

Written Testimony

Hearing on “The Philosophy of AI: Learning from History, Shaping Our Future”

Senate Committee on Homeland Security and Governmental Affairs

Daron Acemoglu

Institute Professor, MIT

November 8, 2023

Chairman Peters, Ranking Member Paul, and Members of the Committee,

Thank you for inviting me to testify today on this important subject. I will start with a summary of my main points and then expand upon them in the rest of the testimony.

Executive Summary: Digital technologies have already ushered in a multifaceted economic, social and political transformation. Artificial intelligence (AI) promises to amplify these epochal changes, for good and bad. Although these tools have tremendous potential to expand our production, communication, and informational capabilities, they also pose major risks to economic prosperity, social cohesion, democracy, and national security — as did many other transformative technologies in the past.

These risks are rooted in three related social changes: (1) economic shifts, especially greater inequality, brought about by new technologies can create social and political tensions; (2) digital tools, including AI, alter who controls information and how that information can be used and manipulated, with direct implications for political behavior and democracy; (3) these technologies also unleash myriad social changes, affecting aspirations and norms, with potentially far-reaching effects. All of these risks apply to both democracy and national security. It is critical to understand them, learn from history about when humanity has and has not managed to develop institutions and norms to deal with similar risks, and chart a clear-eyed regulatory course to guard against the worst eventualities.

My overarching argument is that there is a pro-human (meaning pro-worker and pro-citizen) direction for AI tools that would be much better both for shared prosperity and for democracy, and therefore for national security. We need to take AI risks seriously because—although a pro-human direction for AI could strengthen prosperity, democracy, and security—we are currently on a very different and worrying trajectory.

The AI Revolution: The Possibility of a Pro-Human Direction

By the 2010s, digital technologies had already fundamentally changed our economy and communication, and AI has further amplified these trends. Computers expand our capacity for computation and information processing. When their capabilities are applied to physical processes, they can enable both automation of tasks previously performed by humans (for example, numerically controlled machinery for manufacturing and software systems for office work) and the creation of new tasks for workers (for example, computer-assisted design). They also change how we collect and process information.

Basic computers are pre-programmed with step-by-step instructions to perform basic operations. This keeps their applications limited to routine tasks (that can be performed in a repetitive manner within a stable, predictable environment). AI expands these capabilities by providing greater flexibility, both because not everything needs to be pre-programmed (the learning aspect of AI) and because non-routine tasks can also be tackled by these tools. This flexibility not only amplifies the possibilities for automation, but could also enable better human decision-making and the capacity to perform new and more complex tasks.¹

To improve human performance, we need to think beyond creating AI systems that aim to acquire “higher-level capabilities” or “artificial general intelligence” (AGI) whereby machines can achieve parity with, and even surpass, humans in all mental tasks. Leaving aside whether AGI is even feasible within the next several decades (which, I believe, it is not), my argument is that the emphasis on general intelligence distracts us from the more beneficial uses of digital

¹ Russell, Stuart and Peter Norvig (2009) *Artificial Intelligence: A Modern Approach*, Pearson; Susskind, Daniel (2021) “Technological Unemployment” from *The Oxford Handbook of AI Governance*. Oxford University Press (forthcoming): <https://www.danielsusskind.com/s/Susskind-Handbook-Updated-21-September-2021.pdf>.

technologies to expand human capabilities by providing better information for human problem-solving and decision-making.

Making machines useful to humans and complementary to workers is not a new aspiration. Related ideas were articulated as early as 1949 by MIT mathematician and engineer Norbert Wiener, and put into practice by several leading computer scientists and engineers, such as Douglas Engelbart, inventor of technologies like the computer mouse and hypertext.² Such human-complementary tools, as well as new human tasks and appropriate worker training in the manufacturing and service sectors, were foundational to rapid wage growth and shared prosperity throughout the 1950s, 1960s and early 1970s. Nevertheless, this “machine usefulness” agenda has often been overshadowed by other priorities in today’s tech industry, including a strong focus on automation and monetization models based on data collection and digital ads.

Machine usefulness may be even more promising today than in the past. Many of the routine tasks that workers previously performed have already been automated, so a large fraction of current jobs centrally depends on nonroutine problem-solving and decision-making tasks. Empowering workers to perform these tasks more effectively, and to accomplish even more sophisticated decision-making tasks, will require providing workers with better information and decision-support tools. Recent advances in AI, especially in generative AI tools, are particularly well-suited to this type of information provision.³ An irony of our current digital era is that information is abundant, but useful information is scarce. Generative AI can help by recognizing the relevant context and presenting information that is useful for problem-solving, human decision-making, and performance in new, more complex tasks. For example, an electrician can much more effectively diagnose rare problems and accomplish complex tasks when empowered with AI tools that present information and recommendations based on knowledge accumulated from similar cases. In essence, AI holds great potential for training and provision of useful

² Acemoglu, Daron and Simon Johnson (2023a) *Power and Progress: Our Thousand-Year Struggle over Technology and Prosperity*, Hachette Public Affairs [hereafter *Acemoglu and Johnson (2023a)*], Chapter 9.

³ Acemoglu, Daron, David Autor and Simon Johnson (2023) “Can We Have Pro-Worker AI? Choosing a Path of Machines in Service of Minds” MIT Shaping the Future of Work policy memo: <https://shapingwork.mit.edu/wp-content/uploads/2023/09/Pro-Worker-AI-Policy-Memo.pdf>.

information to support expert workers, such as educators, medical personnel, software developers, and other skilled craft workers such as electricians and plumbers.⁴

These potential pro-worker gains from AI are not hypothetical. Several studies already show that, if used to provide better information or successfully speed up the more-routine parts of tasks for workers, generative AI can provide a significant boost to productivity.⁵

There is an analogue to the pro-worker perspective when it comes to how AI tools are used for communication. What we might call a “pro-citizen” perspective would leverage these tools to provide better information to individuals about the political and social world, and enable them to participate in deliberations without manipulation or undue bias. The opposite approach would be to deploy these tools for surveillance, manipulation, manufacturing false conformity, and fomenting emotional outrage and extremism.

The evolution of social media in the 2010s illustrates the manipulative path, where much evidence shows that algorithmic tools were used to boost user engagement by intensifying echo chambers, spreading extremist or misleading content, and exploiting emotional triggers.⁶ Although there is much we do not know about the systemic political and social effects of social media, existing evidence suggests that they have contributed both to mental health problems and polarization in politics.

The alternative, pro-citizen direction was articulated by a number of optimistic takes on the Internet and social media in the early 2000s, which viewed these tools as expanding democracy and increasing the accountability of authoritarian rulers.⁷ While these hopes were

⁴ Acemoglu, Daron (2021) *Redesigning AI: Work, Democracy, and Justice in the Age of Automation*, Boston Review Forum [hereinafter *Acemoglu 2021*].

⁵ Brynjolfsson, Erik, Danielle Li and Lindsey Raymond (2023) “Generative AI at Work.” NBER Working Paper no. 31161; Noy, Shakked and Whitney Zhang (2023) “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence.” *Science*, 381(6654): 187–192; Peng, Sida, Eirini Kalliamvakou, Peter Cihon and Mert Demirel (2023) “The Impact of AI on Developer Productivity: Evidence from GitHub Copilot.” arXiv Working Paper no. 2302.06590.

⁶ Brady, William, Julian Wills, John Jost, Joshua Tucker and Jay Van Bavel (2017) “Emotion Shapes the Diffusion of Moralized Content in Social Networks” *PNAS* 114(28): 7313–7318; Braghieri, Luca, Ro’ee Levy and Alexey Makarin (2022), *American Economic Review* 112(11): 3660–3693; Wu, Tim (2016) *The Attention Merchants: The Epic Scramble to Get Inside Our Heads*, PRH Knopf; Acemoglu, Daron, Asuman Ozdaglar and James Siderius (2023) “A Model of Online Misinformation,” *Review of Economic Studies* (forthcoming).

⁷ Acemoglu and Johnson (2023a), Chapter 10.

soon dashed, prototypes of more pro-citizen AI tools have been developed successfully, including at significant scale in Taiwan, where they are used to increase the transparency of government processes and services, encourage democratic participation, and enable citizen input on government policies.⁸

With the widely different possibilities that AI tools present, it becomes particularly important to understand which direction we are heading in today, and how different paths might impact our economy and democracy.

Computers, AI and the Economy

It is well known that the American economy has not generated shared prosperity over the last four decades. U.S. labor market inequality has surged since 1980. Many workers—especially men with only a high-school education or less—have experienced significant declines in their real earnings. Even men with a college degree have seen only limited gains. In the meantime, racial inequities have widened, and regional disparities have multiplied. These sweeping changes have many causes, ranging from globalization to the transformation of U.S. labor market institutions.⁹

These trends cannot be separated from the effects of computers. My research indicates that the most important cause of increasing inequality has been automation—the substitution of machines and algorithms for tasks previously performed by workers. Automation is not a recent phenomenon: for example, it was pervasive during the British Industrial Revolution and in the United States during the 1800s.

Two types of evidence illustrate the inequality implications of today’s automation technologies, such as robotics and office software. First, in regions where industrial robots have been adopted faster, employment and wages have fallen more and inequality has increased.¹⁰

⁸ Tang, Audrey (2019) “A Strong Democracy Is a Digital Democracy,” *New York Times*: <https://www.nytimes.com/2019/10/15/opinion/taiwan-digital-democracy.html>; Lanier, Jaron and E. Glen Weyl (2020) “How Civic Technology Can Help Stop a Pandemic,” *Foreign Affairs*: <https://www.foreignaffairs.com/articles/asia/2020-03-20/how-civic-technology-can-help-stop-pandemic>.

⁹ Acemoglu, Daron and David Autor (2011) “Skills, Tasks and Technologies: Implications for Employment and Earnings,” *Handbook of Labor Economics*, Chapter 12, 4(B): 1043–1171; Autor, David (2019) “Work of the Past, Work of the Future,” *AEA Papers and Proceedings*, 109: 1–32.

¹⁰ Acemoglu, Daron and Pascual Restrepo (2020) “Robots and Jobs: Evidence from U.S. Labor Markets,” *Journal of Political Economy*, 128(6): 2188–2244.

Second, demographic groups employed to perform routine tasks in industries undergoing rapid automation have almost uniformly suffered large declines in their real wages. My work finds that automation accounts for more than half, and perhaps as much as three quarters, of the surge in U.S. wage inequality.¹¹ Other trends, such as offshoring and competition from cheaper Chinese imports, have also played a role, but they have been less important than automation.¹²

Automation has contributed to specific dimensions of economic disparities as well. It accounts for 80% of the sizable increase in the college wage premium.¹³ It has also widened the gap between Black and White Americans. Automation, together with globalization, has been a major driver of regional disparities.

Concurrently with the increases in labor market inequality, the gap between capital owners and labor has also widened in the United States. The share of national income distributed as capital (e.g., interest, profits, dividends, capital gains) has risen sharply since the 1990s, as the share to labor has fallen. Among capital owners and managers, the very top have done particularly well, so that the shares of the richest 1% and 0.1% of Americans in national income have skyrocketed.¹⁴ Automation appears to have been an important factor in these trends as well.

What about the benefits of automation? During the mechanization of agriculture and in the three decades following World War II, automation was rapid, and the U.S. economy created millions of good jobs and achieved broadly shared prosperity. The main difference is that new technologies throughout these eras not only automated existing tasks, but also created new tasks for workers. Henry Ford's breakthroughs in automobile manufacturing were due to the application of new electrical machinery that routinized and automated some manual low-skill and artisanal tasks. But they also introduced a range of new technical tasks and critical functions for blue-collar workers. Simultaneously, manufacturing became much more intensive in

¹¹ Acemoglu, Daron and Pascual Restrepo (2022) "Automation, Tasks and the Rise in U.S. Wage Inequality," *Econometrica*, 90(5): 1973–2016 [hereinafter *Acemoglu and Restrepo (2022)*].

¹² Autor, David, David Dorn, and Gordon Hanson (2013) "The China Syndrome: Local Labor Market Effects of Import Competition in the United States" *American Economic Review* 103(6): 2121–2168; Acemoglu, Daron, David Autor, David Dorn, Gordon Hanson and Brendan Price (2016) "Import Competition and the Great U.S. Employment Sag of the 2000s," *Journal of Labor Economics* 34(S1.2): S141–S198 [hereinafter *Acemoglu et al. (2016)*]; Acemoglu and Restrepo (2022).

¹³ Acemoglu and Restrepo (2022).

¹⁴ Piketty, Thomas and Emmanuel Saez (2003) "Income Inequality in the United States, 1913–1998," *Quarterly Journal of Economics* 118(1): 1–41; Moll, Benjamin, Lukasz Rachel and Pascual Restrepo (2022) *Econometrica* 90(6): 2645–2683.

information activities, including design, planning, inspection, quality control and a variety of back-office tasks. These new tasks were critical for the expansion of employment in the automobile industry and throughout the U.S. economy during this era.¹⁵

The same forces were foundational during the decades that followed World War II, which witnessed both robust productivity growth and rapid wage growth that was broadly shared across demographic groups. While automation of certain manufacturing and white-collar tasks was rapid during this growth episode, so was the introduction of a range of new worker tasks, which were supported by novel technologies and training. Put simply, the 1950s, 1960s and early 1970s ushered in rapid growth and shared prosperity because automation went hand-in-hand with new tasks that increased workers' contributions to the production process.¹⁶ There has been a breakdown of shared prosperity since the 1980s because digital technologies have been used predominantly for automation and insufficiently for creating new tasks.

It is not certain whether the automation focus of early computers and other digital technologies was because their capabilities made automation easier (but did not enable enough new task creation) or if, instead, there was an excessive focus on automation that neglected alternative, potential pro-worker uses of these technologies. My research suggests that the latter explanation is more likely. Two pieces of evidence support this interpretation. First, when incentives and priorities were aligned, computers were used in more pro-worker ways, and relatedly, cross-country differences indicate that when institutions and policies encourage more pro-worker outcomes, technologies respond as well.¹⁷ Second, despite the great promise of digital tools, productivity gains from computers and other digital technologies have been disappointing, and there is some evidence that this has been because of an overemphasis on automation.¹⁸

¹⁵ Acemoglu and Johnson (2023a), Chapter 7; Hounshell, David (1985) *From the American System to Mass Production, 1800–1932: The Development of Manufacturing Technology in the United States*, Johns Hopkins University Press [hereinafter *Hounshell (1985)*].

¹⁶ Acemoglu, Daron and Pascual Restrepo (2019) "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, 33(2): 3-30; Autor, David, Caroline Chin, Anna Salomons and Bryan Seegmiller (2022) "New Frontiers: The Origins and Content of New Work, 1940–2018" NBER Working Paper no. 30389.

¹⁷ Acemoglu and Johnson (2023a), Chapter 7; Hounshell (1985).

¹⁸ Acemoglu et al. (2016); Gordon, Robert (2016) *The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War*, Princeton University Press; Acemoglu, Daron, David Autor and Christina Patterson (2023) "Bottlenecks: Sectoral Imbalances and the U.S. Productivity Slowdown," NBER Macroeconomics Annual 2023(38).

The evidence I have presented so far pertains to pre-AI digital technologies. AI can be used in even more versatile ways than early computers and opens up new pathways for pro-worker technologies. Nevertheless, recent research also suggests that emerging AI applications are targeting automation. For instance, establishments with jobs and tasks that can be automated by AI have been at the forefront of AI adoption in the late 2010s and the same establishments have concurrently cut back on their hiring.¹⁹ There is also evidence suggesting that AI tools are often used for surveillance and intensive monitoring, which are costly for workers and do not generate much productivity growth.²⁰ The emphasis on surveillance is much more intense in China, triggered by the Communist Party's preoccupation with controlling the population's political beliefs and activities. The available evidence suggests that surveillance and censorship tools produced in China are spreading throughout the world.²¹

The evolution of the market structure for AI providers will matter for how these tools are incorporated into business applications and who captures the gains. Two very different paths appear possible. In one, open-source models and the entry of many small generative AI providers would create a fairly competitive environment.²² This is not a panacea against many of the problems I have pointed out, but would at least prevent monopolization of these tools, as well as monopoly control of the massive amounts of data that the tech industry is collecting. The alternative is a much more oligopolistic structure, where a few companies that have a lead in data collection and in building large, all-purpose generative AI models, such as GPT-4, will become dominant. Along this path, these leading companies would capture a disproportionate share of the profits from new AI tools. This oligopolistic model would intensify inequality (as a few companies earn most of the profits), and could increase the influence of these players on the direction of technology and how centralized control of information plays out (an issue I return to below).²³

¹⁹ Acemoglu, Daron, David Autor, Jonathon Hazell and Pascual Restrepo (2022) "Artificial Intelligence and Jobs: Evidence from Online Vacancies," *Journal of Labor Economics* 40(S1): S293-S340.

²⁰ Acemoglu and Johnson (2023a), Chapter 9; Hounshell (1985).

²¹ Acemoglu and Johnson (2023a), Chapters 9 and 10; Beraja, Martin, Andrew Kao, David Yang and Noam Yuchtman (2023) "Exporting the Surveillance State via Trade in AI," NBER Working Paper no. 31676.

²² Masiello, Betsy and Derek Slater (2023) "Will Open Source AI Shift Power from 'Big Tech'? It Depends," *Tech Policy Press*: <https://techpolicy.press/will-open-source-ai-shift-power-from-big-tech-it-depends/>.

²³ Acemoglu, Daron and Simon Johnson (2023b) "Big Tech Is Bad, Big AI Will Be Worse," *New York Times*: <https://www.nytimes.com/2023/06/09/opinion/ai-big-tech-microsoft-google-duopoly.html>.

In summary, digital technologies and AI present a range of different possibilities for our economic future. A pro-worker direction would leverage these tools to make human workers more productive and provide them with better information for problem-solving and critical decision-making. The alternative, using them for automation and surveillance, would not bring the same benefits for workers and would tend to boost inequality. The evidence we currently have suggests that digital automation has been at the root of the spectacular increase in U.S. inequality since 1980, and early AI technologies are going in the same direction of automation and surveillance.

Social and Democratic Implications of Inequality

A thesis dating back at least to Alexis de Tocqueville maintains that democracy depends on limiting social and economic inequality.²⁴ People need to participate in democratic decision-making as citizens, and this becomes more likely when they are active socially, when they believe that they have a chance to succeed economically, and when they are getting their fair share of the economic benefits. Economic crises often lead to political instability, and recent evidence from the United States and from around the world supports these concerns.

In the United States, areas that were most severely disrupted by the rise in Chinese competition, which led to business closures and greater joblessness, show a sharp increase in political polarization and support for more extremist political candidates.²⁵ These trends are visible through both changes in general attitudes and in voting behavior. There is a similar increase in polarization and support for more extremist candidates in places where the introduction of robots led to declines in employment and wages.²⁶

Additionally, my research across a large number of countries shows that high inequality undermines support for democracy among the population, and this lack of support makes democracies more unstable and less capable of dealing with adverse shocks.²⁷

²⁴ de Tocqueville, Alexis (1835) *Democracy in America*, Eds. Harvey Mansfield and Delba Winthrop (2000), University of Chicago Press.

²⁵ Autor, David, David Dorn, Gordon Hanson and Kaveh Majlesi (2020) "Importing Political Polarization? The Electoral Consequences of Rising Trade Exposure," *American Economic Review* 110(10): 3139–3183.

²⁶ Frey, Carl, Thor Berger and Chinchih Chin (2018) "Political Machinery: Did Robots Swing the 2016 U.S. Presidential Election?" *Oxford Review of Economic Policy* 34(3): 418–442.

²⁷ Acemoglu, Daron, Nicolás Ajzenman, Cevat Aksoy, Martin Fiszbein and Carlos Molina (2023) "(Successful) Democracies Breed Their Own Support," *Review of Economic Studies* (forthcoming).

This body of work thus suggests that the inequality effects of digital automation have likely been a major factor in the rise of anti-democratic views in parts of the U.S. electorate and the more general political polarization of the country. By fueling extremism, these changes are raising new challenges for national security.

If AI deepens inequality, the dangers for U.S. democracy and national security could become intensified as well. The risks may be even more pronounced because several decades of non-shared prosperity could potentially have much bigger effects on grievances and discontent than shorter-term trends. This perspective suggests that it is particularly important to find a path for new AI tools that would help build shared prosperity, especially after the huge increases in inequality in the country over the last four decades.

Control over Information and the Future of Democracy

The survival of any political regime depends on the beliefs and values of its people and on how information is controlled and presented. Authoritarian rulers have understood this for ages. The top-down, highly repressive rule that the Qin Dynasty imposed in China after the Warring States period in third century B.C. was accompanied by much more intense control over information, including book burnings. The Qin rulers were sufficiently worried about information spreading in the population that they are reported to have executed the people who wrote the books, as well as those who could rewrite them.²⁸

We also see the importance of control of information, and its breakdown, in the history of Europe. The printing press, which broke the monopoly of political and religious elites on information, was not just critical for the Protestant Reformation, but was a direct contributor to the momentous political changes that followed. Pamphlets, newspapers and books that could be quickly printed and disseminated were important in the English Civil War and the run-up to the French Revolution. Many rulers, including Thomas Cromwell in his efforts to promote Henry

²⁸ Chan, Lois Mai (1972) "The Burning of the Books in China, 213 B.C.", *Journal of Library History*, 7(2): 101-108.

VIII's centralization project and Napoleon Bonaparte in his efforts to inculcate the French identity, also used the printing press and control over information that it enabled.²⁹

As information tools, computers have opened myriad possibilities for citizens to access information and communicate. They have also created new pathways for surveillance and centralized control of information. These trends are clearly visible in Iran, Russia and China. For example, Internet censorship and digital surveillance have become the bulwark of the Chinese Communist Party's rule over the last two decades. Part of the reason why early hopes about the democratizing effects of social media and the Internet have been dashed is precisely because governments have been able to mobilize significant resources to establish their control over online information.³⁰

While the parallels between the current and past authoritarian efforts to control information are most evident, the issues are relevant to democracies as well, for two related reasons. First, it is not only the survival of nondemocratic regime that depends on who controls information. Threats from extremist groups against many democratic governments in the interwar years were intensified when information manipulation through print media and radio became widespread—for example, by Father Coughlin in the United States and the insurgent Nazi Party during the Weimar Republic in Germany.³¹ Second, the political and social consequences of concentration (or even monopolization) of information in the hands of a few private corporations, which has been a defining feature of the digital era, are poorly understood.

Powerful media empires have been a mainstay of the 20th century. But the extent to which digital platforms today have access to a vast amount of public and private information is unprecedented. Google can track the metadata of hundreds of millions of users from their search, email, driving and shopping activity. Amazon has even more detailed data on purchasing behavior and evolving preferences for its massive user base, and social media platforms have unparalleled data about the emotional, political and social states of billions of people. There is

²⁹ Einstein, Elizabeth (1993) *The Printing Revolution in Early Modern Europe*, Cambridge University Press; Darnton, Robert (1995) *The Forbidden Best-Sellers of Pre-Revolutionary France*, W.W. Norton & Co; Raymond, Joad (2005) *The Invention of the Newspaper: English Newsbooks 1641–1649*, Oxford University Press.

³⁰ Acemoglu and Johnson (2023a), Chapters 9 and 10; Beraja, Martin, David Yang and Noam Yuchtman (2023) "Data-Intensive Innovation and the State: Evidence from AI Firms in China," *Review of Economic Studies* 90(4): 1701–1723.

³¹ Acemoglu and Johnson (2023a), Chapter 10.

mounting evidence that these platforms have sometimes manipulated this information for their own private benefits, as recent lawsuits against Meta, Amazon and Alphabet attest.

Reining in anti-competitive behavior is important for a fair and functioning market economy. But the monopolized control of information may have even more sweeping implications for democracy and national security.

The business model of most digital platforms in the industrialized world today depends on collection of personal data, which is then monetized via individualized digital ads. This business model encourages a particular way in which data is collected, processed and used. Most importantly, digital ads are more profitable when users are more frequently on the platform and are more intensively engaged with the offered content. This has prompted the development of algorithmic methods for boosting user engagement, often by presenting more provocative, sensational and emotionally charged material. A growing body of work in social psychology documents—and, in fact, teaches businesses—how platforms can increase engagement by manipulating user perceptions and presenting them with varying stimuli, often in the form of emotional cues.³² These cues take different forms for different subpopulations—related to envy and insecurity for teenagers and based on extremist and sensationalist material for politically engaged populations. Regardless of its exact form, this business model appears to have contributed to polarization and political disengagement, as well as to mental health problems within the population.

Algorithms have played an important role in the emergence of these new online business models. More powerful AI capabilities are significantly expanding what platforms can do, as they are able to sift through more massive amounts of data and identify new person-specific vulnerabilities.³³

Although the evidence is not conclusive, a reasonable reading of prior studies is that generative AI tools will amplify the manipulative capabilities of digital platforms, and depending on how the industry evolves, the degree to which control of information becomes more

³² Eyal, Nir (2014) *Hooked: How to Build Habit-Forming Products*, PRH Portfolio.

³³ Haidt, Jonathan and Eric Schmidt (2023) "AI Is About To Make Social Media (Much) More Toxic: We Must Prepare Now," *The Atlantic*: <https://www.theatlantic.com/technology/archive/2023/05/generative-ai-social-media-integration-dangers-disinformation-addiction/673940/>.

monopolized could increase as well. This perspective suggests that risks to democracy and national security from how AI develops are first order.

As in the realm of production technologies, there are critical choices about the path of AI in communication. The pro-citizen path would empower individuals, provide safeguards against misinformation and offer new tools for community-building and pro-democracy activities; the manipulation path would be controlled by a few corporate players and government entities that can mislead users as they choose. Which one we choose will have major implications for democracy and national security.

Broader Social Implications

New technologies sometimes inculcate new values and practices. When they do so, they can deeply impact our political and social arrangements. The printing press is a case in point.

In the case of AI, an important set of issues will center on how people’s social priorities and perceptions of their communal duties evolve.³⁴ If we proceed along the path that sidelines workers and disempowers citizens, we may also expect a greater level of disengagement and passivity within the population.³⁵ Although there is little evidence on these issues, the rise in mental health problems and withdrawal of younger generations into online environments—away from real-world social networks, community-level activities and traditional political participation—may auger more systemic social changes, with potential effects on democracy. Historical evidence suggests that economic and social means of encouraging active citizenry is critical for democratic institutions.³⁶

The direction of AI technologies—between pro-worker vs. automation-focused and between pro-citizen vs. top-down manipulative—will almost certainly impact norms and aspirations. If new technologies sideline workers, discourage political participation and promote

³⁴ Sandel, Michael (1998) *Democracy's Discontent: America in Search of a Public Philosophy*, Harvard University Press Belknap.

³⁵ Wike, Richard and Alexandra Castillo (2018) “Many Around the World Are Disengaged From Politics” *Pew Research Center Report*; Gale, Steven and Mat Burrows (2022) “Disillusioned Youth: A Danger to Democracy” *Wilson Center*: <https://www.newsecuritybeat.org/2022/06/disillusioned-youth-danger-democracy/>.

³⁶ Acemoglu, Daron and James A. Robinson (2012) *Why Nations Fail: The Origins of Power, Prosperity and Poverty*, PRH Crown Business.

manipulation on online platforms, they may raise new risks to democracy and national security. We would be wise to study and track how political and social values, means of information consumption and political participation will change in the age of AI.

Learning from History

AI is a new technology. But there have been other transformative technologies, and history offers important clues about how to manage them. The most critical lesson is that democratic governance, both in the context of how to best develop and use new technologies and how to confront the social changes that they are unleashing, is critical.

The Industrial Revolution, which started in Britain sometime in the middle of the 18th century, is today remembered as the beginning of our modern age, ushering in greater prosperity, better health and unparalleled comfort to current generations. This is true, but is only part of the story. The first 100 years of the Industrial Revolution were simply awful for the working classes. Real incomes stagnated, working hours increased, working conditions worsened in the modern factories that were highly disciplinarian and inhospitable, and life expectancy and health conditions deteriorated in the face of intensifying pollution and uncontrolled epidemics. These pitiful conditions cannot be understood without recognizing that a lot of early industrial technologies focused on automation and worker control, and there were essentially no institutions protecting workers economically or socially.³⁷

The more positive developments after 1850 stand in stark contrast to this earlier era. But there was nothing automatic about this turnaround in events. It took a major redirection of technology away from a singular focus on automation and toward the goals of raising worker productivity and increasing their roles in production, and this was embedded in fundamental political and social changes. In Britain, only the very privileged could participate in politics at the beginning of the 19th century, until the adoption of universal male suffrage and then universal adult suffrage between 1867 and 1928.³⁸ Trade unions, which had been banned and heavily

³⁷ Acemoglu and Johnson (2023a), Chapter 6.

³⁸ Acemoglu, Daron and James A. Robinson (2006) *Economic Origins of Dictatorship and Democracy*, Cambridge University Press.

prosecuted, became legal, and workers began to campaign for higher wages and better working conditions. Draconian “master and servant” laws that made it difficult for workers to quit their jobs were abolished. Investment in public infrastructure and public health started cleaning up the cities and brought diseases under control.³⁹ Mass education and basic social safety net provisions improved the skill level and welfare of the working classes.

The general lesson is that a pro-human direction for technology, democracy and shared prosperity were symbiotic during the Industrial Revolution and similar periods of history. This lesson may have particular importance today, as it reminds us that risks against shared prosperity and democracy are tightly entangled.

Choosing the Right Path for AI

I have presented my take on the state of current knowledge about digital technologies and their economic, social and political effects. We have widely different paths ahead of us on how AI technologies will develop and their implications for productivity, inequality, social norms and democracy.

Given the fast-changing nature of AI technologies, it is impossible to be certain on any of these issues. Nevertheless, there is considerable evidence that a pro-human direction of AI—pro-worker in production, pro-citizen in communication—would be a better foundation on which to build renewed shared prosperity and would likely create a far better environment for democracy.

Is our current trajectory likely to become pro-human? I believe the answer is no. There is ample evidence that the automation focus of AI is continuing and that the most pernicious uses of centralized information by private and government actors are intensifying. It is important to understand the reasons for these global trends and develop regulatory, institutional and policy responses.

Let me first elaborate on why we are currently heading in the wrong direction. Many U.S. corporations are focused on cost-cutting due to competitive pressures or to target short-run

³⁹ Steinfeld, Robert (1991) *The Invention of Free Labor: The Employment Relation in English and American Law and Culture, 1350–1870*, University of North Carolina Press.

performance metrics. This often implies that increasing the contribution of employees to long-run performance does not receive as much attention as it deserves. Additionally, the tech industry has been dominated by the business models of the largest corporations, which prioritize the development of automation tools or monetization based on digital ads, as I have explained. Finally, the emphasis on artificial general intelligence and more broadly on reaching “human parity” have also become key metrics in the tech industry, sidelining the objective of maximizing machine usefulness.⁴⁰

If this diagnosis is correct, then government policy may be necessary to redirect AI in a more pro-human direction. There is no silver bullet for doing so, and just recognizing that a pro-human direction is both feasible and desirable may be an already important step. But there are also some policy ideas worth considering.

First, the government can set goals that are more pro-worker, and this should start with additional support for the research and development of human-complementary AI technologies. Though it is hard to target the direction of AI in the abstract, there are many areas in which opportunities for human-complementary work abound. These include education and training, healthcare, and technologies for better decision-making in skilled craft work.⁴¹ Just as DARPA orchestrated investments and competitions to foster the development of self-driving cars and dexterous robotics, sufficient investment from the federal government can foster breakthroughs in pro-worker AI.⁴² A new federal agency (perhaps “National Institutes of Computing” modeled on the National Institutes of Health) may be necessary to coordinate research effort, funding and priorities in this area.

Second, the U.S. tax code currently places a heavier burden on firms that hire labor than on those that invest in algorithms to automate work. Firms and workers jointly face a tax burden

⁴⁰ Acemoglu, Daron, Danielle Allen, Kate Crawford, James Evans, Michael Jordan, Divya Siddarth and E. Glen Weyl (2021) “How AI Fails Us,” *Edmond J. Safra Center for Ethics, Harvard: Justice, Health, and Democracy Impact Initiative*: https://ethics.harvard.edu/files/center-for-ethics/files/howai_fails_us_2.pdf; Brynjolfsson, Erik (2022) “The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence,” *Stanford Digital Economy Lab: Human-Centered Artificial Intelligence (HAI) Lab*: <https://digitaleconomy.stanford.edu/news/the-turing-trap-the-promise-peril-of-human-like-artificial-intelligence/>.

⁴¹ Acemoglu, Daron (2021) *Redesigning AI: Work, Democracy, and Justice in the Age of Automation*, Boston Review Forum.

⁴² Gruber, Jonathan and Simon Johnson (2019) *Jump-Starting America: How Breakthrough Science Can Revive Economic Growth and the American Dream*, Hachette PublicAffairs.

of 25–30% for labor, while firms and capital owners pay only about 5% net when investing in equipment and computers.⁴³ A more symmetric tax structure, where marginal taxes for hiring and training labor and for investing in equipment and software are equated, may provide a better foundation upon which to build pro-worker technologies. A fairer tax system can be achieved by reducing or eliminating payroll taxes, lessening capital income tax deductions, or increasing corporate income taxes.

Third, well-functioning data markets may be an important step for pro-human AI.⁴⁴ Large-scale data collection is at the root of the manipulative use of information in online platforms. Moreover, the current path of generative AI tools is partly shaped by the possibility of using freely available data with few constraints, which also discourages investment in higher-quality data. Legislation can support a functioning data market by establishing property rights to workers and citizens for their data, which could then enable companies to invest in high-quality data. Such an environment could be a better basis for a more pro-worker trajectory of AI and may also discourage the dominance of business models based solely on maximizing user attention and engagement.

Fourth, digital ad taxes may be needed to discourage most manipulative online practices.⁴⁵ Such taxes can also have the pro-competitive effect of opening up online markets to alternative business models, such as those based on subscription (like Netflix and several other streaming platforms) or voluntary contributions (like Wikipedia), which are often crowded out because of the ability of ad-based models to grow rapidly.

Fifth, it is important to redirect technology away from an excessive focus on monitoring and surveillance, which is already receiving a big boost from large-scale AI investments in China. U.S. government leadership and funding can again play a central role here, but other regulations are also worth considering. OSHA could be tasked with regulation of surveillance in

⁴³ Acemoglu, Daron, Andrea Manera and Pascual Restrepo (2020) “Does the U.S. Tax Code Favor Automation?” *Brookings Papers on Economic Activity*.

⁴⁴ Lanier, Jaron and E. Glen Weyl (2018) “A Blueprint for a Better Digital Society” *Harvard Business Review*: <https://hbr.org/2018/09/a-blueprint-for-a-better-digital-society>.

⁴⁵ Acemoglu and Johnson (2023a), Chapter 11; Romer, Paul (2021) “Taxing Digital Advertising”: <https://adtax.paulromer.net/>.

workplaces to reduce negative impacts on workers from excessive monitoring, and new laws and regulations can clarify the limits to personal data collection and surveillance on digital platforms.

Finally, with any new technology, new ideas, entrepreneurs, companies and researchers are critical to chart different directions. Creating a more pro-competitive environment, by reducing the dominance of a few firms by means of antitrust and other regulatory measures, may be important as well.