REDESIGNING AI

WORK, DEMOCRACY, AND JUSTICE IN THE AGE OF AUTOMATION

in collaboration with
The Partnership on AI

made possible by a generous grant from
The William and Flora Hewlett Foundation

Boston Review
PO Box 390568
Cambridge, MA 02139

ISSN: 0734-2306 / ISBN: 978-1-946511-62-1

d_ro

CONTENTS

EDITORS' NOTE

FORUM

- REDESIGNING AI

FORUM RESPONSES

- A WORLD WITH LESS WORK
- CENTERING WORKERS' POWER AND RIGHTS
- THE PANDEMIC BOLSTERED SUPPORT FOR NECESSARY REFORMS
- TECHNOLOGY-FOCUSED SOLUTIONS WON'T WORK
- DECOLONIZING AI
- THE PROBLEM IS WAGES, NOT JOBS
- BEYOND THE AUTOMATION-ONLY APPROACH
- BETWEEN DYSTOPIA AND UTOPIA
- THE FRONTIER OF AI SCIENCE SHOULD BE IN UNIVERSITIES
- THE MEANS OF PREDICTION
- IT IS NOT TOO LATE

ESSAYS

- STOP BUILDING BAD AI
- WORKPLACE TRAINING IN THE AGE OF AI
- MEDICINE'S MACHINE LEARNING PROBLEM

THE PAST AND FUTURE OF AI
CODING CARE
CONTRIBUTORS

EDITORS' NOTE

Joshua Cohen & Deborah Chasman

OUR WORLD is increasingly powered by artificial intelligence. The singularity is not here, but sophisticated machine-learning algorithms are—revolutionizing medicine and transport; transforming jobs and markets; reshaping where we eat, who we meet, what we read, and how we learn. At the same time, the promises of AI are increasingly overshadowed by its perils, from unemployment and disinformation to powerful new forms of bias and surveillance.

Leading off a forum that explores these issues, economist Daron Acemoglu argues that the threats—especially for work and democracy—are indeed serious, but the future is not settled. Just as technological development promoted broadly shared gains in the three decades following World War II, so AI can create inclusive prosperity and bolster democratic freedoms. Setting it to that task won't be easy, but it can be achieved through thoughtful government policy, the redirection of professional and industry norms, and robust democratic oversight.

Respondents to Acemoglu—economists, computer scientists, labor activists, and others—broaden the conversation by debating the role new technology plays in economic inequality, the range of algorithmic harms facing workers and citizens, and the additional steps that can be taken to ensure a just future for AI. Some ask how we can transform the way we design AI to create better jobs for workers. Others urge that we need new participatory methods in research, development, and deployment to address the unfair burdens AI bias has already imposed on vulnerable and marginal populations. Others argue that changes in social norms won't happen until workers have a seat at the table.

Contributions beyond the forum expand the aperture, exploring the impact of new technology on medicine and care work, the importance of workplace training in the AI economy, and the ethical case for not building certain forms of AI in the first place. In “Stop Building Bad AI,” Annette Zimmermann challenges the belief that something designed badly can later be repaired and improved, an industry-wide version of the Facebook motto to “move fast and break things.” She questions whether companies will police themselves, and instead calls for new frameworks for determining what kinds of AI are too risky to be designed in the first place.

What emerges from this remarkable mix of perspectives is a deeper understanding of the current challenges of AI and a rich, constructive, morally urgent vision for redirecting its course.

FORUM

REDESIGNING AI

Daron Acemoglu

ARTIFICIAL INTELLIGENCE (AI) is not likely to make humans redundant. Nor will it create superintelligence anytime soon. But, like it or not, AI technologies and intelligent systems will make huge advances in the next two decades—revolutionizing medicine, entertainment, and transport; transforming jobs and markets; enabling many new products and tools; and vastly increasing the amount of information that governments and companies have about individuals. Should we cherish and look forward to these developments, or fear them?

There are reasons to be concerned. Current AI research is too narrowly focused on making advances in a limited set of domains and pays insufficient attention to its disruptive effects on the very fabric of society. If AI technology continues to develop along its current path, it is likely to create social upheaval for at least two reasons. For one, AI will affect the future of jobs. Our current trajectory automates work to an excessive degree while refusing to invest in human productivity; further advances will displace workers and fail to create new opportunities (and, in the process, miss out on AI's full potential to enhance productivity). For another, AI may undermine democracy and individual freedoms.

Each of these directions is alarming, and the two together are ominous. Shared prosperity and democratic political participation do not just critically reinforce each other: they are the two backbones of our modern society. Worse still, the weakening of democracy makes formulating solutions to the adverse labor market and distributional effects of AI much more difficult. These dangers have only multiplied during the COVID-19 crisis. Lockdowns, social distancing, and workers' vulnerability to the virus have given an additional boost to the

drive for automation, with the majority of U.S. businesses reporting plans for more automation.

None of this is inevitable, however. The direction of AI development is not preordained. It can be altered to increase human productivity, create jobs and shared prosperity, and protect and bolster democratic freedoms—if we modify our approach. In order to redirect AI research toward a more productive path, we need to look at AI funding and regulation, the norms and priorities of AI researchers, and the societal oversight guiding these technologies and their applications.

Our Modern Compact

THE POSTWAR ERA witnessed a bewildering array of social and economic changes. Many social scientists in the first half of the twentieth century predicted that modern economies would lead to rising inequality and discontent, ultimately degenerating into various types of authoritarian governments or endless chaos.

The events of the interwar years seemed to confirm these gloomy forecasts. But in postwar Western Europe and North America—and several other parts of the globe that adopted similar economic and political institutions—the tide turned. After 1945 industrialized nations came to experience some of their best decades in terms of economic growth and social cohesion—what the French called *Les Trente Glorieuses*, the thirty glorious years. And that growth was not only rapid but also broadly shared. Over the first three decades after World War II, wages grew rapidly for all workers in the United States, regardless of education, gender, age, or race. Though this era was not without its political problems (it coincided with civil rights struggles in the United States), democratic politics worked: there was quite a bit of bipartisanship when it came to legislation, and Americans felt that they had a voice in politics. These two aspects of the postwar era were critical for social peace—a large fraction of the population understood that they were benefiting from the economic system and felt they had a

voice in how they were governed.

How did this relative harmony come about? Much of the credit goes to the trajectory of technological progress. The great economist John Maynard Keynes, who recognized the fragility of social peace in the face of economic hardship more astutely than most others, famously predicted in 1929 that economic growth would create increasing joblessness in the twentieth century. Keynes understood that there were tremendous opportunities for industrial automation—replacing human workers with machines—and concluded that declining demand for human labor was an ineluctable consequence of technological progress. As he put it: "We are being afflicted with a new disease of which . . . [readers] . . . will hear a great deal in the years to come—namely, technological unemployment."

Yet the technologies of the next half century turned out to be rather different from what Keynes had forecast. Demand for human labor grew and then grew some more. Keynes wasn't wrong about the forces of automation; mechanization of agriculture—substituting harvesters and tractors for human labor—caused massive dislocation and displacement for almost half of the workforce in the United States. Crucially, however, mechanization was accompanied by the introduction of new tasks, functions, and activities for humans. Agricultural mechanization was followed by rapid industrial automation, but this too was counterbalanced by other technological advances that created new tasks for workers. Today the majority of the workforce in all industrialized nations engages in tasks that did not exist when Keynes was writing (think of all the tasks involved in modern education, health care, communication, entertainment, back-office work, design, technical work on factory floors, and almost all of the service sector). Had it not been for these new tasks, Keynes would have been right. They not only spawned plentiful jobs but also generated demand for a diverse set of skills, underpinning the shared nature of modern economic growth.

Labor market institutions—such as minimum wages, collective bargaining, and regulations introducing worker protection—greatly contributed to shared prosperity. But

without the more human-friendly aspects of technological change, they would not have generated broad-based wage growth. If there were rapid advances in automation technology and no other technologies generating employment opportunities for most workers, minimum wages and collective wage demands would have been met with yet more automation. However, when these institutional arrangements protecting and empowering workers coexist with technological changes increasing worker productivity, they encourage the creation of "good jobs"—secure jobs with high wages. It makes sense to build long-term relationships with workers and pay them high wages when they are rapidly becoming more productive. It also makes sense to create good jobs and invest in worker productivity when labor market institutions rule out the low-wage path. Hence, technologies boosting human productivity and labor market institutions protecting workers were mutually self-reinforcing.

Indeed, good jobs became a mainstay of many postwar economies, and one of the key reasons that millions of people felt they were getting their fair share from the growth process—even if their bosses and some businessmen were becoming fabulously rich in the process.

Why was technology fueling wage growth? Why didn't it just automate jobs? Why was there a slew of new tasks and activities for workers, bolstering wage and employment growth? We don't know for sure. Existing evidence suggests a number of factors that may have helped boost the demand for human labor. In the decades following World War II, U.S. businesses operated in a broadly competitive environment. The biggest conglomerates of the early twentieth century had been broken up by Progressive Era reforms, and those that became dominant in the second half of the century, such as AT&T, faced similar antitrust action. This competitive environment produced a ferocious appetite for new technologies, including those that raised worker productivity.

These productivity enhancements created just the type of advantage firms were pining for in order to surge ahead of their rivals. Technology was not a gift from the heavens, of course.

Businesses invested heavily in technology and they benefited from government support. It wasn't just the eager investments in higher education during the Sputnik era (lest the United States fall behind the Soviet Union). It was also the government's role as a funding source, major purchaser of new technologies, and director and coordinator for research efforts. Via funding from the National Science Foundation, the National Institutes of Health, research and development tax credits, and perhaps even more importantly the Department of Defense, the government imprinted its long-term perspective on many of the iconic technologies of the era, including the Internet, computers, nanotechnology, biotech, antibiotics, sensors, and aviation technologies.

The United States also became more democratic during this period. Reforms during the Progressive and New Deal Eras reduced the direct control of large corporations and wealthy tycoons over the political process. The direct election of senators, enacted in 1913 in the Seventeenth Amendment, was an important step in this process. Then came the cleaning up of machine politics in many northern cities, a process that took several decades in the first half of the century. Equally important was the civil rights movement, which ended some of the most undemocratic aspects of U.S. politics (even if this is still a work in progress). Of course there were many fault lines, and not just Black Americans but many groups did not have sufficient voice in politics. All the same when political scientist Robert Dahl set out to investigate “who governs” local politics in New Haven, the answer wasn't an established party or a well-defined elite. Power was pluralistic, and the involvement of regular people in politics was key for the governance of the city.

Democracy and shared prosperity thus bolstered each other during this epoch. Democratic politics strengthened labor market institutions protecting workers and efforts to increase worker productivity, while shared prosperity simultaneously increased the legitimacy of the democratic system. And this trend was robust: despite myriad cultural and institutional differences, Western Europe, Canada, and Japan followed remarkably similar trajectories to that of the United States,

based on rapid productivity growth, shared prosperity, and democratic politics.

The World Automation is Making

WE LIVE in a very different world today. Wage growth since the late 1970s has been much slower than during the previous three decades. And this growth has been anything but shared. While wages for workers at the very top of the income distribution—those in the highest tenth percentile of earnings or those with postgraduate degrees—have continued to grow, workers with a high school diploma or less have seen their real earnings fall. Even college graduates have gone through lengthy periods of little real wage growth.

Many factors have played a role in this turnaround. The erosion of the real value of the minimum wage, which has fallen by more than 30 percent since 1968, has been instrumental in the wage declines at the bottom of the distribution. With the disappearance of trade unions from much of the private sector, wages also lagged behind productivity growth. Simultaneously, the enormous increase in trade with China led to the closure of many businesses and large job losses in low-tech manufacturing industries such as textiles, apparel, furniture, and toys. Equally defining has been the new direction of technological progress. While in the four decades after World War II automation and new tasks contributing to labor demand went hand-in-hand, a very different technological tableau began in the 1980s—a lot more automation and a lot less of everything else.

Automation acted as the handmaiden of inequality. New technologies primarily automated the more routine tasks in clerical occupations and on factory floors. This meant the demand and wages of workers specializing in blue-collar jobs and some clerical functions declined. Meanwhile professionals in managerial, engineering, finance, consulting, and design occupations flourished—both because they were essential to the success of new technologies and because they benefited from


image
not
available

standard algorithms used by social media sites such as Facebook reduce users' exposure to posts from different points of view, further contributing to the polarization of the U.S. public.

Other emerging applications of AI may be even more threatening to democracy and liberty around the world. Basic pattern recognition techniques are already powerful enough to enable governments and companies to monitor individual behavior, political views, and communication. For example, the Chinese Communist Party has long relied on these technologies for identifying and stamping out online dissent and opposition, for mass surveillance, and for controlling political activity in parts of the country where there is widespread opposition to its rule (such as Xinjiang and Tibet). As Edward Snowden's revelations laid bare, the U.S. government eagerly used similar techniques to collect massive amounts of data from the communications of both foreigners and American citizens. Spyware programs—such as Pegasus, developed by the Israeli firm NSO Group, and the Da Vinci and Galileo platforms of the Italian company Hacking Team—enable users to take control of the data of individuals thousands of miles away, break encryption, and remotely track private communications. Future AI capabilities will go far beyond these activities.

Another area of considerable concern is facial recognition, currently one of the most active fields of research within AI. Though facial recognition technology has legitimate uses in personal security and defense against terrorism, its commercial applications remain unproven. Much of the demand for this technology originates from mass surveillance programs.

With AI-powered technologies already able to collect information about individual behavior, track communications, and recognize faces and voices, it is not far-fetched to imagine that many governments will be better positioned to control dissent and discourage opposition. But the effects of these technologies may well go beyond silencing governments' most vocal critics. With the knowledge that such technologies are monitoring their every behavior, individuals will be discouraged from voicing criticism and may gradually reduce their