

Written Comments for AI Workforce Forum, November 2, 2023

Daron Acemoglu (Institute Professor, MIT)

Executive summary: Generative AI is potentially transformative, not because it can write Shakespeare-like sonnets, but because it has a unique capability to produce human-complementary, pro-worker tools. If these tools are developed and adopted widely, AI can boost both productivity and wage growth. But we are currently heading toward a very different future, in which AI technologies are focused on automation and surveillance. On this path, AI will increase inequality, limit wage growth and probably fail to deliver on its potential to raise productivity. Government policies can play an important role in redirecting AI in a more pro-worker direction.

Current advances are intensifying concerns about AI replacing jobs and expanding inequality. Indeed, if these tools are used primarily for automation, they will amplify inequalities in the labor market, and if they are monopolized by a few companies, they will likely widen the gap between workers and capital owners.

An alternative, “pro-worker” path is available for AI.¹ This path can improve human decision-making across a range of production tasks, rather than just automate work; reduce inequality, rather than further widening it; and more likely contribute to productivity growth. We need policymakers to be concerned about these issues because, though pro-worker AI is feasible and desirable, we are not currently heading in that direction and are unlikely to benefit from its potential without a course correction.²

The background context of the US labor market today is generally well known. Inequality has surged since 1980. Workers, and especially men with high-school education or less, have experienced significant declines in their real earnings. Even men with a college degree have seen only limited gains. These epochal changes have many causes, ranging from globalization to the transformation of US labor market institutions.

My research indicates that the most important factor has been automation — that is, the substitution of machines and algorithms for tasks previously performed by workers. Although

¹ Acemoglu, Daron, David H. 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 Initiative policy memo: <https://shapingwork.mit.edu/wp-content/uploads/2023/09/Pro-Worker-AI-Policy-Memo.pdf>

² Acemoglu, Daron (2021) *Redesigning AI: Work, Democracy, and Justice in the Age of Automation*, Boston Review Forum.

automation is not a recent phenomenon, it has accelerated over the last four decades with rapid advances in digital technologies.³ My work estimates that automation from industrial robots, other advanced automated machinery, and various software systems accounts for about 50-70% of the increase in wage inequality between different demographic groups in the US labor market. Other trends, such as offshoring, competition from cheaper Chinese imports and changes in labor market institutions and norms, have also played a role, but they have been less important than automation.⁴

While most of this evidence pertains to pre-AI digital technologies, recent work suggests that emerging AI applications are targeting automation. For instance, establishments with jobs and tasks that can be automated by AI were at the forefront of AI adoption in the late 2010s and the same establishments concurrently cut back on their hiring.⁵ 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. (This surveillance path is also getting a boost from China, where these technologies have been a major priority for government investment.)⁶

Job displacement and the inequality costs of new technologies can be counterbalanced by productivity gains, and there is no doubt that some automated machinery, including industrial robots, has boosted productivity. But overall, the evidence suggests that productivity benefits from digital automation technologies and early AI have been somewhat disappointing.⁷ One reason for this may be a failure to use these technologies to make workers more productive, for example, by expanding the set of complex tasks that they perform and improving their contribution to production in more tangible ways.

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,

³ Acemoglu, Daron and Pascual Restrepo (2020) "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, 128(6): 2188-2244.

⁴ Acemoglu, Daron and Pascual Restrepo (2022) "Automation, Tasks and the Rise in US Wage Inequality," *Econometrica*, 90(5): 1973–2016.

⁵ Acemoglu, Daron, David H. Autor, Jonathon Hazell, and Pascual Restrepo. 2022. "AI and Jobs: Evidence from Online Vacancies." *Journal of Labor Economics* 40(S1): S293-S340.

⁶ Acemoglu, Daron and Simon Johnson (2023) *Power and Progress: Our-Thousand-Year Struggle over Technology and Prosperity*, Hachette PublicAffairs.

⁷ Acemoglu, Daron, David Dorn, Gordon H. Hanson, and Brendan Price (2014) "Return of the Solow Paradox? IT, Productivity, and Employment in US Manufacturing." *American Economic Review* 104(5): 394-99; and Gordon, Robert (2016) *The Rise and Fall of American Growth*, Princeton University Press.

inventor of the mouse and hypertext, among other things.⁸ Human-complementary tools like these, 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, '60s and early '70s. Nevertheless, this “machine usefulness” agenda has often been overshadowed by other priorities in today’s tech industry, including automation and monetization models based on data collection and digital ads, that do not offer the same opportunity for productivity gains.

Many of the routine tasks that workers previously performed have already been automated, so a large fraction of current jobs requires problem-solving and decision-making tasks. Empowering workers to perform these tasks more effectively, and to accomplish even more sophisticated decision-making tasks, will necessitate providing workers with better information and decision-support tools. Generative AI is particularly well-suited to this type of information provision. 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 handle complex tasks when empowered with AI tools that present information and recommendations on the basis of the accumulated knowledge from similar cases in the past. In essence, AI holds great potential for training and retraining expert workers, such as educators, medical personnel, software developers, and other “skilled craft workers” such as electricians and plumbers.

These potential pro-worker gains from AI are not hypothetical. Several studies already show that—if used to provide better information to workers or to take over routine tasks in a way that humans with expertise can build upon—generative AI can provide a significant boost to productivity.⁹

If feasible and socially desirable, why isn’t pro-worker AI already the centerpiece of most industry efforts? There are several reasons for this. First, many US corporations are focused on cost-cutting due to the pressures of competition or short-run performance metrics. This often means that increasing the contribution of employees to long-run performance does not receive as much attention as

⁸ Acemoglu, Daron and Simon Johnson (2023) *Power and Progress: Our-Thousand-Year Struggle over Technology and Prosperity*, Hachette PublicAffairs.

⁹ 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 Demirer (2023) “The Impact of AI on Developer Productivity: Evidence from GitHub Copilot.” arXiv Working Paper no. 2302.06590.

it deserves. Second, the tech industry has been dominated by the business models of the largest corporations, which prioritize the development of automation tools and monetization based on digital ads. Third, 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 machine usefulness.

All of these factors push us towards more automation, rather than the pro-worker path. They are also amplified by claims that the next stage of automation itself can be inequality-reducing and projections that generative AI will lead to a huge productivity boom (thus obviating the need to making workers more productive). Both of these claims are unsupported. The idea that automation of well-paid office jobs will reduce inequality is not convincing in light of existing evidence; previous office software systems and early AI have not done this, and even if some high-skill tasks, such as accounting or financial analysis, are automated, workers previously performing these tasks will then compete with less-skilled workers for jobs. This will still transfer some of the burden of automation onto lower-skill Americans (as demonstrated by previous waves of automation). Moreover, while generative AI has tremendous potential, massive productivity benefits from AI-based automation are unlikely for at least two key reasons: first, because these tools can only automate a subset of tasks that humans perform; and second, because some of the tasks that will be automated—especially those involving social skills—are already performed quite productively by existing workers, which limits any opportunity for revolutionary productivity improvements. Even if there were significant productivity gains from automation, these gains would not accrue directly to workers, and building shared prosperity by relying mainly on redistributive policies. For example, higher minimum wages could encourage even more automation in an environment where generative AI is providing additional automation tools.

The direction of technology is never predetermined, and there are many different paths for developing and using almost all technologies. This is doubly true for generative AI, which is a versatile technological platform that can be used in many distinct ways. Government policy can play a role in encouraging a more beneficial trajectory for AI.

A number of policy levers are 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 industries with abundant opportunities for human-complementary work. These include education and training, healthcare, and modern craft work. 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.¹⁰ 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 US 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 of 25–30% for labor, while firms and capital owners pay only about 5% net when investing in equipment and computers.¹¹ A more symmetric tax structure, in which the marginal taxes for hiring and training labor and for investing in equipment and software are equated, can better incentivize firms to implement pro-worker technologies. This tax balance can be achieved by reducing or eliminating payroll taxes, lessening capital income tax deductions, or increasing corporate income taxes.

Third, well-functioning data markets can contribute to the pro-worker technology objective. The current path of large language models is partly shaped by the possibility of using freely available data with few constraints, which discourages investment in higher-quality data. Legislation that clarifies property rights over data and supports a functioning data market could also help by enabling companies to invest in high-quality data.

Fourth, charting a course of new technology that limits excessive surveillance and monitoring is also important, and not only because this would reduce China’s influence over the future of AI. OSHA could be tasked with regulation of surveillance in workplaces, and new legislation should clarify limits to data collection and surveillance on digital platforms.

¹⁰ Gruber, Jonathan and Simon Johnson (2019) *Jump-Starting America: How Breakthrough Science Can Revive Economic Growth and the American Dream*, Hachette PublicAffairs.

¹¹ Acemoglu, Daron, Andrea Manera, and Pascual Restrepo (2020) “Does the U.S. Tax Code Favor Automation?” Brookings Papers on Economic Activity.