, "Why Speculate?"

14. Carie Mullins, "Retrospective Analysis of Technology Forecasting: In-Scope Extension," Tauri Group, August 13, 2012, https://apps.dtic.mil/sti/citations/ADA568107.

15. Mullins, "Retrospective Analysis of Technology Forecasting."

16. Carl Benedikt Frey and Michael Osborne, "The Future of Employment: How Susceptible Are Jobs to Computerisation?" (working paper, Oxford Martin Programme on Technology and Employment, Oxford, UK, September 17, 2013), https://www.oxfordmartin.ox.ac.uk/publications/the-future-of-employment.

17. Michael Chui, James Manyika, and Mehdi Miremadi, "Four Fundamentals of Workplace Automation," McKinsey Digital, November 1, 2015, https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/four-fundamentals-of-workplace-automation.

18. Ambrose Evans Pritchard, "Robots May Shatter the Global Economic Order Within a Decade," *Telegraph*, November 5, 2015, https://www.telegraph.co.uk/finance/economics/11978542/Robots-may-shatter-the-global-economic-order-within-a-decade.html.

19. James Manyika et al., "Harnessing Automation for a Future That Works," McKinsey Global Institute, January 12, 2017, https://

www.mckinsey.com/featured-insights/digital-disruption/harnessing-automation-for-a-future-that-works.

20. Robert Atkinson, "Nine Years On, Predicted AI Job Loss Hasn't Happened," IndustryWeek, October 3, 2022, https://www.industryweek.com/talent/labor-employment-policy/article/21252036/nine-years-on-predicted-ai-job-loss-hasnt-happened.

21. US Department of Labor, Bureau of Labor Statistics, *Tomorrow's Manpower Needs: National Manpower Projections and a Guide to Their Use as a Tool in Developing State and Area Manpower Projections*, vol. 1, February 1969, https://fraser.stlouisfed.org/title/tomorrow-s-manpower-needs-national-manpower-projections-a-guide-use-a-tool-developing-state-area-manpower-projections-4863/tomorrow-s-manpower-needs-volume-i-498427.

22. James Bessen, *Learning by Doing: The Real Connection Between Innovation, Wages, and Wealth* (New Haven, CT: Yale University Press, 2015).

23. Matt Welsh, "The End of Programming," *Communications of the ACM* 66, no. 1 (January 2023): 34–35, https://cacm.acm.org/magazines/2023/1/267976-the-end-of-programming/fulltext.

24. "We must remember that artificial intelligence will be training humans much better than humans are trained now." Adam Ozimek, "The Paradox of Robots Taking All Our Jobs," *Forbes*, October 28, 2017, https://www.forbes.com/sites/modeledbehavior/2017/10/28/the-paradox-of-robots-taking-all-our-jobs.

25. Thomas H. Davenport and Steven M. Miller, *Working with AI: Real Stories of Human-Machine Collaboration* (Cambridge, MA: MIT Press, 2022), xiv.

26. Davenport and Miller, *Working with AI*, 197.

27. Katherine C. Kellogg and Melissa A. Valentine, "Five Mistakes Managers Make When Introducing AI—and How to Fix Them," *Wall Street Journal*, November 5, 2022, https://www.wsj.com/articles/managers-mistakes-artificial-intelligence-11667596940.

28. Oliver Hsu, "Building the American Workforce," Andreessen Horowitz, November 3, 2022, https://a16z.com/2022/11/03/building-the-american-workforce.

29. Lola Fadulu, "Why Is the U.S. So Bad at Worker Retraining?" *The Atlantic*, January 4, 2018, https://www.theatlantic.com/education/archive/2018/01/why-is-the-us-so-bad-at-protecting-workers-from-automation/549185; and David Muhlhausen, "So Far, Federal Job-Training Programs Have Been Outright Failures," *The Hill*, March 14, 2017, https://thehill.com/blogs/pundits-blog/economy-budget/323885-thus-far-federal-job-training-programs-have-been-an.

30. White House, Council of Economic Advisers, *Government Employment and Training Programs: Assessing the Evidence on Their Performance*, June 2019, https://trumpwhitehouse.archives.gov/wp-content/uploads/2019/06/Government-Employment-and-Training-Programs.pdf.

31. Scott L. Wyatt and Allen C. Guelzo, "College Doesn't Need to Take Four Years," *Wall Street Journal*, February 2, 2023, https://www.wsj.com/articles/college-doesnt-need-to-take-four-years-higher-education-credentials-university-degree-training-credits-students-11675370630.

32. US Department of Labor, Commission on Workforce Quality and Labor Market Efficiency, *Investing in People: A Strategy to Address America's Workforce Crisis*, September 1989, https://eric.ed.gov/?id=ED317664.

33. John Lettieri and Kenan Fikri, *The Case for Economic Dynamism and Why It Matters for the American Worker*, Economic Innovation Group, 2022, 6, https://eig.org/dynamism.

34. Eli Lehrer, "The Future of Work," *National Affairs*, no. 28 (Summer 2016): 35–52, https://www.nationalaffairs.com/publications/detail/the-future-of-work.

35. *The Economist*, "Robots Threaten Jobs Less Than Fearmongers Claim," April 10, 2021, https://www.economist.com/special-report/2021/04/08/robots-threaten-jobs-less-than-fearmongers-claim.

36. Daphne Luchtenberg, "The Fourth Industrial Revolution Will Be People Powered," McKinsey & Company, January 7, 2022, https://www.mckinsey.com/capabilities/operations/our-insights/the-fourth-industrial-revolution-will-be-people-powered.

37. US Bureau of Labor Statistics, "Occupational Outlook Handbook," 2023, https://www.bls.gov/ooh.

38. Daron Acemoglu et al., "Artificial Intelligence and Jobs: Evidence from Online Vacancies," *Journal of Labor Economics* 40, no. S1 (2022), https://www.journals.uchicago.edu/doi/abs/10.1086/718327.

39. Edward W. Felten, Manav Raj, and Robert Seamans, "A Method to Link Advances in Artificial Intelligence to Occupational Abilities," *AEA Papers and Proceedings* 108 (May 2018): 54–57, https://www.aeaweb.org/articles?id=10.1257/pandp.20181021; and Edward Felten, Manav Raj, and Robert Channing Seamans, "The Effect of Artificial Intelligence on Human Labor: An Ability-Based Approach," *Academy of Management Proceedings* 1 (August 2019), https://journals.aom.org/doi/10.5465/AMBPP.2019.140.

40. Erik Brynjolfsson, Tom Mitchell, and Daniel Rock, "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?" *AEA Papers and Proceedings* 108 (May 2018): 43–47; and Erik Brynjolfsson, Tom Mitchell, and Daniel Rock, "Machine Learning and Occupational Change" (unpublished manuscript, 2019).

41. Michael Webb, "The Impact of Artificial Intelligence on the Labor Market," (unpublished manuscript, January 2020), https://www.michaelwebb.co/webb_ai.pdf.

42. Acemoglu et al., "Artificial Intelligence and Jobs."

43. Daron Acemoglu, Andrea Manera, and Pascual Restrepo, "Does the US Tax Code Favor Automation?" *Brookings Papers on Economic Activity* (Spring 2020): 231–85, https://www.brookings.edu/articles/does-the-u-s-tax-code-favor-automation.

44. Daron Acemoglu and Simon Johnson, *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity* (New York: PublicAffairs, 2023).

45. Acemoglu and Johnson, *Power and Progress*.

46. Erik Brynjolfsson, "The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence," Stanford Digital Economy Lab, January 12, 2022, https://digitaleconomy.stanford.edu/news/the-turing-trap-the-promise-peril-of-human-like-artificial-intelligence; and Daron Acemoglu, Michael I. Jordan, and E. Glen Weyl, "The Turing Test Is Bad for Business," *Wired*, November 8, 2021, https://www.wired.com/story/artificial-intelligence-turing-test-economics-business.