

NBER WORKING PAPER SERIES

BUILDING PRO-WORKER ARTIFICIAL INTELLIGENCE

Daron Acemoglu
David Autor
Simon Johnson

Working Paper 34854
<http://www.nber.org/papers/w34854>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2026

This manuscript was written for the Hamilton Project of the Brookings Institution. We thank Julie Gnagy and Julianna Quattrocchi for expert research assistance, Lauren Bauer, Aviva Aron-Dine, Alison Hope, Eileen Powell, and Julia Regier for suggestions that improved the paper, and the Hewlett Foundation, the Smith Richardson Foundation, and the James and Cathleen D. Stone Foundation for research support. Autor also acknowledges research support from the Google Technology and Society Visiting Fellows Program, the NOMIS Foundation, and the Schmidt Sciences AI2050 Fellowship. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w34854>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Daron Acemoglu, David Autor, and Simon Johnson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Building Pro-Worker Artificial Intelligence
Daron Acemoglu, David Autor, and Simon Johnson
NBER Working Paper No. 34854
February 2026
JEL No. J23, J24, J31, M50, O33

ABSTRACT

This paper defines pro-worker technologies, including Artificial Intelligence, as technologies that make human skills and expertise more valuable by expanding worker capabilities. Our conceptual framework distinguishes among five categories of technological change: labor-augmenting, capital-augmenting, automating, expertise-leveling, and new task-creating. Only the last category is unambiguously pro-worker, generating demand for novel human expertise rather than commodifying it. We illustrate these distinctions through hypothetical and real-world examples spanning aviation maintenance, electrical services, custodial work, education, patent examination, and gig delivery. While AI's capacity to automate work is substantial, we argue that its potential to serve as a collaborator, by extending human judgment, enabling new tasks, and accelerating skill acquisition, is equally transformative and currently underexploited. We identify market failures, including misaligned firm and developer incentives, path dependence, and a pervasive pro-automation ideology, that may lead to underinvestment in pro-worker AI. We consider nine policy directions that would change incentives, including targeted investments in health care and education, tax code reform, antitrust enforcement, and intellectual property protections for worker expertise.

Daron Acemoglu
Massachusetts Institute of Technology
Department of Economics
and NBER
daron@mit.edu

Simon Johnson
Massachusetts Institute of Technology
MIT Sloan School of Management
and NBER
sjohnson@mit.edu

David Autor
Massachusetts Institute of Technology
Department of Economics
and NBER
dautor@mit.edu

1 Introduction

The majority of U.S. workers—52 percent of those surveyed in 2024—are worried about how artificial intelligence (AI) will affect their jobs (Lin and Parker, 2025). This is not a simple case of fear of the unknown: 42 percent of workers who currently use AI at work believe that it will reduce their future job opportunities. That number is substantially higher than the 30 percent of AI non-users who express the same opinion.

This pessimism is understandable. AI is frequently hailed as the greatest automation technology in human history. OpenAI, maker of ChatGPT, defines Artificial General Intelligence (AGI) as “highly autonomous systems that outperform humans at most economically valuable work” (OpenAI, n.d.). Achieving AGI is central to the charters of OpenAI and several other leading AI companies.

If, as seems likely, AI will automate many core roles in customer service, marketing, driving, translating foreign languages, coding, medical diagnosis, and myriad creative tasks such as illustrating, animating, and writing, this will devalue the skill sets from which many workers draw their livelihoods—and the pain will be magnified if this automation occurs rapidly.

The consequences could be momentous. Labor market inequality in the United States has already soared over the past four and a half decades. This development has been accompanied by weak—and in some decades negative—real wage growth among noncollege workers and a substantial decline in labor’s share of national income and its share in manufacturing value-added (Acemoglu, 2026). Though the rise in U.S. inequality is extreme, much of the industrialized world has experienced similar trends: mounting earnings inequality, stagnant earnings levels among noncollege workers, and a declining labor share of national income (Acemoglu and Johnson, 2023). There is now widespread concern among policymakers, scholars, technologists, and social commentators that the era of AI will accelerate these worrisome trends.

This paper argues that there is another direction that can push back against these trends, one that uses AI constructively to expand labor market opportunities and increase the value of labor. We refer to this direction as “Pro-Worker AI” (PWAI). We define pro-worker technologies—including AI—as *technologies that make human skills and expertise more valuable by expanding worker capabilities*.

While our definition of pro-worker technologies reads as innocuous—who could not want human expertise to rise in value?—we will explain why this is a demanding definition that many technologies fail to meet. Technologies that do not meet this definition are not necessarily

socially undesirable, but we will explain why they pose additional labor market trade-offs, measured in earnings, inequality, and opportunity. Given the shared goal of many technologists, policymakers, worker representatives, and the public to identify, foster, and invest in technologies that are beneficial for workers and society at large, we see an urgent need to define pro-worker AI and figure out how to advance its development.

In this essay, we explain how pro-worker AI could affect jobs, wages, and the value of worker expertise, why pro-worker AI is technically feasible and socially desirable, why the private market is likely to under-invest in pro-worker AI, and how specific public and private policies might alter investment incentives.

2 AI's pro-worker potential

AI's capacity to automate work and displace workers is substantial. Simultaneously, we believe that AI has equally transformative potential to act as a force-multiplier for human skills and expertise. This potential arises from AI's capacity to collaborate with workers, enabling them to (a) perform more-sophisticated tasks related to their jobs, (b) tackle new tasks, and (c) acquire new expertise. This multimodal collaborative capacity is what gives AI its potency as a pro-worker technology.

Where does this collaborative capacity come from? Prior to the advent of modern AI, traditional computing's superpower was its ability to follow clearly-specified rules to carry out routine tasks with speed and accuracy. This made traditional computing almost perfectly suited to automation of routine tasks that could be reduced to the mechanistic performance of the same prespecified steps in a stable environment ([Autor, Levy, Murnane, 2003](#)).

But five decades into the computer revolution, most human work is not mechanistic. Instead, it requires discretion, improvisation, and judgment—capabilities that are required precisely when prespecified rules are not enough. There is no single best way, for example, to care for a cancer patient, write a legal brief, remodel a kitchen, or develop a lesson plan. But the skill, judgment, and ingenuity of human decision-making determine outcomes in many of these tasks, sometimes dramatically so. Making the right call means exercising judgment.

This is where AI could be highly effective. Rather than relying on hard-coded procedures, AI learns by example, synthesizes unstructured information, and acquires capabilities that it was not explicitly engineered to possess. Like a human expert, AI can weave formal knowledge (rules) with acquired experience to make—or support—one-off, high-stakes decisions. This capacity to master knowledge from unstructured data, and to improvise based

on this acquired knowledge, makes AI an invaluable tool for advising, coaching, supporting, and enhancing the work of decision-makers—and all workers are decision-makers. Whereas traditional computing stopped at the water’s edge when workers needed to swim in the uncertain waters of judgment, discretion, and improvisation, AI can wade into those waters with them to advise, support, coach, and enhance their work.

Human-machine collaboration is present when machine capabilities are a force multiplier for human expertise and is distinct from the much more common vision of using machines to automate human capabilities and hence make human expertise superfluous (Aitor and Manyika, 2025). Rather than seeking to automate expertise entirely, delegate decisions wholly to machines, or assume either humans or AI will invariably excel, the objective of collaboration is to integrate human judgment with rapidly advancing computational capabilities to achieve superior outcomes.

The idea of human-machine collaboration has deep roots in the history of computing. In an iconic paper, MIT computer scientist Licklider (1960) wrote, “The hope is that, in not too many years, human brains and computing machines will be coupled together very tightly, and that the resulting partnership will think as no human brain has ever thought and process data in a way not approached by the information-handling machines we know today.”

This vision was a far-fetched aspiration six decades ago. But contemporary AI has changed this equation. A modern AI system can ingest drone imagery and soil sensor data from a farm’s every acre, the complete sensor logs from a building’s HVAC system, or the detailed vital signs of a single patient observed over many months to support workers making high-stakes decisions. Drawing on this pretraining, AI tools can think alongside workers, identify relevant context, generate well-informed responses to questions, and present lucid, well-structured data to support decision-making. This is collaboration.

You may object: “If AI can behave like an expert, can’t it simply replace experts, thus automating their expertise into irrelevance?” In some cases, the answer is yes. But in many more cases, we think the answer is no: AI will prove more effective at collaboration than at automation. Precisely because AI is not rule-bound, it is less trustworthy as an autonomous actor than a conventional computer system, and more valuable as a collaborator (Narayanan and Kapoor, 2025).

To be useful, an automation tool must deliver near-flawless performance almost all the time. You would not tolerate a spreadsheet that hallucinated values, a robotic surgeon that glitched-out during bypass surgery, an agentic investing tool that squandered your money while you were not paying attention, or an AI-powered vending machine that gave away

PlayStations and stocked live fish at the behest of persuasive customers (Stern, 2025). For most of these tasks, the stakes are too consequential and the decisions too nuanced to be fully delegated to an automatic system that acts on its own discretion. The AI needs human expertise.

A collaboration tool does not need to be anywhere close to infallible to be useful. A doctor with a stethoscope can better diagnose a patient than the same doctor without one, and a contractor can pitch a squarer house frame with a laser level than they could by eyeballing it. These tools do not need to work flawlessly, because they do not promise to replace the expertise of their user. They make experts better at what they do—and extend their expertise to places it could not go unassisted. Rather than making expertise unnecessary, they render expertise more valuable by extending its efficacy and scope. It is this complementarity between machine capacity and human expertise that we believe imbues AI with vast pro-worker potential.

To be clear, automation is not immoral or intrinsically undesirable. It has been with us since the beginning of the Industrial Revolution and will continue to be an important part of what technologies do. In the recent past, extensive expertise was required to drive an automobile, manage an investment portfolio, or measure one’s blood glucose levels (necessary for diabetics), making these activities inconvenient, expensive, or simply infeasible for many. Automation now makes these tasks inexpensive, near effortless, and broadly available.

Looking forward, many tasks that do not currently require human judgment, dexterity, or contextual understanding will be ripe targets for automation. Because the core tasks of educators, health-care workers, crafts workers and many others do not fit this description, those tasks are ill-suited for full automation, at least at present. Workers performing these jobs, however, can benefit substantially from better technology. The opportunity that AI offers, as we argue below, is not primarily to automate this work but rather to serve as a collaborator for the workers who perform these vital roles.

The next section explains what makes a technology pro-worker by laying out the causal channels by which new technologies affect the labor market. Applying this framework, we discuss several promising examples of pro-worker AI that are already in the field. However, we also underscore that pro-worker AI is not the norm, and we explain the reasons why it is not. The first is incentives: leading firms perceive greater economic return to building and deploying technologies that automate expertise than those that create new tasks or workers and increase the value of skills and expertise. The second reason is ideology. Even though it is a technical and highly quantitative field, computer science—and the AI community in general—is nonetheless gripped by an ideological vision that places AGI, meaning machines

that exceed all human capabilities, as its highest possible pursuit. Since we are economists and not philosophers, we have little to say about how to change the ideology of a scientific field. But we have a lot to say about how to steer AI development and deployment in a pro-worker direction.¹

3 What makes a technology pro-worker?

As a prelude to our formal discussion, we offer a stylized example of how different approaches to using AI—or potentially any technology—may have sharply different implications for human skills and expertise, even when used for the same goals.

Our example concerns the occupation of Aircraft and Avionics Equipment Mechanics and Technicians (AMTs). In 2024 there were 161,000 AMTs in the United States with median pay of \$79,000 (Bureau of Labor Statistics, 2025a). In the same year, Delta Airlines’ technical operations division alone employed 9,600 AMTs at 51 maintenance stations worldwide (Wikipedia, 2025). Becoming an AMT requires specific vocational training and federal certification. Most AMTs get their start either in the military or by attending a Federal Aviation Administration (FAA)-approved aviation maintenance technician program. Completing an approved program qualifies AMTs to sit for FAA exams, which they must pass to earn certification. Non-certified AMTs work under the supervision of a certified Airframe and Powerplant Mechanic until they have accrued sufficient experience to sit for the FAA exams (Federal Aviation Administration, n.d.a).

Airlines have a strong incentive to run their inspection, maintenance, and repair operations as reliably and cost-effectively as possible. We next present three hypothetical AI tools that airlines might adopt to support those objectives and outline how each would differ in its implications:

1. **Aviation Maintenance Automator (AMA).** This AI tool purports to make aircraft maintenance so easy that “any kid with a screwdriver” can do it. Outfitted with hard-hat-mounted cameras, microphones, lidar, and thermal sensing gear, the AMA guides even uncertified technicians to inspect and maintain aircraft. Although the long-term goal of AMA developers is to fully replace workers with dexterous robots that perform last-mile physical tasks, technicians are presently still needed. Fortunately, these workers will require no specialized training, just proficiency with basic

¹ We recognize that the terms “socially desirable” and “pro-worker” are not in all cases synonymous. We briefly explain below why pro-worker AI has many socially desirable implications.

mechanics' tools and a willingness to follow directions diligently.

2. **AMT Assistant.** This tool is also designed to assist technicians who maintain and troubleshoot commercial planes. Different from the AMA Automator, the technician working with the AMT Assistant is more than a pair of hands; experience and training are also required. But the AMT Assistant extends the technician's capabilities. From the aircraft bay, a technician can upload diagnostic data, photographs, and field notes. Analyzing this information and asking supplemental questions, the AMT Assistant makes recommendations on what tests to run and which procedures to implement. The software also advises on regulatory compliance, completes required paperwork, and provides a real-time forecast on when the plane will again be airworthy. The AMT Assistant is pre-trained using basic engineering know-how, technical information on all relevant aircraft types, and a vast set of field cases encountered by experienced technicians. Additionally, the AMT Assistant can assist with supplemental training. By generating realistic use cases that technicians can tackle interactively, the AMT Assistant functions like a flight simulator for technicians. Field trials have demonstrated that the AMT Assistant enables technicians to make better diagnoses and perform repairs that are more cost-effective. It also effectively supports skill acquisition: technicians using the AMT Assistant attain expert level performance more rapidly than those using only conventional tools and training resources.
3. **AMT Liftoff.** This tool, a successor to the AMT Assistant, assists Airplane Maintenance Technicians to transition their skill sets to the rapidly growing spaceflight sector. This sector includes firms such as SpaceX, United Launch Alliance (Boeing Company and the Lockheed Martin Corporation), Blue Origin, and Rocket Lab USA. Inspection, maintenance, and repair of spacecraft is a rapidly evolving, technically demanding, and mission-critical line of work, so new expertise is badly needed. Airplane Maintenance Technicians (including those who have worked with the AMT Assistant) have an applicable foundational skill set to build on. Using flight simulator-like features, AMT Liftoff provides a virtual hands-on learning environment where trainees can practice maintenance, diagnosis, and repair tasks on different spacecraft. Although AMT Liftoff's virtual learning environment is not a complete substitute for direct supervision and field experience, the tool accelerates the learning process, reduces costs for the trainee, and substantially shortens the on-ramp for AMTs seeking to enter the spaceflight sector. As technicians move from the virtual environment to the field, AMT Liftoff provides real-time support, much like the AMT Assistant, to support problem-solving, to ensure quality, and to accelerate learning.

Each of these hypothetical technologies helps with the critical task of inspecting, maintaining, and repairing aerospace craft. But they do so differently. With the first technology, the Aviation Maintenance Automator’s ultimate objective is to de-skill the AMT job. Technicians serve as little more than eyes and hands for the AMA, which directs their work. While this simplification might sound like good news for technicians—after all, they will need less training to do the work—it also inevitably means lower earnings. If, going forward, basic mechanical proficiency is the sole requirement for this job, airlines would be foolish to pay AMT wages for skills that they no longer need. Even if, hypothetically, dexterous robots proved unable to perform manual maintenance tasks, and, moreover, airlines began to hire more maintenance technicians considering their higher productivity, the pay of AMTs would decline. In effect, the AMA commodifies AMT’s scarce expertise, undercutting the market value of that skill set by providing a cheap machine substitute.

The AMT Assistant presents a sharp contrast. Rather than commodifying AMTs’ expertise, this tool extends that knowledge base by enabling AMTs to do more-difficult diagnostic and repair work, plausibly with higher quality and greater efficiency. By advising on regulatory compliance and completing paperwork, the AMT Assistant also enables AMTs to focus on core competencies. While this tool will be particularly useful for junior AMTs because it enables them to perform the sophisticated tasks typically reserved for technicians who are much more seasoned, it also enables all technicians to perform new tasks, such as dealing with rare problems and complex troubleshooting activities.

Underscoring that the details of how technology is designed, developed, and used matter, it is possible that the AMT Assistant may not be unalloyed good news for experienced technicians who might now face greater competition from novice AMTs in some of the tasks that were previously reserved for them. This concern is not frivolous, and highlights that there is a difference between tools that enable workers to branch into completely new tasks and those that enable the workers to perform some of the more sophisticated tasks previously assigned to workers with greater experience or expertise.

AMT Liftoff has much in common with the AMT Assistant, but with one important difference: it enables workers to expand their expertise into truly new tasks by mastering a greenfield where there is both rising demand and a rapidly expanding set of expert skills to be mastered. Akin to the AMT Assistant, this tool enables workers to acquire and deploy new expertise to branch into genuinely novel tasks, as opposed to tasks that are new to these workers but are already being performed by workers who are more expert.

Few technologies create only winners or losers, as these examples underscore. The first technology devalues human expertise, but also presumably reduces costs for airlines and hence,

ultimately, customers. The second technology extends worker expertise but, depending on how it is used, can also increase the supply of expertise. That has clear benefits for airlines and their customers, but could come at a competitive cost to existing technicians. The third technology benefits entrants, firms, and customers. And because there are essentially no workers currently performing these tasks, this technology does not come at a competitive cost to other workers.

That makes it an unusual case, however. In a few years' time, today's trainees will be tomorrow's incumbents. At that point, they may be unhappy that AMT Liftoff II makes it even faster and easier for new entrants to become highly capable aerospace technicians. This concern worries us less. While almost everyone loves free markets in theory, most people dislike competition in practice. This is a truth that pro-worker AI cannot resolve.

With this concrete example in mind, we next outline formally how new technologies interact with human skills and expertise, then circle back to their practical applications.

4 Our conceptual framework for pro-worker AI

We now introduce our conceptual framework and more formally define features of pro-worker technologies and how they differ from other types of technological advances.

4.1 The value of expertise

To understand how technologies change the values of skills and expertise, we should begin by considering what gives them labor market value. Although many forces affect wage setting—such as minimum wage laws, collective bargaining regimes, and social norms—workers' pay in market economies is determined, to a significant extent, by their capabilities in performing tasks that produce goods or services that themselves have market value.² A task is a notional concept—one can define tasks broadly or narrowly—but it is useful to think of tasks as the building blocks of work. Most tasks require specific capabilities or know-how. We will refer to that know-how as expertise. Specific task expertise might be specialized, highly trained, and time-consuming to master (e.g., heart surgery), or it might be commonplace or easily mastered (e.g., shelf stocking).

Although the words “skills” and “expertise” are often used interchangeably, they have distinct meanings here. Skills are broad capabilities, such as analytical reasoning, leadership,

² We speak here only of market value, not of intrinsic, spiritual, or ethical value.

dexterity, creativity, and so on. Because they are useful across many domains, broad skills tend to have robust and durable market value—although of course that value can evolve over time (e.g., physical strength in the U.S. is probably less valuable than it used to be). By contrast, many job tasks—typing, coding, welding, baking, translating, sewing, navigating, and so on—require specific, formal expertise that is applicable primarily within a narrow domain.

Reflecting this specificity, the market value of any particular form of expertise may be altered rapidly by technological changes, either because the expertise is rendered obsolete or because it is made newly relevant and applicable. For example, as AI-based language translation tools have become proficient and ubiquitous, employment of professional language translators in the United States has modestly declined and is expected to fall further (Abril, 2025). While mastery of foreign languages is a relatively broad skill that will likely remain useful for the foreseeable future, the labor market value of language translation expertise is clearly threatened.³

The occupation of data scientist provides an example where expertise is made more valuable by advances in digital technology. Data science is a form of novel, specialized expertise that was essentially brought into being by the advent of big data. The title of “data scientist” was first coined in 2008 in an article in the *Harvard Business Review* (Davenport and Patil, 2012); it was not formerly recognized by U.S. statistical agencies until 2018 (Bureau of Labor Statistics, 2018, occupation 15-2050).⁴ Yet, in 2024, there were 246,000 data scientists working in the United States, with median annual earnings of \$113,000 (Bureau of Labor Statistics, 2025b).⁵

Not all expertise is equally valuable, as these examples suggest. To possess significant labor market value, expertise must have two characteristics: (1) it must be required for some economically valuable activity, and (2) it must be scarce—meaning not abundantly available or readily produced. To see why scarcity is important, consider the related occupations of school crossing guards or flaggers (people who control traffic around construction sites) and air traffic controllers. Although the safety of schoolchildren and other pedestrians depends on them, crossing guards and flaggers are low-paid workers, typically earning around \$40,000 annually (Bureau of Labor Statistics, 2024). Conversely, air traffic controllers typically earn

³ It is likely that the most expert translators will still be in demand for high-stakes applications such as diplomatic communications, book translations, etc. This is a common pattern wherein automation eliminates the inexpert components of a job, leaving behind a subset of jobs that are more specialized, better paid, and less numerous (Autor and Thompson, 2025b).

⁴ See occupation 15-2050 in the Standard Occupational Classification of 2018.

⁵ See occupation 15-2050 in the Standard Occupational Classification of 2018.

about \$145,000 annually, three-and-a-half times what crossing guards do, even though they are effectively crossing guards for the sky (Bureau of Labor Statistics, 2025b). What explains this pay differential? The principal answer is the relative scarcity of expertise. Almost any able-bodied adult can become a crossing guard or flagger with one to three days of training. Becoming a certified air traffic controller, however, requires two to four years of college-level classroom training, followed by thousands of hours of apprenticeship (Federal Aviation Administration, n.d.b). This difference in expertise makes air traffic controllers hard to replace. If elementary schools suddenly ran short of crossing guards, they could readily train teachers, parents, retirees, or even air traffic controllers to do the work. If airports ran short of air traffic controllers, however, crossing guards could not fill in.

4.2 How technologies reshape the value of skills and expertise

We defined pro-worker technologies as those that expand human capabilities in a way that makes human skills and expertise more valuable. Here we consider how new technologies—including machines, algorithms, or other innovations—affect these valuations. It is helpful to group technologies into five broad, economically distinct categories, following and extending Acemoglu, Kong, Restrepo (2025): (1) labor-augmenting technologies, (2) capital-augmenting technologies, (3) automating technologies, (4) expertise-leveling technologies that enable workers to perform more sophisticated tasks currently done by others; and (5) new task-creating technologies.⁶

As we explain below, only one of these five categories, new task-creating technologies, is unambiguously pro-worker, meaning that it creates only winners. Closely related are expertise-leveling technologies that enable less expert workers to perform tasks that previously demanded expertise from other domains. While these technologies have many potential benefits, they also create both winners and losers in the labor market by opening new opportunities for some and intensifying competition for others. For this reason, we do not classify them as unambiguously pro-worker.

Although these five categories are conceptually distinct, we note that most real-world technologies do not fall neatly into one category but rather combine elements from several. We discuss these nuances next. Table 1 summarizes our discussion.

⁶ Our discussion of expertise here and throughout the paper integrates insights from Autor and Thompson (2025b).

Table 1: Types of technologies and their labor market consequences

	Example technology	Labor productivity	Value of human expertise	Labor's share of national income	Pro-Worker?
Labor-augmenting technologies	Electric cable stripper replaces hand stripper	+ More output per hour	\leq Expertise more relevant/useful, but higher output could lower price	≈ 0 No task reallocation	Ambiguous
Capital-augmenting technologies	Lighter, faster electric cable stripper	+ More output per hour	\leq Expertise more relevant/useful, but higher output could lower price	≈ 0 No task reallocation	Ambiguous
Automation technologies	Cable installing robot	+ More output per hour	- Existing expertise made obsolete	- Fewer tasks done by labor, more by capital	Not pro-worker
New task-creating technologies	Ethernet, fiber optics, occupancy sensing	+ More output per hour	+ New expertise needed	+ More tasks done by labor, fewer by capital	Unambiguously pro-worker
Expertise-leveling technologies	Blood oximeter	+	\leq Entrants benefit. Incumbent expertise potentially devalued	\leq Ambiguous due to offsetting effects	Ambiguous

4.2.1 Labor-Augmenting Technologies

We define labor-augmenting technologies as those that make workers more effective at the tasks they already do.⁷ For example, an electric cable stripper is a labor-augmenting tool: an electrician who swaps their manual wire-stripper for an electric cable stripper will be able to complete electrical installations faster.

Expertise, productivity, and wages Since they neither automate existing tasks nor expand what this group of workers does, labor-augmenting technologies do not directly in-

⁷ More generally, labor-augmenting technologies could allow groups to become more productive across a range of tasks, including those that they are not currently performing, enabling them to compete against others for marginal tasks and increasing their task shares. See [Acemoglu, Kong, Restrepo \(2025\)](#).

crease the value of expertise. The net effect on worker earnings is accordingly ambiguous, depending on how prices adjust as worker output rises. Although higher productivity tends to lower prices (benefiting consumers), raise wages (benefiting workers), or yield some combination of the two, labor-augmenting technologies are not guaranteed to boost labor demand or earnings.

Labor share Labor-augmenting technologies—which increase the productivity of workers in the tasks they were already performing—do not, to a first approximation, change the division of tasks between workers and machines, and, as a result, have only small effects on the labor share of value-added.⁸

4.2.2 Capital-Augmenting Technologies

What labor-augmenting technologies do for workers, capital-augmenting technologies do for machines (e.g., algorithms, processes, innovations): they make them better, cheaper, or faster at performing their current tasks. For our purposes, labor- and capital-augmenting technologies are not fundamentally distinct. Replacing an electrician’s hand pliers with an electric cable stripper (a labor-augmenting change) or upgrading that cable stripper with a better model (a capital-augmenting change) both make electricians more productive.

Productivity and wages A new machine that costs the same as its predecessor but works twice as rapidly will necessarily raise productivity and typically lower prices, both of which create aggregate benefits. Because of this price effect, however, the net impact on worker earnings is ambiguous. Concretely, if the price of cable installation falls faster than output rises—which is perfectly possible—cable installers will see their wages fall.⁹ Because of this ambiguity, we do not classify capital-augmenting technologies as unambiguously pro-worker.

⁸ We focus here on narrow labor-augmenting and narrow capital-augmenting technologies. If capital-augmenting technologies made machines and algorithms more productive in all tasks, then they would have a first-order effect on the allocation of tasks to workers—somewhat akin to automation technologies, discussed next.

⁹ We note that, although wages of workers who are directly affected by capital-augmenting advances can fall (i.e., if prices fall faster than output rises), labor can benefit in aggregate, especially if there is enough competition for labor. This net benefit occurs through a general equilibrium effect: an increase in the effective supply of capital (due to capital accumulation can raise total payments to labor because capital and labor are productive complements. Similarly, labor-augmenting technologies necessarily raise payments to capital, even though labor may or may not benefit. These subtleties are discussed in [Acemoglu, Kong, Restrepo \(2025\)](#) and in [Jones and Liu \(2024\)](#).

Expertise and labor share As with labor-augmenting technologies, pure capital-augmenting technologies neither make existing forms of expertise redundant through automation nor make new forms of expertise relevant through new task creation. Because they do not fundamentally alter the division of tasks between capital and labor, these technologies have generally neutral effects on the labor share—they expand the economic pie without shifting the relative size of the slices paid to capital and labor.

4.2.3 Automation Technologies

Distinct from labor- or capital-augmenting technologies, automation technologies directly reshape the division of labor between workers and machines by substituting machinery or algorithms for tasks that were previously performed by workers (see [Acemoglu and Autor, 2011](#); [Acemoglu and Restrepo, 2018, 2019](#); [Autor, Levy, Murnane, 2003](#)). Consider, for example, the introduction of a dexterous robot that autonomously installs electrical cabling at commercial construction sites. Rather than augmenting workers (or machines) in the tasks that they already perform, this machine would primarily take over tasks currently performed by workers.

Productivity, expertise, and wages Automation raises productivity by substituting cheaper machines and algorithms for human labor. Nevertheless, it also devalues expertise and could undercut wages.¹⁰ The first channel for its potential negative effect on wages is via a reduction in labor demand: even if demand for cable installation soars as robotic automation lowers prices, electricians will not be needed to do this work anymore.

The second channel is the commodification of expertise. By replacing a specialized skill set with an inexpensive automated process, automation technologies reduce the value of expertise. If, for example, prior to robotics, the installation of electrical cabling required specialized expertise and certification, then automating that capability would diminish or eliminate the economic value of that expertise.¹¹ Concretely, no matter what job cable installers do next, their expertise in cable installation will have little relevance or economic value. Conversely, if cable installation were an inexpert task, readily mastered in only a few hours, then little expertise would be commodified if this task were automated. It therefore matters not just whether automation displaces tasks, but also which tasks it displaces ([Autor](#)

¹⁰ Automation can also be used *unproductively* to reduce worker bargaining power. This raises firm profits at the expense of workers' earnings without any positive effect on output or efficiency ([Acemoglu and Restrepo, 2026](#)).

¹¹ [Acemoglu and Restrepo \(2020\)](#) document that the adoption of industrial robots reduced the earnings of workers specializing in blue-collar tasks in the labor markets where robot adoption was concentrated.

and Thompson, 2025b). While all automation is labor saving, automation that displaces workers from tasks that demand specialized expertise has an additional cost: it renders that specific stock of human expertise obsolete.

This distinction has practical relevance. Workers instinctively understand the difference between expertise-displacing (commodifying) automation and merely labor-saving automation. Hollywood screenwriters have experienced two major waves of automation over the past five decades: traditional computing (e.g., word processing and editing tools) and, more recently, AI. The first technology was welcomed, the second resisted, including with a prolonged Writers’ Guild strike in 2023. Given that word processors and AI are automation technologies, one might naïvely guess that screenwriters would have resisted both. But from the screenwriters’ perspective, they are fundamentally different: word processors assist screenwriters to apply their screenwriting expertise more efficiently—they are merely labor-saving—whereas AI threatens to turn that expertise into a low-cost commodity. Ultimately, screenwriters are paid for their expertise in crafting stories—precisely what AI threatens to undermine—and not simply for the hours they spend writing.¹²

We again underscore that automation has costs and benefits. Productive automation technologies will typically lower prices or raise wages (for those who remain employed), or both. Even so, automation that displaces the expertise from which workers draw their livelihoods almost always has concentrated costs on those whose expertise is devalued, even while creating beneficiaries elsewhere in the economy.

Labor share This ambiguous impact on wages notwithstanding, by reallocating tasks from workers to machines, automation unambiguously reduces labor’s share of value-added in the sectors in question. This reallocation effect is generally a minor concern at the level of any single activity (e.g., electrical installations), but is broadly consequential when the economywide labor share falls, as has happened in the United States and many other countries over the past two decades (Karabarbounis, 2024).

4.2.4 New Task-Creating Technologies

The first three technology categories do not meet our definition of pro-worker. That is because, in their simplest form, they neither expand human capabilities nor necessarily make existing human skills and expertise more valuable. By contrast, new task-creating technologies do both. Precisely opposite to automation technologies, these technologies

¹² We thank Anna Salomons of Tilburg University for this example.

generate demand for novel human expertise.

Returning to our example of electrical installation: innovations such as ethernet networks, fiber optic cabling, and occupancy-aware heating and lighting systems have vastly increased the quantity and complexity of electrical cabling required in modern buildings. Befitting this complexity, workers now require specialized expertise to plan, install, and maintain these systems—expertise that was neither relevant nor prevalent prior to the introduction of these technologies.

Productivity, expertise, and wages New tasks increase productivity by expanding what humans can productively do, and they do so specifically by creating forms of expertise. To see why this matters, it is important to distinguish between new work and more work (Acemoglu and Restrepo, 2018, 2019; Autor, Chin, Salomons, Seegmiller, 2025). Many economic forces can create more work, meaning a larger quantity of work that workers do already. What makes new task-creating technologies distinctive is that they add to work that demands new expertise (i.e., new work) and reinstate workers into jobs that value expertise. Over the course of decades, the flow of new work replenishes the demand for expertise that automation tends to otherwise erode. Quantitatively, this effect is large. Recent work estimates that more than six out of ten workers in 2018 were employed in occupational specialties that did not yet exist in 1940 (Autor, Chin, Salomons, Seegmiller, 2024).

Like the other technology categories discussed above, new task-creating technologies also raise productivity, in this case by satisfying new demands rather than by doing the same things more cheaply. Distinct from labor- and capital-augmenting technologies, they necessarily increase both the variety and the quantity of work demanded. Distinct from automation technologies, they tend to increase the value of the expertise required to perform these new tasks.¹³

Of course, not all new expertise is of equal economic value. Between 2000 and 2024, Amazon.com created more than 1.5 million new jobs, with the vast majority in logistics—meaning warehousing and delivery (Prakash, 2025). Although this work was arguably “new,” giving rise to distinctive occupations (e.g., Amazon warehouse picker), the required expertise was either already abundant or readily acquired. Thus, this work was closer to more work than to new work, and its pay was accordingly low. Ironically, many of these same Amazon logistics jobs are now themselves threatened by automation—perhaps because their low expertise levels made them susceptible to the next wave of warehouse automation (Weise, 2025).

¹³ Increasing the quantity of work—of all types—tends to increase wages by boosting labor demand. In addition to boosting labor demand, new tasks imbue new forms of expertise with market value.

Labor share and expertise Like automation technologies, new task-creating technologies also alter the division of labor. But they shift it in the opposite direction from automation, expanding the set of tasks performed by workers rather than by machines, thus raising the labor share of value-added. By increasing the quantity of workers demanded, making novel forms of human expertise valuable, and increasing labor’s share of economic activity, these technologies are unambiguously pro-worker: they boost the value of worker skills and expertise by expanding human capabilities.

As our discussion below highlights, new expertise does not necessarily mean more high-tech jobs requiring higher degrees. Expertise is just as relevant in blue-collar work, in health-care delivery, and in personal services, among many other domains, as it is in typical college-degreed white-collar jobs. In fact, we suspect that pro-worker AI has even more potential in these non-degreed jobs than it does in conventional white-collar work. Our examples below will make clear why this is so.

4.2.5 Expertise-Leveling Technologies

Closely related to but distinct from technologies that create new tasks are those that enable a new set of workers to perform tasks that previously demanded expertise from another domain(s). For example, a medical technician equipped with a pulse oximeter can quickly read a patient’s blood oxygen levels. Prior to availability of this tool, the same task required a phlebotomist to draw blood, a lab technician to analyze that blood, and a doctor or nurse to interpret the results. Nevertheless, this technology did not create a new task: it simply made medical technicians newly able to perform the existing task.

Productivity, expertise, wages, and labor share Expertise-leveling technologies typically raise productivity by reducing costs. While firms almost always benefit from lower costs, the effects on workers of expertise-leveling technologies are usually mixed. By expanding the effective capabilities of less-expert (typically less-well-paid) workers, expertise-leveling technologies frequently make these workers’ expertise more valuable and hence boost their earning power.

By the same token, expertise-leveling technologies typically reduce the scarcity value of those who were previously performing these tasks—unless these latter workers can now themselves allocate their time toward tasks that require more expertise (e.g., medical doctors can focus on diagnosing more-complex cases). Given these offsetting effects on different worker groups, expertise-leveling technologies have an ambiguous effect on the labor share. Overall, whether

these technologies are pro-worker or not will depend across contexts.

4.3 Why the direction of technological change matters

The direction of technological change—specifically the balance between automation and new task creation—matters not only for wages but also for its far-reaching social consequences. Over the past five decades, the United States has experienced rising inequality along two dimensions: widening wage gaps among workers across education levels and demographic groups, and a growing divide between workers and capital owners (Acemoglu and Restrepo, 2022; Piketty, Saez, Zucman, 2018).

Automation exacerbates both trends. Automation increases labor market inequality by displacing specific groups of workers from their tasks and jobs—robots replacing skilled blue-collar workers in painting, welding, and sorting, for example. And automation increases overall inequality by raising the share of value-added accruing to capital. Because capital ownership is distributed far more unequally than labor income, any force that expands the capital share also widens the gap between households at the top and everyone else. Rising concentration of wealth, in turn, poses steep challenges for the quality and stability of democratic governance (Acemoglu and Robinson, 2019).

Meanwhile, sharp falls in employment and wages—whether driven by technology, trade, or other economic disruptions—carry social consequences that extend well beyond the labor market. Job loss, especially when concentrated within communities, contributes to premature mortality, family dissolution, and substance abuse (Autor, Dorn, Hanson, 2019; Black, McKinnish, Sanders, 2003; Case and Deaton, 2022; Sullivan and von Wachter, 2009; Wilson, 2011).

In contrast, technologies that create new tasks are likely to ameliorate these social challenges by expanding employment and increasing wages. In short, while all five forms of technological change discussed above are likely to raise productivity, new task-creating technologies have additional virtues for workers and for labor markets broadly. And, of course, pro-worker AI, as we define it, is a new task-creating technology.

4.4 Automation isn't everything

The technology categories detailed above have all played a prominent role in recent and distant history. And yet automation stands out. Many of the breakthrough technologies of the early British Industrial Revolution were automation technologies, first in spinning and

then in weaving. Office computers, software systems, computer numerical control machinery, and industrial robots have been at the forefront of the technological advances of the past six decades. These technologies have many uses, but one thing they have done is to automate core job tasks that were previously carried out by workers on factory floors and in offices.

The role of automation is especially evident in manufacturing. Labor’s share of value-added in U.S. manufacturing fell from 74 percent to 46 percent between 1981 and 2016, while the share of labor in U.S. national income also fell, though less precipitously, from 58 percent to 52 percent (Acemoglu, 2026). Automation is associated with declines in the labor share in value-added within manufacturing and across the entire U.S. economy; the international evidence is consistent with this pattern (see Acemoglu and Restrepo, 2022; Restrepo, 2024; Acemoglu, Kong, Restrepo, 2025). Studies have also shown how automation has displaced those working in routine tasks that can be taken over by machines (Acemoglu and Restrepo, 2020; Autor, Levy, Murnane, 2003).

But automation isn’t everything.

Given the rapid march of automation—from manufacturing, to commerce, to finance—one might have anticipated a secular downward trend in labor’s share from the outset of the First Industrial Revolution in the late 1700s to the present. This is not what has happened. Instead, the share of labor increased during the first eight decades of the 20th century in the United States and similar places, though it has fallen over the past four decades (Abramovitz and David, 2000; Autor et al., 2020; Acemoglu, 2026; Budd, 1960).

Adding to the paradox, rich industrialized countries have substantially higher labor shares than low-income, developing countries, despite their more-extensive automation and greater capital wealth (see Figure 1 from Our World in Data, 2020). If more-advanced technology and greater automation inevitably meant a lower labor share, then we would expect the opposite, with a lower labor share in richer countries and a higher labor share in poorer countries.

Why hasn’t automation driven labor’s share relentlessly toward zero?

One potential explanation is that automation makes workers so much more productive that rising wages offset automation’s labor-share eroding effect. This is unlikely: while automation does not necessarily lower wages, it lowers the labor share. During the Industrial Revolution, for example, major technological advances in cotton spinning, such as the spinning jenny and the water frame, vastly increased labor productivity in yarn production: fewer workers using more machines could produce more spun cotton in less time. Higher productivity meant lower prices, which in turn increased the quantity of weavers needed (to use the

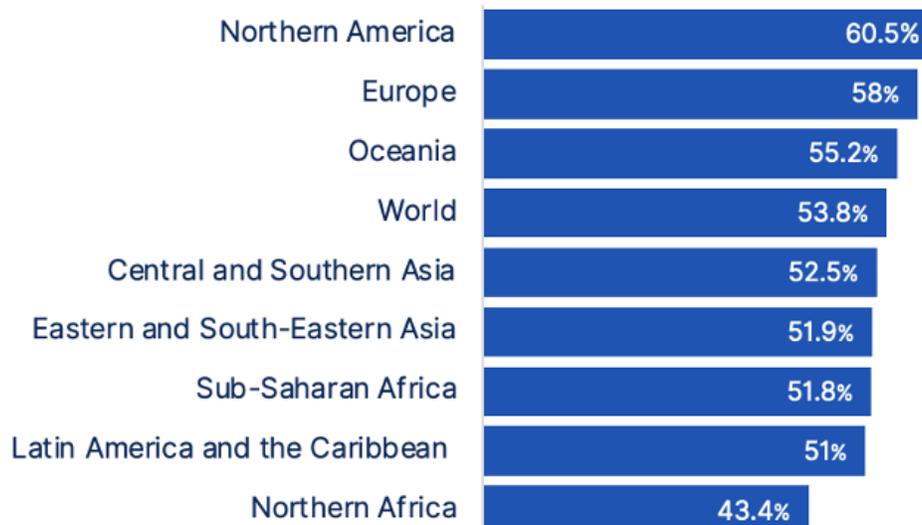


Figure 1: Labor’s share of Gross Domestic Product was larger in higher-income countries in 2020

Source: [Our World in Data](#) (2020).

increased output of spun cotton). But this alone likely did not restore labor share since the division of labor in textile production had shifted decisively toward machines.

In addition, technological advances ultimately proved economically wrenching for weavers also. As economic historian Joel Mokyr and colleagues wrote in 2015, “the handloom weavers and frame knitters with their little workshops were quite rapidly wiped out by factories after 1815” (Mokyr, Vickers, Ziebarth, 2015). Earnings of weavers fell by a striking two-thirds in the early decades of the 19th century (Acemoglu and Johnson, 2024). Instead of skilled weavers, textile mills primarily hired indentured children and unmarried women to perform inexpert, low-paid work under polluted and often gruesomely hazardous conditions—the “dark Satanic Mills” referenced in Blake’s (1804) poem.

Why didn’t vastly higher productivity translate into higher wages and better working conditions? It might have, had textile factories required scarce, specialized expertise that only skilled spinners or weavers could supply. But the opposite was true: dexterous fingers and a high tolerance for discomfort became the only skill requirements, and the poor, and even children, supplied these inexpert capabilities in abundance.

A second explanation for why automation has not further eroded the role of labor in industrialized economies is that the indirect effects of automation offset its direct effects. For example, while industrialization produced inexpert, low-paid jobs for some workers, it indirectly created new demand in other expertise-intensive activities, such as machine design,

installation, and maintenance (Goldin and Katz, 1998; Katz and Margo, 2014). This hypothesis is surely correct: automation that eliminates the need for expertise in one domain often creates demand for novel expertise in another—often in the sector that supplies the automation technology. The importance of this indirect channel is challenging to quantify, however. Evidence from the post-World War II United States tends to find a relatively small contribution of indirect effects that operate across sectors (Acemoglu and Restrepo, 2019; Autor and Salomons, 2018).

By far the largest factor that counterbalances the impact of automation is the simultaneous creation of new expert tasks. In the 19th century, British wages eventually rose and the labor share recovered due to new tasks, such as more-technical work in new manufacturing industries; railways, paying higher wages than prevailed on average in the economy; and new clerical and accounting positions, that finally enabled wages to grow and the labor share to recover (Acemoglu and Johnson, 2023; Allen, 2009). Throughout the 20th century, new tasks typically play an important role in generating new labor demand, as documented by Lin (2011), Autor (2015), Acemoglu and Restrepo (2018), Autor, Chin, Salomons, Seegmiller (2024, 2025), and Acemoglu, Kong, Restrepo (2025).

But even if new task creation is broadly pro-worker, it may or may not benefit the workers most directly affected by automation. In the first four decades after World War II, the flow of new expert work largely replicated the stock of existing work across the occupational distribution, as shown in Figure 2 from Autor, Chin, Salomons, Seegmiller (2024). As automation accelerated in manufacturing and office work between 1940 and 1980, new forms of skilled factory work, such as numerically controlled machining, and skilled office work, such as stenographers and Dictaphone operators, proliferated. But, in the next four decades, the flow of new task creation shifted out of these activities. A disproportionate share of work displaced by automation after 1980 was in production, operative, clerical, and administrative jobs. Yet, much of the new work appearing in these decades was in professional, technical, and managerial occupations—jobs that frequently require far more formal education. New work creation did not, for the most part, open new pathways for the non-college-educated workers displaced by automation.

5 Pro-worker AI in practice

Our framework provides a template for identifying what types of technological changes are likely to be pro-worker. Our main criterion is this: pro-worker technologies make human skills and expertise more useful rather than less necessary. Pro-worker technologies do this

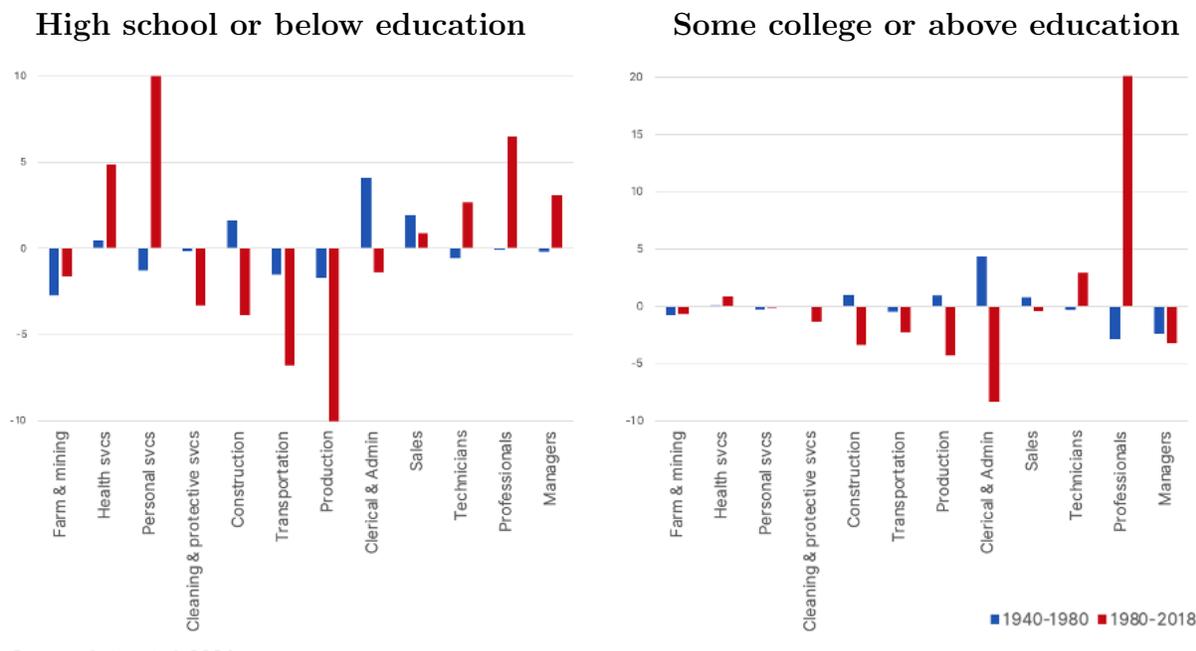


Figure 2: New work creation shifted from the middle of the occupational distribution between 1940 and 1980 to the tails of the distribution between 1980 and 2018

Source: [Autor, Chin, Salomons, Seegmiller \(2024\)](#).

Note: This figure reports the difference in the share of employment in new work versus the share of employment in existing work separately by education group and period. Each panel reports this difference across 12 broad occupational categories ordered from lowest to highest paying. The first set of bars represent the average difference in employment shares over 1940–2018, while the second set of bars represent the difference over 1980–2018.

by extending the relevance, range, and applicability of human expertise, thereby enabling workers to use their expertise more effectively, and, ideally, enabling them to acquire new expertise more efficiently. Workers and technologies collaborate to leverage the value of workers’ expertise in new and in more-sophisticated existing tasks ([Acemoglu, 2021](#); [Acemoglu, Autor, Johnson, 2023](#); [Autor and Manyika, 2025](#)). Pure automation technologies, by contrast, do the opposite of collaborating: they commodify human expertise, rendering it less valuable and potentially superfluous.

We stress that a pro-worker AI tool is not a tool that allows anyone to do anything without need of expertise, judgment, or training. Such a tool, if it existed, would have many virtues. But it would also fully commodify labor, such that its scarcity value would be reduced or eliminated. This would create profound societal challenges, as discussed, for example, in [Bell and Korinek \(2023\)](#), [Autor and Thompson \(2025a\)](#), and [Restrepo \(2025\)](#).

Here we apply this pro-worker framework to the hypothetical Airplane Maintenance and

Repair Technician tools discussed above.

1. The **Aviation Mechanic Automator** is a pure automation tool. It commodifies the expertise that AMTs possess, and it might ultimately render their labor entirely redundant. Even before that happens, this tool will surely reduce the earnings power of AMTs by diluting or eliminating the scarcity value of their specialized knowledge.
2. The **AMT Assistant** is much closer to pro-worker AI. It leverages and extends workers' existing expertise and enables workers to acquire new expertise more efficiently. Like many technologies, however, the AMT Assistant combines technological elements from multiple categories. The AMT Assistant has unambiguously pro-worker potential for enabling workers to perform novel tasks, meaning tasks that were previously infeasible. But the AMT Assistant is also an expertise-leveling technology, enabling less-experienced mechanics and technicians to tackle tasks that previously required workers with more training and experience. There are many potential benefits to expertise-leveling technologies, as noted above, including lowering prices and opening new opportunities for novices. But these capabilities are a mixed bag for workers.
3. **AMT Liftoff** is an unequivocally pro-worker AI because it extends expertise without simultaneously diluting the value of expertise held by others. AMT Liftoff does not create competition for incumbents because, by assumption, there are almost no incumbent aerospace technicians. This might also make AMT Liftoff a relatively rare new technology.

Although our focus is on unambiguously pro-worker AI applications, we note that even the AMT Automator might not be a pure negative for workers. This is true in two senses. First, it likely makes the expertise of other workers in aviation even more valuable—such as the pilots and aircraft crew who will spend less time grounded.¹⁴ Second, and more profoundly, by lowering the expertise barrier to becoming an Airplane Maintenance Technician, the AMT Automator likely opens this occupation to workers who would otherwise have taken other, lower-paying work.

This case is not anomalous. Consider, for example, how the advent of app-based ride-hailing (e.g., Uber and Lyft) transformed the market for taxi and chauffeur drivers. Taxi and chauffeur driver employment in the United States rose by around 170 percent between

¹⁴ This technology should also reduce travel delays and ultimately ticket costs for consumers—though that is not our focus here.

2000 and 2018.¹⁵ Simultaneously, the average earnings of such primarily urban drivers fell relative to the economywide average.¹⁶ The rollout of app-based ride-hailing likely explains both trends. By eliminating the need for expert knowledge of places and routes (as well as for a taxi medallion), ride-hailing opened the taxi-driving occupation to essentially all adults with a car, smartphone, driver’s license, and minimal criminal history. This created stiff competition for incumbent taxi drivers—likely lowering their wages—but also opened new opportunities for millions of workers. This was clearly unfavorable for incumbents, but it was favorable for entrants, who now vastly outnumber those incumbents.¹⁷

We do not categorize ride-hailing as a pro-worker technology. It commodified the services provided by taxi drivers, created a flood of new competition, and, in the process, reduced drivers’ wages (Berger et al. 2018). But we would be remiss if we did not acknowledge that it simultaneously benefited consumers and, by lowering barriers to worker entry, created substantial new employment opportunities. Still, the next phase in the ride-hailing sector appears to be fully autonomous vehicles (e.g., Waymo, Zoox), which are rolling out across U.S. metro areas. Autonomous ride-hailing services are arguably even more pro-consumer than human-piloted ride-hailing. In the longer run, they will be more convenient, safer, and perhaps less expensive than their predecessors. And surely some workers will benefit—those who use the service for commuting, for example. From the perspective of ride-hailing drivers, however, this technology could mean the end of the road.

6 Concrete examples of pro-worker AI

Recall [Licklider’s \(1960\)](#) vision is of a partnership between mind and machine that will “think as no human brain has ever thought and process data in a way not approached by the information-handling machines we know today.” That is our vision, too, and what a notion of pro-worker AI encapsulates. Here we provide five illustrative examples of pro-worker AI tools that are in use or under development, and we explain why we believe that each constitutes pro-worker AI.

¹⁵ According to authors’ calculations using the Current Population Survey ([Flood et al., 2025](#)), there were 288,971 taxi drivers in chauffeurs in 2000 and 780,371 in 2018.

¹⁶ According to the authors’ calculations using the Current Population Survey’s Annual Social and Economic Supplements ([Flood et al., 2025](#)), in 2000 taxi drivers/chauffeurs earned 67% of the economywide average annual wage, salary, and business income; in 2018 they earned 55%.

¹⁷ Ride-hailing work offers relatively low pay and little economic security. Yet millions of workers flooded into this occupation as entry barriers fell—indicating a clear preference over available alternatives. [Garin, Jackson, Koustas, Miller \(2025\)](#) estimate that there were slightly more than 2 million gig workers in the United States in 2018, with nearly all that employment in ride-hailing and delivery platforms.

These concrete examples highlight the broad reach of pro-worker AI across domains—from food delivery to janitorial services, skilled trades, education, and the abstract technical evaluation work of patent examiners. These pro-worker AI applications are in the field, either as prototypes or as fully deployed tools. What they have in common is that they aim to serve as a force magnifier for human expertise. They do not “dumb down” work, but rather enable workers to better leverage their knowledge and judgment, to apply that knowledge to new problems, and, ideally, to acquire new knowledge to tackle new problems.

6.1 Electrician’s assistant

The Electrician’s Assistant (EA) is akin to our AMT Assistant above. Developed by the global electrical services provider Schneider Electric, EA supports electricians and electrical engineers to identify problems with electrical machinery and circuitry. The tool uses a large language model to provide real-time assistance to engineers and electricians performing specific repair and troubleshooting tasks. Workers queue EA by visually and manually entering inputs such as pictures of the relevant hardware. By matching these inputs to a large database of machinery and existing problems, most of which have been previously documented by experienced electricians, the AI model makes recommendations and suggests iterative steps to support field technicians.

Prior to the introduction of the tool, engineers with graduate educations manually reviewed factories’ machinery sensor data and help tickets, wrote reports detailing recommended maintenance and repair actions, and translated deliverables into the worksite client’s language (Godemel, 2024).¹⁸ Trained on these documents, the EA tool summarizes input data, drafts maintenance recommendations, and translates reports into clients’ native language. Among engineers using the tool, the average time for completing maintenance reports is reported to have been halved.

What makes this tool pro-worker? It leverages the expertise of electricians, field technicians, and engineers, enabling them to maintain, troubleshoot, and repair equipment more effectively. The engineer using the tool remains in the loop to modify AI-generated recommendations for accuracy; that is, the engineer is working collaboratively with the tool, not simply subserviently following instructions. Aided by this tool, they can accomplish existing assignments more efficiently and reliably, and tackle more-challenging assignments that were previously out of reach—which we interpret as carrying out new tasks. Tools akin to EA could be readily built to support many additional trade and modern craft workers, such as

¹⁸ Based on this source and company interviews.

plumbers, building contractors, and health-care workers.

There are no data available that would allow us to directly assess the real wage effects of EA. But we expect that it will have a positive impact on the demand for engineers and electricians, which should boost wages. We note, however, that this tool could also increase the effective supply of workers who can take on these technical tasks. And it might heighten competition between less-experienced and more-experienced technicians. As above, it is a rare case in which boosting one set of workers' capabilities does not also introduce new competition for other workers. S. What makes EA likely to be pro-worker, however, is that it makes worker expertise more effective.

6.2 Service worker's assistant

Another type of potentially pro-worker AI tool would be one that helps decision-making in the labor market for low-wage workers, engaged in factory and service jobs.

One such tool, developed through a collaboration between a start-up and a social enterprise employing people with disabilities in the Pacific Northwest, is already in existence. This tool, Empowerment Companion, is a personalized conversational AI assistant for low-wage workers, aiming to provide advice and information for both work tasks and other labor market strategies. The objective is to give continuous guidance and decision support to workers at a nominal cost to firms, potentially transforming how firms hire, train, and support employees.

The AI assistant currently furnishes on-demand task management, feedback, and auditing capabilities to custodial workers. It is based on a large language model that analyzes work orders and then provides task prompts tailored to each worker based on the specific work orders, and repeats them as needed during task completion. For example, a worker might receive guidance on the most appropriate cleaning product and technique to remove a persistent grease stain, complete with video instructions. Using computer vision, the tool also confirms that tasks have been completed according to contract specifications and, if not, redirects workers to address remaining problems. We view this tool as potentially pro-worker because it supports workers in performing their work more effectively, documenting their results, and potentially accomplishing a broader set of tasks. Employers may also benefit from such AI assistants via reduced training costs, managerial time and effort, and monitoring.

However, even seemingly benign technologies can be harnessed for different ends. For example, Amazon.com hires, monitors, and terminates (“deactivates”) its Amazon Flex drivers via a phone-based app called Flex (Crispin, 2021). What is the difference between Empower-

ment Companion and Amazon Flex? Technologically, we would guess that the answer is not much. Yet we would not call Amazon Flex a pro-worker technology. The difference between these tools is intention rather than technology. Empowerment Companion is designed to support worker autonomy and build self-efficacy. Amazon Flex is apparently designed to enforce centralized monitoring and control, reducing worker autonomy in the process. Unfortunately, we have no means to assess which design is more effective and for whom, and no means to forecast which model will win out over the longer run. Is there a pro-worker alternative to Amazon Flex that would deliver equally good or better results for Amazon and its Flex drivers? Alternatively, will employers ultimately conclude that Empowerment Companion would be even more profitable if it were used like Flex to reduce autonomy and intensify monitoring? We cannot, at present, answer either question with confidence.

6.3 Teacher’s AI assistant

A similar collaborative model can be applied to supporting educators and their students. Consider a teaching support tool, trained on relevant curricular material and past classroom experiences, that receives as inputs the responses to short quizzes that students take periodically (e.g., weekly), as well as visual and other real-time input from the classroom. Parsing these materials, the AI assesses what specific difficulties each student is having with the material and recommends to the teacher potential reorganizations of the classroom into smaller groups and menus of lesson plans for each group, depending on the exact challenges of each student, and the group’s composition.¹⁹ (Teachers could deviate from the recommendations of the AI tool.)

What part of this technology is pro-worker (pro-teacher)? While numerous automated teaching offerings attempt to reduce the need for teachers, a successful rollout of an AI tool of the sort might increase teacher efficacy, enabling them to focus their scarce attention in places where it is most needed, and to engage the students at the level that is most helpful for them. The tasks that this tool will enable teachers to perform are more sophisticated and potentially novel, satisfying one part of our definition. Teachers will also have to develop new expertise for harnessing this tool and for mastering a more flexible approach to teaching, thus satisfying the second part of our definition for pro-worker AI. This form of classroom education would likely be still more effective if additional teachers, working in tandem with

¹⁹ Such a tool does not currently exist, but a subset of experts, including the authors, is developing it in collaboration with several small-scale tech companies and nonprofits. The intent is to develop this tool and then evaluate it in several countries, starting in one large state in India, with plans of doing so in Kenya and Nigeria as well.

the tool, were available to support subgroups of students according to their learnings needs. Schools using this technology also might ultimately want to hire more teachers.

This collaborative education paradigm could be readily extended to support workers engaged in team leadership and training, whether as coaches, managers, crew supervisors, or corporate trainers.

6.4 Decision support for patent examiners

Another AI application that has pro-worker potential is search and classification tools that enable workers to perform analyses that are more in-depth. One example is found in the United States Patent and Trademark Office (USPTO). In 2021 USPTO began to integrate AI search tools into its prior art search software, enabling examiners to more precisely conduct prior art searches and reduce examination times using a tool called More Like This ([United States Patent and Trademark Office, n.d.](#)). Developed by Accenture Federal Services in collaboration with USPTO, More Like This enables patent examiners to both process applications faster and, more relevantly from the pro-worker AI perspective, perform better assessments of applications for new patents based on better research on existing patents.

Prior to AI assistance, patent examiners relied on manual document review and Boolean searches (e.g., “Author last name begins with ‘A’”; “Year of issuance is 1964”). These search methods poorly distinguish between relevant and tangentially related prior art. Examiners might spend a full day sifting through thousands of documents while performing a single patent examination.

Distinct from Boolean search, More Like This suggests documents that are similar to those an examiner has identified as relevant ([United States Patent and Trademark Office, 2022](#)). When reviewing search results, examiners can request additional documents that share conceptual relationships with an applicable reference, even when the documents use different terminology. One feature allows examiners to weight the importance of portions of a character string or concept in their searches. Examiners can emphasize novel aspects of an invention while de-emphasizing common industry terminology. Queries to the tool are answered in seconds whereas the prior search tool took hours to process.

As of June 2024, almost 80 percent of USPTO patent examiners had used one or more AI-based search features to supplement existing search methods ([United States Patent and Trademark Office, 2025](#)). Reduced search time allows examiners to dedicate more attention to analyzing the most relevant documents. More importantly, examiners are now able to do

a better job of finding relevant references and additional information for patents.

Is this a pro-worker tool? Our tentative answer is yes. If this tool reduces time spent on low-value supporting tasks (e.g., patent search) and enables examiners to focus on high-value specialized tasks (e.g., evaluation, judgment), it might make the expertise of skilled patent examiners even more valuable. Moreover, if the tool enables patent examiners to tackle novel, previously infeasible tasks (e.g., deeper comparisons between new patent applications and prior art), then this tool would almost surely be pro-worker by our definition. Conversely, however, if this tool does nothing more than enable patent examiners to do their existing tasks more quickly, it may simply save labor without amplifying the value of examiners' expertise.

6.5 Hearing Aids

In 2025 there were more than 200 million gig workers in China, a number that exceeded total U.S. employment in that same year by about 30 million ([The Economist, 2025](#)). A substantial fraction of gig workers does food delivery, and a subset of those workers have hearing impairments ([Chen, Hoffman, Xu, Yuan, 2025](#)). Hearing impairments are surprisingly consequential in this line of work. When pulling up with their delivery, gig workers in China are expected to call recipients to gain entry to access-controlled buildings or to locate units within large complexes.

The requirement to communicate with customers via voice calls puts hearing-impaired workers at a disadvantage. As documented in [Chen, Hoffman, Xu, Yuan \(2025\)](#), hearing-impaired workers complete fewer orders and receive more bad reviews than others with similar platform experience. Despite these potentially discouraging results, hearing-impaired workers typically log more hours and spend more months on the platform than non-hearing-impaired workers, suggesting that their options elsewhere in the economy are comparatively limited.

In 2024 software developers at one gig platform in China took notice of this problem, and built a simple technological fix to address it: a voice chatbot embedded in the delivery app that performs real time text-to-speech and speech-to-text communications for workers with hearing impairments. As [Chen, Hoffman, Xu, Yuan \(2025\)](#) document, this simple tool closed the customer review gap between hearing-impaired and other drivers, and opened a positive performance gap between hearing-impaired and non-hearing-impaired workers. With the communication barrier removed, hearing-impaired workers performed at a level commensurate with their greater experience.

This instance of pro-worker AI is so straightforward that one may wonder if it even fits our definition. It does so, because this technology makes human skills and expertise more valuable. The core service of gig delivery workers is moving food from restaurants to customers. Communicating with customers is a trivial, inexpert task for most delivery workers, those who have good hearing, but it is nevertheless essential to success. The addition of a simple voice-to-text-to-voice tool to the delivery app enabled this capability for hearing-impaired workers—with transformational consequences. If this pro-worker application of AI seems prosaic, recall that real-time voice-to-text and text-to-voice services are AI-enabled tools that were virtually unknown on smartphones before Apple introduced its voice assistant, Siri, in 2011.

7 Why isn't pro-worker AI everywhere?

Pro-worker tools are arguably pro-firm tools as well. They enable workers to produce more value with the same inputs, or the same value with fewer inputs.²⁰ Alongside saving costs, these tools might also enable firms to create and offer new goods and services. Given these virtues, one should expect the market to capitalize on these opportunities. Is that occurring?

Our reading of the evidence is that it is not. Although no authoritative estimates exist, anecdotal evidence suggests that a great deal of the current AI focus is on task automation and development of high-level capabilities in line with AGI, which does not cohere with our pro-worker definition, (since, as discussed above, AGI would ultimately aim to replace most things that workers do, rather than to create new tasks and capabilities for them).

What is wrong with automation? Nothing, fundamentally. Automation drives efficiencies that lower costs and raises profits, which is not a bad thing. Some of these efficiencies benefit consumers. As we discussed previously, automation may also raise wages, even if it drives down the labor share and displaces workers, and we are of course supportive of efficiency improvements.

But incremental automation—doing existing tasks better, cheaper, and faster—is somewhat overrated. The most consequential innovations are those that enable new capabilities rather than merely reducing the costs of existing capabilities. Iconic companies of the early 20th century, such as the Ford Motor Company, are not remembered because they reduced costs

²⁰ More precisely, anything that raises productivity creates surplus that must accrue to some combination of workers, firms, and consumers. Our definition of pro-worker AI stipulates conditions under which some of the surplus is likely to go to workers. Although the division of the remainder between firms and consumers is indeterminate, it is reasonable to assume that both would benefit on average, though not in each case.

by a few percentage points, but because they introduced new products that revolutionized their industries and improved the lives of their consumers.

These innovations depended on state-of-the-art tools. But they also fundamentally leveraged human creativity and ingenuity. By unlocking new possibilities, such innovations also generated new employment and novel demands for expertise. There were no automotive mechanics, aircraft crews, household plumbers, geneticists, or television actors until supporting innovations created the need for these specialized skill sets. While it is commonplace to think of new technologies as work-replacing, they are also one of the principal sources for the emergence of new work.

7.1 What is the opposite of pro-worker technology?

While we have argued above that automation is a mixed bag, some workplace applications of AI are far more pernicious. One such application is workplace surveillance, which AI has made increasingly cheap and powerful. There are now hundreds of AI-based workplace surveillance systems.

Some degree of monitoring is obviously useful for firms. However, aggressive workplace monitoring can border on coercive control, leaving workers fearful of even taking lunch and bathroom breaks. This is not a hypothetical possibility. Research documents adverse impacts of surveillance on self-reported worker mental health, while undercover reporting documents that warehouse workers skip bathroom breaks to keep their jobs ([American Psychological Association, 2023](#); [Liao, 2018](#)). Coercive control can also enable firms to reduce sharing of profits with workers that would have previously been used for at incentivizing and motivating workers. Intuitively, if managers can use surveillance to enforce worker compliance, they have less need to motivate workers through generosity, goodwill, or mission alignment.

Surveillance is not limited to the office. One of the world’s largest call center subcontractors in Colombia, Teleperformance Group , for example, has started using AI tools to analyze real-time audio-video feeds from workers’ homes to enforce company standards. In its current incarnation, a computer vision system analyzes audio and visual data from cameras installed in employees’ home workspaces. Potential contract violations are flagged based on worker movement and computer activity inputs. If workers do not register keyboard or mouse activity within defined intervals, the AI logs them as “idle.” When the system detects infractions such as an employee eating at their desk or looking at a phone, it sends managers screenshots for review.

Recognizing the challenge in making empirically grounded comparisons in this arena, we are nevertheless comfortable with stating that much of the energy and investment in workplace AI is not flowing in a pro-worker direction.

7.2 What is the market failure?

Why is the market not going all-in on pro-worker AI?

One reason firms might not be adopting pro-worker AI is that, contrary to our arguments, it is not technically feasible for them—or, even if feasible, it is less profitable than alternative applications. A related possibility is that firms expect these investments to be overtaken by events. If AGI can truly be achieved within a few years, as some AI evangelists forecast, why bother investing in pro-worker technologies at all? We are skeptical that AGI is anywhere close to imminent, but organizations that hold this view may see little point in pro-worker AI. We doubt this is the whole story, however.

A second class of explanation, and one that we find compelling, is market failure, by which we mean misaligned incentives. Here, we distinguish between misaligned incentives for those who purchase and deploy AI and misaligned incentives for those who develop and sell AI technologies.

Misaligned firm incentives

Since higher productivity should translate into higher profits in a competitive labor market, we should expect managers of profit-maximizing firms to deploy whichever technologies most increase productivity. This logic says that firms will adopt pro-worker AI tools if they are available and more profitable than the alternatives. Yet, our informal assessment is that they are not doing so. If there is an incentive problem, what is it?

One possibility is that such tools might not yet be widely available. If AI developers are pursuing primarily pure automation tools, perhaps because they believe that customers are laser focused on cost-cutting, then pro-worker AI tools may be slow to appear, immature, and not yet ready for prime time. In that case, we would expect the market to catch up with these opportunities eventually.

A second consideration is that managers could have objectives other than maximizing overall productivity (and might even deviate from pure profit maximization). In unionized firms, especially where firm-union relations are conflictual, managers may wish to adopt automation

technologies to reduce their dependence on unionized labor, strengthen their bargaining hand, and perhaps ultimately weaken or decertify unions.²¹ Since pro-worker AI is likely to make a firm’s workforce more rather than less essential to its success, firms with conflictual worker-management relations might, in general, be reluctant to invest in such tools.

Even in non-unionized firms, managers could still wish to adopt technologies that reduce workers’ surplus and redistribute the surplus to shareholders. This adoption can occur if workers earn wages in excess of their outside options, or “rents.” Imagine, for example, a firm with a highly profitable hit product, the production of which depends on the specialized knowledge of the firm’s workforce. Because of their centrality, those production workers will likely earn rents. These rents provide managers with an incentive to make these workers more replaceable, thereby capturing a larger share of those profits for shareholders and themselves. While this idea might seem far-fetched, evidence suggests that firms target automation toward “high-rent” tasks, and that this targeting works to the detriment of worker earnings (Acemoglu and Restrepo, 2026). Ironically, the incentives created by these rents can themselves create economic distortions: firms may choose rent-dissipating automation technologies that lower overall productivity, as long as the managerial and shareholder slices of the economic pie expand by more than the pie shrinks.

Misaligned developer incentives

Technology developers face their own set of incentives. If purchasers of technology products prefer automation technologies (e.g., for the reasons above), then tech firms will have a strong incentive to prioritize them. Again, in the longer run we might expect a market correction, if indeed more pro-worker technologies offer unexploited opportunities.

A second concern is path dependence rooted in incumbent firms’ business models. These models determine how companies earn profits and where they concentrate organizational resources and expertise. Because leading AI firms have built deep capabilities in developing and marketing automation software to enterprise customers, their business models naturally orient them toward further automation. Combined with their market dominance, this creates strong momentum toward automation technologies. In the extreme, incumbents may view pro-worker alternatives as cannibalizing their core business or jeopardizing their primary profit stream. Start-ups face similar pressures: if building on incumbents’ platforms and selling their new technologies to incumbents proves to be more profitable than competing alternatives, start-ups, too, will be discouraged from pursuing pro-worker tools. The prob-

²¹ Evidence suggests that firms close unionized plants and shift their business to non-unionized plants. See Wang and Young (2024).

lem intensifies when pro-worker technologies require years of investment while automation solutions are already market ready.

Workers may also resist pro-worker AI tools that require them to acquire new expertise and adjust their work habits. If workers are reluctant to make these investments or lack the foundational skills that are required to do so, firms could be discouraged from developing or adopting pro-worker AI.

Finally, we suspect that a specific vision of AI, one focused on AGI, shapes the trajectory of technology development. Since its inception, the field of AI has focused on making machines that can outthink humans at all cognitive tasks. The computer science and AI communities are gripped by an ideological vision that places AGI as its highest possible pursuit, where AGI implies machines that exceed all human capabilities. [Acemoglu and Johnson \(2023\)](#) detail how this perspective has come to dominate the AI field and how it affects the direction of AI today.

While ideology does not ultimately determine what technologies are feasible, it does shape how firms, researchers, and funding agencies direct their research and development efforts. If firms are single-mindedly pursuing AGI-driven automation, this makes it less likely that they will build pro-worker AI.

8 Policies that would promote pro-worker AI

Pro-worker AI is not a fundamentally distinct technology. Rather, it is a distinct vision for how to use existing and emerging technologies. The goal of pro-worker AI—and pro-worker technologies more generally—is to enable workers to accomplish tasks more effectively, to tackle new tasks, and to master new expertise required for new work.

We have argued that the market is likely undersupplying pro-worker AI due to misaligned incentives, path dependence in innovation, and a pervasive pro-automation, pro-AGI mindset. Are there ways to channel some of the abundant talent, creativity, and resources flowing into the AI sector in a pro-worker direction? We see nine ways that public policy could channel advances in AI toward being pro-worker if desired.²²

²² See [Korinek and Stiglitz \(2025\)](#) for a theoretical discussion of policy tools for steering the direction of technical progress.

8.1 Shaping pro-worker AI in health care and education

It is challenging to target human-complementary work in the abstract. It is, however, feasible to focus on specific sectors and activities where opportunities are already abundant. Two sectors where pro-worker AI has enormous potential to improve the quality of work and simultaneously create enormous social value are health care and education.

In 2023 the U.S. spent approximately 18 percent of GDP on health care and around 6 percent of GDP on education ([Organisation for Economic Co-operation and Development, 2025](#); [Peter G. Peterson Foundation, 2025](#)). These sectors command such a large share of GDP because they employ vast numbers of highly trained decision-makers (e.g., teachers, medical workers) whose work is largely artisanal: educators and health-care workers provide individualized services to students and patients. Arguably because this artisanal work has proved ill-suited to the industrial model that, over the past century, transformed factories and offices, productivity growth is notoriously slow in both sectors ([Chandra and Skinner, 2012](#); [Hoxby, 2004](#)). Because AI can support expert decision-making at scale—as many of our examples above indicate—we believe that AI offers great potential to raise the productivity and expand the capabilities of educators and care-workers.

Health care and education are also two sectors where public policy has the great leverage. In 2023 the U.S. spent \$4.9 trillion on health care, and nearly half (43 percent) of that total, \$2.1 trillion, was paid by government insurance programs, primarily Medicare and Medicaid ([Peter G. Peterson Foundation, 2025](#)). The public share of education dollars is even larger. The [Organisation for Economic Co-operation and Development \(2025\)](#) estimates that 92 percent of K-12 expenditure in the United States and 39 percent of tertiary expenditure is publicly funded. Through purchasing, agenda-setting, and regulation, the public has outsized capacity to influence how AI is developed and deployed in these two enormous, vital sectors.

Indeed, the public sector already heavily shapes the path of technology in health care and education. For example, the federal Health Information Technology for Economic and Clinical Health (HITECH) Act of 2009 dramatically accelerated electronic health record adoption in U.S. hospitals through financial incentives and penalties. Within less than a decade, the United States went from approximately 10 percent of hospitals with electronic health records to near-universal adoption ([Office of the National Coordinator for Health Information Technology, 2017](#)). In a similar vein, the federal schools and libraries universal support (E-Rate) program, established by the Telecommunications Act of 1996, provides ongoing subsidies to schools and libraries for Internet connectivity. As of 2021, 95 percent of U.S. public school classrooms had WiFi ([Munson, 2023](#)).

Due to the applicability of AI in health care and education, there is immense potential for using AI to complement these sectors' workforces while improving outcomes for their customers (i.e., students, patients). The fact that the public is already the majority stakeholder means that there is also ample opportunity to shape the development and adoption of pro-worker potential in both sectors.

8.2 Building AI expertise in government

How might the public sector could use this leverage if desired? We do not envision that the federal government would become a leading developer of AI for schools and health care, nor that it would dictate precisely how AI is used. But the federal and state governments could set goals for AI deployment: they could design objectives for how AI can be deployed to support health practitioners and classroom educators, procure supports that target these objectives, and even set Medicare and Medicaid reimbursement policies for health-care delivery by non-physician professionals (such as nurse practitioners) supported by AI tools.

To realize any of these objectives, AI expertise will be badly needed within the federal and state governments. Those experts, in collaboration with relevant stakeholders (i.e., schools, hospitals, parents, patients, teacher organizations, medical professions), could lead this process.

More broadly, AI will touch every area of government investment, regulation, and oversight, including, but not limited to, transportation, energy production, labor conditions, health care, education, environmental protection, public safety, and military capabilities. Seizing this opportunity will require state capacity that we currently lack. One means to remediate this shortfall is to develop a consultative AI division within the federal government, with the goal of supporting the many agencies and regulators who can influence AI development.

8.3 Using grant-making to support worthwhile AI investments

Until very recently, the federal government funded about one in five U.S. R&D dollars ([National Science Board, 2025](#)). Deployed by the National Science Foundation, National Institutes of Health, and the Departments of Defense, Energy, and Agriculture, among others, federal funding has fundamentally shaped the path of research in both academia and the private sector since World War II ([Gross and Sampat, 2023](#)). Beyond the short-term cycle of grant applications and funding decisions, federal grant funding shapes the path of research by building depth, capacity, and regional centers of excellence.

Grant-making could be directed to shape the course of (some) AI development in a pro-worker direction. One emerging area of research is particularly relevant: how to design AIs to collaborate effectively with workers to support rather than undermine sound decision-making, and to foster rather than thwart learning. Abundant evidence demonstrates that getting such designs right is challenging. While some human-AI combinations perform better than either humans or AIs alone, others degrade human-alone or AI-alone performance (Vaccaro, Almaatouq, Malone, 2024; Agarwal, Moehring, Rajpurkar, Salz, 2023; Buçinca et al., 2025). We are confident that research can advance quickly in this arena and that this research will prove crucial for successful design of pro-worker AI tools.

It is undeniable that the locus of frontier AI research has shifted decisively toward the private sector, with close to 70 percent of U.S. AI PhDs hired by industry in 2020 versus 21 percent in 2004 (Ahmed, Wahed, Thompson, 2023). Still, the federal government invested close to \$3 billion per year in AI R&D between 2022 and 2024, which is surely enough money to move the needle (U.S. Networking and Information Technology Research and Development Program, n.d.).

8.4 Fostering competition for excellence

A drawback of using public funding for advancing industry is that doing so could short-circuit the competitive process that determines which ideas succeed and which fail. Recognizing this danger, we would note that R&D and procurement activities in this domain could use a Defense Advanced Research Projects Agency (DARPA)-style competitive prize model rather than a traditional procure-to-specifications model. Notably, China fosters intense competition among upstart entrants in sectors that it targets with industrial policy (Chen and Yuan 2025). Simultaneously, while many of China’s world-renowned firms have emerged from this economic-Darwinism-as-industrial-policy model, the intensity of competition means that most publicly-subsidized competitors that China stands up will ultimately fail (Boeing, Eberle, Howell, 2022). It is not clear, however, that the U.S. political system is able to tolerate the competitive failure of large numbers of publicly funded firms, even if this strenuous competition fosters an ecosystem from which excellence emerges.

8.5 Shifting investment incentives in the tax code

The current U.S. tax code places a much heavier burden on firms that hire labor than those that invest in algorithms to automate work (Acemoglu, Manera, Restrepo, 2020). Policy

could aim to create a more symmetric tax structure, where marginal taxes for hiring and training workers and for investing in equipment and/or software are equated. This could shift incentives toward pro-worker technological choices by reducing the bias of the tax code toward physical capital over human capital (Acemoglu, Manera, Restrepo, 2020).

8.6 Fostering private sector competition

Antitrust enforcement could foster greater competition in technology development. The concentration of power among a handful of dominant AI firms extends beyond typical competition concerns: the power of network effects in these markets is so great that these firms arguably propagate their business models and technology choices across an entire ecosystem. Start-ups naturally gravitate toward technologies they can sell to large incumbents or that position them as attractive acquisition targets. When dominant firms derive profits primarily from automation and data monetization through digital advertising, this dynamic discourages development of pro-worker alternatives. Antitrust action could counteract these pressures. Stricter scrutiny of mergers and acquisitions, along with thwarting predatory pricing aimed at eliminating competitors, could create space for new business models that are more conducive to pro-worker AI.

8.7 Harnessing worker voice

AI will profoundly affect all workers, yet workers as a group currently lack meaningful voice in shaping AI's development and deployment. Establishing institutional mechanisms for worker input could prove valuable. Civil society organizations, including labor unions, can help advance this goal by articulating workers' needs and advocating for protective frameworks at local, state, and federal levels. As one example, health and safety rules may need to be updated to place limits on deployment of untested (or insufficiently tested) AI for applications that could put workers at risk, including those used for intrusive workplace monitoring and surveillance.

8.8 Establishing worker ownership of expertise

Current intellectual property law was designed for an earlier era and offers little protection against AI's industrial-scale harvesting of human expertise. AI systems freely scrape content from websites, social media, YouTube, newspapers, Wikipedia, and blogs, then statistically

recombine this material and sell access to the results. Authors, journalists, visual artists, musicians, translators, and countless other creators find their work appropriated as training data, with no compensation or control (Kasy, 2025).

The challenge extends beyond publicly available content. Firms increasingly train AI models on their own employees' expert performance—an appealing shortcut with dangerous implications for workers. Few employees would willingly train an apprentice designed to replace them, and yet this is precisely what happens when companies use worker expertise to build automation systems. Once training is complete, the AI can replicate the performance while the human expert becomes expendable.

We do not assert that all AI training constitutes intellectual property theft. But neither do we accept the opposite extreme—now widely assumed in the AI industry—that creators should receive no compensation when their work is harvested by machines that could ultimately imitate and compete with them. The current arrangement resembles the Napster era of music piracy, when technological change outpaced legal frameworks and creators bore the costs. Left unaddressed, this arrangement will prove profoundly anti-worker. Building legal frameworks that support workers' ownership of their capabilities and creative output would give workers greater control over how their expertise is used, reduce the commodification of human knowledge, and preserve incentives to invest in skill development and innovation.

8.9 Addressing licensure barriers

A central promise of pro-worker AI is that it can enable workers to accomplish a broader array of tasks by leveraging their expertise with better tools. If this vision is correct, it almost inevitably means that AI will also instigate additional competition among workers of different expertise levels (what we termed expertise-leveling above). If they are using better tools, nurse practitioners can do some tasks traditionally done by doctors, paralegals can do tasks previously assigned to lawyers, junior craft workers can do tasks performed by senior craft workers, and non-PhD researchers can do tasks done by PhD researchers; this shift will create frictions as entrants battle for occupational turf and incumbents erect barriers to thwart their entrance.

Such frictions are commonplace in the professions, in the trades, and in many services, where a thicket of licensure requirements and scope of practice boundaries serve to soften competition and protect incumbents from upstart entrants. *The American Medical Association*, for example, has for decades resisted calls to expand the scope of practice of nurse practitioners, despite their capacity to handle many diagnostic and treatment tasks that were once limited

to medical doctors (Avi-Yonah, 2023). A large literature studies the political economy and economic consequences of occupational licensure and generally finds that it stifles competition, raises prices, and has, at best, mixed effects on service quality (Allensworth, 2025; Kleiner, 2016; Kleiner and Soltas, 2023).

These challenges will likely become more severe as AI provides tools that enable entrants to better compete with established experts in offering valuable services. Policymakers could attend closely to this dynamic, lest the potential gains from newly-empowered human experts are thwarted by incumbents.

9 Conclusion

This essay defines pro-worker technologies—including AI—as technologies that make human skills and expertise more valuable by expanding human capabilities. While AI’s capacity to automate work and displace workers is beyond doubt, we simultaneously believe that, used well, AI has equally momentous potential to act as a force-multiplier for human skills and expertise. This potential arises from AI’s capacity to collaborate with workers, to enable them to be more effective at their existing tasks, to tackle new tasks, and to acquire new expertise. This collaborative capacity is what gives AI its potency as a pro-worker technology.

Our focus on pro-worker AI might seem anachronistic or naïve to those who believe that we are on the cusp of AGI—a technology that threatens to render all human expertise superfluous. Although we cannot definitively reject the AGI thesis, we believe that the AGI goal is distant. Moreover, we are certain that banking on the unproven assumption that AGI will soon deliver us from the need for human expertise is grossly misguided. AGI’s expected arrival can be (and frequently is) invoked as an argument against almost any long-term investment other than in AGI itself. Why build pro-worker AI if AGI will supersede it shortly? Why advise workers to invest in their own expertise if the labor market for those capabilities is about to collapse? In our view, it is foolish to treat the AGI future as certain. We should instead build judiciously for numerous possible futures—many of which do not include AGI.

The case for pro-worker AI is not merely precautionary: it is practical. AI’s demonstrated prowess in narrow domains has fostered inflated expectations about what pure automation can achieve. The logic runs as follows: if AI has already surpassed human performance in specific tasks, it must be capable of replicating everything experts do—just without the experts themselves. This assumption shapes thinking across the spectrum, from enthusi-

asts championing AI’s potential to critics warning of its dangers. Both camps envision a similar trajectory: AI first mimics expert capabilities, then exceeds human performance, and ultimately renders experts obsolete. In this framing, AI does not enhance expertise—it eliminates the need for it entirely.

Although this vision may notch occasional successes, we believe that it will fail more often than not. AI is not ready to automate most expert work. The stakes are too high and the decisions too nuanced in much human work to fully delegate these roles to opaque systems that operate on their own discretion.

Though full automation remains a distant prospect for much decision-making work, the opportunities for human-AI collaboration are immediately available. AI excels as a partner precisely because its strengths—inexhaustible memory, fast pattern recognition, and continuous operation—compensate for human limitations. Where experts struggle to recall every relevant precedent, consider all possible scenarios, or synthesize insights across disparate fields, AI can fill these gaps. In doing so, it often enhances distinctively human capacities: interpreting context, making ethical judgments, generating novel solutions, and understanding how individual tasks advance larger objectives.

The choice between automation and collaboration need not be absolute, and the division of labor between human judgment and AI capabilities will shift as the technology matures. AI is and will be transformative—a reality already evident in many domains. The crucial question is therefore not whether to deploy AI but rather how to deploy it wisely. Do we pursue maximum automation wherever possible? Or do we develop AI systems that learn from human decision-making, augment human judgment, and work alongside workers to improve outcomes? The correct answer, of course, is both. Getting this balance right across capabilities is a formidable and ever-evolving challenge. Building tools that make human skills and expertise more valuable should be one principal strategy for responding to this challenge.

References

Moses Abramovitz and Paul David. American macroeconomic growth in the era of knowledge-based progress: The long-run perspective. In Robert E. Gallman and Stanley L. Engerman, editors, *The Cambridge Economic History of the United States: Volume 3: The Twentieth Century*, volume 3 of *Cambridge Economic History of the United States*. Cambridge University Press, 2000. doi: 10.1017/CHOL9780521553087.002.

Danielle Abril. AI is taking on live translations. But jobs and meaning are getting lost.

- The Washington Post*, September 2025. URL <https://www.washingtonpost.com/business/2025/09/26/ai-translation-jobs/>.
- Daron Acemoglu. AI's future doesn't have to be dystopian. *Boston Review*, 2021. URL <https://www.bostonreview.net/forum/ais-future-doesnt-have-to-be-dystopian/>.
- Daron Acemoglu. *What happened to liberal democracy? Remaking a politics of shared prosperity*. Penguin Random House, 2026.
- Daron Acemoglu and David Autor. Skills, tasks and technologies: Implications for employment and earnings. In David Card and Orley Ashenfelter, editors, *Handbook of Labor Economics*, volume 4. Elsevier, 2011. doi: 10.1016/S0169-7218(11)02410-5.
- Daron Acemoglu and Simon Johnson. *Power and progress: Our thousand-year struggle over technology and prosperity*. PublicAffairs, 2023.
- Daron Acemoglu and Simon Johnson. Learning from Ricardo and Thompson: Machinery and labor in the early Industrial Revolution and in the age of Artificial Intelligence. *Annual Review of Economics*, 16(1):597–621, 2024. doi: 10.1146/annurev-economics-091823-025129.
- Daron Acemoglu and Pascual Restrepo. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542, 2018. doi: 10.1257/aer.20160696.
- Daron Acemoglu and Pascual Restrepo. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30, 2019. doi: 10.1257/jep.33.2.3.
- Daron Acemoglu and Pascual Restrepo. Robots and jobs: Evidence from U.S. labor markets. *Journal of Political Economy*, 128(6):2188–2244, 2020. doi: 10.1086/705716.
- Daron Acemoglu and Pascual Restrepo. Tasks, automation, and the rise in U.S. wage inequality. *Econometrica*, 90(5):1973–2016, 2022. doi: 10.3982/ECTA19815.
- Daron Acemoglu and Pascual Restrepo. Automation and rent dissipation: Implications for wages, inequality, and productivity. *The Quarterly Journal of Economics*, 2026. doi: 10.1093/qje/qjag006.
- Daron Acemoglu and James A. Robinson. *The narrow corridor: States, societies, and the fate of liberty*. Penguin Press, 2019. URL <https://www.penguinrandomhouse.com/books/555400/the-narrow-corridor-by-daron-acemoglu-and-james-a-robinson/>.
- Daron Acemoglu, Andrea Manera, and Pascual Restrepo. Does the U.S. tax code favor automation? *Brookings Papers on Economic Activity*, (Spring):231–300, 2020.
- Daron Acemoglu, David Autor, and Simon Johnson. Can we have pro-worker AI? Choosing a path of machines in service of minds. Policy brief, MIT, 2023. URL

<https://shapingwork.mit.edu/wp-content/uploads/2023/09/Pro-Worker-AI-Policy-Memo.pdf>.

Daron Acemoglu, Fredric Kong, and Pascual Restrepo. Tasks at work: Comparative advantage, technology and labor demand. In Christian Dustmann and Thomas Lemieux, editors, *Handbook of Labor Economics*, volume 6. Elsevier, 2025. doi: 10.1016/bs.heslab.2025.08.003.

Nikhil Agarwal, Alex Moehring, Pranav Rajpurkar, and Tobias Salz. Combining human expertise with Artificial Intelligence: Experimental evidence from radiology. Working paper, MIT, 2023.

Nur Ahmed, Muntasir Wahed, and Neil C. Thompson. The growing influence of industry in AI research. *Science*, 379(6635):884–886, 2023. doi: 10.1126/science.ade2420.

Robert C. Allen. Engels’ pause: Technical change, capital accumulation, and inequality in the British Industrial Revolution. *Explorations in Economic History*, 46(4):418–435, 2009. doi: 10.1016/j.eeh.2009.04.004.

Rebecca Haw Allensworth. *The licensing racket: How we decide who is allowed to work, and why it goes wrong*. Harvard University Press, 1st edition, 2025.

American Psychological Association. 2023 work in America™ survey: Artificial intelligence, monitoring technology, and psychological well-being, 2023. URL <https://www.apa.org/pubs/reports/work-in-america/2023-work-america-ai-monitoring>.

David Autor. Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3):3–30, 2015. doi: 10.1257/jep.29.3.3.

David Autor and James Manyika. A better way to think about AI. *The Atlantic*, August 2025. URL <https://www.theatlantic.com/technology/archive/2025/08/ai-job-loss-human-enhancement-google/683963/>.

David Autor and Anna Salomons. Is automation labor-displacing? Productivity growth, employment, and the labor share. *Brookings Papers on Economic Activity*, 2018. URL <http://www.jstor.org/stable/26506212>.

David Autor and Neil Thompson. Beyond job displacement: How AI could reshape the value of human expertise. *The Digitalist Papers*, December 2025a. URL <https://www.digitalistpapers.com/vol2/autorthompson>.

David Autor and Neil Thompson. Expertise. *Journal of the European Economic Association*, 23(4):1203–1271, 2025b. doi: 10.1093/jeea/jvaf023.

David Autor, Frank Levy, and Richard Murnane. The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, 118(4): 1279–1333, 2003. doi: 10.1162/003355303322552801.

David Autor, David Dorn, and Gordon Hanson. When work disappears: Manufacturing

- decline and the falling marriage market value of young men. *American Economic Review: Insights*, 1(2):161–178, 2019. doi: 10.1257/aeri.20180010.
- David Autor, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. The fall of the labor share and the rise of superstar firms. *Quarterly Journal of Economics*, 135(2):645–709, 2020.
- David Autor, Caroline Chin, Anna Salomons, and Bryan Seegmiller. New frontiers: The origins and content of new work, 1940–2018. *The Quarterly Journal of Economics*, 139(3):1399–1465, 2024. doi: 10.1093/qje/qjae008.
- David Autor, Caroline Chin, Anna Salomons, and Bryan Seegmiller. What makes new work different from more work? *Annual Review of Economics*, 2025. Forthcoming.
- Shera Avi-Yonah. Medicine without doctors? State laws are changing who treats patients. *The Washington Post*, August 2023. URL <https://www.washingtonpost.com/health/2023/08/20/nurse-doctor-scope-medical-titles/>.
- Stephanie A. Bell and Anton Korinek. AI’s economic peril. *Journal of Democracy*, 34(4): 151–161, 2023.
- Dan A. Black, Terra G. McKinnish, and Seth G. Sanders. Does the availability of high-wage jobs for low-skilled men affect welfare expenditures? Evidence from shocks to the steel and coal industries. *Journal of Public Economics*, 87(9):1921–1942, 2003. doi: 10.1016/S0047-2727(02)00014-2.
- William Blake. Jerusalem [“and did those feet in ancient time”], 1804. URL <https://www.poetryfoundation.org/poems/54684/jerusalem-and-did-those-feet-in-ancient-time>. Poetry Foundation.
- Philipp Boeing, Jonathan Eberle, and Anthony Howell. The impact of China’s R&D subsidies on R&D investment, technological upgrading and economic growth. *Technological Forecasting and Social Change*, 174:121212, January 2022. doi: 10.1016/j.techfore.2021.121212.
- Zana Buçinca, Siddharth Swaroop, Amanda E. Paluch, Finale Doshi-Velez, and Krzysztof Z. Gajos. Contrastive explanations that anticipate human misconceptions can improve human decision-making skills. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–25, April 2025. doi: 10.1145/3706598.3713229.
- Edward C. Budd. Factor shares, 1850–1910. In *Trends in the American economy in the nineteenth century*. Princeton University Press, 1960. URL <https://www.nber.org/books-and-chapters/trends-american-economy-nineteenth-century/factor-shares-1850-1910>.
- Bureau of Labor Statistics. 2018 Standard Occupational Classification System, 2018. URL https://www.bls.gov/soc/2018/major_groups.htm.

- Bureau of Labor Statistics. 33-9091: Crossing guards and flaggers, April 2024. URL <https://www.bls.gov/oes/2023/may/oes339091.htm>.
- Bureau of Labor Statistics. Aircraft and avionics equipment mechanics and technicians, 2025a. URL <https://www.bls.gov/ooh/installation-maintenance-and-repair/aircraft-and-avionics-equipment-mechanics-and-technicians.htm>.
- Bureau of Labor Statistics. Data scientists, 2025b. URL <https://www.bls.gov/ooh/math/data-scientists.htm>.
- Anne Case and Angus Deaton. *Deaths of despair and the future of capitalism*. Princeton University Press, 2022.
- Amitabh Chandra and Jonathan Skinner. Technology growth and expenditure growth in health care. *Journal of Economic Literature*, 50(3):645–680, 2012. doi: 10.1257/jel.50.3.645.
- Yanyou Chen, Mitchell Hoffman, Huilan Xu, and Zhe Yuan. Empowering inclusive work. Preprint, 2025.
- Jessa Crispin. Welcome to dystopia: Getting fired from your job as an Amazon worker by an app. *The Guardian*, July 2021. URL <https://www.theguardian.com/commentisfree/2021/jul/05/amazon-worker-fired-app-dystopia>.
- Thomas H. Davenport and DJ Patil. Data scientist: The sexiest job of the 21st century. *Harvard Business Review*, 2012. URL <https://hbr.org/2012/10/data-scientist-the-sexiest-job-of-the-21st-century>.
- Federal Aviation Administration. Become an aviation mechanic, n.d.a. URL <https://www.faa.gov/mechanics/become>.
- Federal Aviation Administration. Air traffic controller qualifications, n.d.b. URL <https://www.faa.gov/air-traffic-controller-qualifications>.
- Sarah Flood, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, Daniel Backman, Etienne Breton, Grace Cooper, Julia A. Rivera Drew, Stephanie Richards, David Van Riper, and Kari C. W. Williams. IPUMS CPS: Version 13.0 [dataset], 2025.
- Andrew Garin, Emilie Jackson, Dmitri Koustas, and Alicia Miller. The impact of third-party reporting on tax compliance: Evidence from gig workers. SSRN Scholarly Paper 5239694, Social Science Research Network, May 2025.
- Frédéric Godemel. Three ways generative AI is helping our services experts become superheroes. *Schneider Electric Blog*, May 2024. URL <https://blog.se.com/services/2024/05/14/three-ways-generative-ai-is-helping-our-services-experts-become-superheroes/>.
- Claudia Goldin and Lawrence F. Katz. The origins of technology-skill complementarity. *The Quarterly Journal of Economics*, 113(3):693–732, 1998. doi:

10.1162/003355398555720.

Daniel P. Gross and Bhaven N. Sampat. America, jump-started: World War II R&D and the takeoff of the U.S. innovation system. *American Economic Review*, 113(12): 3323–3356, 2023. doi: 10.1257/aer.20221365.

Caroline M. Hoxby. Productivity in education: The quintessential upstream industry. *Southern Economic Journal*, 71(2):209–231, 2004. doi: 10.2307/4135289.

Benjamin F. Jones and Xiaojie Liu. A framework for economic growth with capital-embodied technical change. *American Economic Review*, 114(5):1448–1487, 2024. doi: 10.1257/aer.20221180.

Loukas Karabarbounis. Perspectives on the labor share. *Journal of Economic Perspectives*, 38(2):107–136, 2024. doi: 10.1257/jep.38.2.107.

Maximilian Kasy. *The means of prediction: How AI really works (and who benefits)*. University of Chicago Press, Chicago, 2025.

Lawrence F. Katz and Robert A. Margo. Technical change and the relative demand for skilled labor: The United States in historical perspective. In Leah Platt Boustan, Carola Frydman, and Robert A. Margo, editors, *Human Capital in History: The American Record*. University of Chicago Press, 2014.

Morris M. Kleiner. Battling over jobs: Occupational licensing in health care. *American Economic Review*, 106(5):165–170, 2016. doi: 10.1257/aer.p20161000.

Morris M. Kleiner and Evan J. Soltas. A welfare analysis of occupational licensing in U.S. states. *The Review of Economic Studies*, 90(5):2481–2516, 2023. doi: 10.1093/restud/rdad015.

Anton Korinek and Joseph E. Stiglitz. Steering technological progress. SSRN Scholarly Paper 5279584, Social Science Research Network, May 2025. URL <https://papers.ssrn.com/abstract=5279584>.

Shannon Liao. Amazon warehouse workers skip bathroom breaks to keep their jobs, says report. *The Verge*, 2018. URL <https://www.theverge.com/2018/4/16/17243026/amazon-warehouse-jobs-worker-conditions-bathroom-breaks>.

J. C. R. Licklider. Man-computer symbiosis. *IRE Transactions on Human Factors in Electronics*, HFE-1:4–11, 1960.

Jeffrey Lin. Technological adaptation, cities, and new work. *The Review of Economics and Statistics*, 93(2):554–574, 2011.

Luona Lin and Kim Parker. U.S. workers are more worried than hopeful about future AI use in the workplace. Technical report, Pew Research Center, 2025. URL https://www.pewresearch.org/wp-content/uploads/sites/20/2025/02/ST_2025.2.25_AI-Workers_REPORT.pdf.

- Joel Mokyr, Chris Vickers, and Nicolas L. Ziebarth. The history of technological anxiety and the future of economic growth: Is this time different? *Journal of Economic Perspectives*, 29(3):31–50, 2015. doi: 10.1257/jep.29.3.31.
- Emilie Munson. 95 percent of public-school classrooms have Wi-Fi. *Government Technology*, December 2023. URL <https://www.govtech.com/education/k-12/95-percent-of-public-school-classrooms-have-wi-fi>.
- Arvind Narayanan and Sayash Kapoor. AI as normal technology. *Knight First Amendment Institute at Columbia University*, 2025. URL <http://knightcolumbia.org/content/ai-as-normal-technology>.
- National Science Board. Discovery: R&D activity and research publications. Technical Report NSB-2025-7, 2025. URL <https://nces.nsf.gov/pubs/nsb20257>.
- Office of the National Coordinator for Health Information Technology. Non-federal acute care hospital electronic health record adoption, September 2017. URL <https://www.healthit.gov/data/quickstats/non-federal-acute-care-hospital-electronic-health-record-adoption>. Health IT Quick-Stat #47.
- OpenAI. Openai charter, n.d. URL <https://openai.com/charter/>. Accessed January 5, 2026.
- Organisation for Economic Co-operation and Development. Education at a glance 2025: United States. Technical report, OECD Publishing, 2025. URL https://www.oecd.org/en/publications/education-at-a-glance-2025_1a3543e2-en/united-states_784df67f-en.html.
- Our World in Data. Labor share of gross domestic product (gdp), 2020, 2020. URL <https://ourworldindata.org/grapher/labor-share-of-gdp?tab=discrete-bar&time=latest>.
- Peter G. Peterson Foundation. How does government healthcare spending differ from private insurance?, 2025. URL <https://www.pgpf.org/article/how-does-government-healthcare-spending-differ-from-private-insurance/>.
- Thomas Piketty, Emmanuel Saez, and Gabriel Zucman. Distributional national accounts: Methods and estimates for the United States. *The Quarterly Journal of Economics*, 133(2):553–609, 2018. doi: 10.1093/qje/qjx043.
- Snehil Prakash. Amazon employees: How many people work at Amazon? (2024 update). *How to Buy SAAS*, March 2025. URL <https://www.howtobuysaas.com/blog/amazon-employees/>.
- Pascual Restrepo. Automation: Theory, evidence, and outlook. *Annual Review of Economics*, 16:1–25, 2024. doi: 10.1146/annurev-economics-090523-113355.

- Pascual Restrepo. We won't be missed: Work and growth in the AGI world. In Ajay K. Agrawal, Erik Brynjolfsson, and Anton Korinek, editors, *The economics of transformative AI*. University of Chicago Press, 2025. URL <https://www.nber.org/books-and-chapters/economics-transformative-ai/we-wont-be-missed-work-and-growth-agi-world>. Preprint Chapter.
- Joanna Stern. We let AI run our office vending machine. It lost hundreds of dollars. *Wall Street Journal*, December 2025. URL <https://www.wsj.com/tech/ai/anthropic-claude-ai-vending-machine-agent-b7e84e34>.
- Daniel Sullivan and Till von Wachter. Job displacement and mortality: An analysis using administrative data. *The Quarterly Journal of Economics*, 124(3):1265–1306, 2009. doi: 10.1162/qjec.2009.124.3.1265.
- The Economist. China's 200m gig workers are a warning for the world. *The Economist*, September 2025. URL <https://www.economist.com/leaders/2025/09/18/chinas-200m-gig-workers-are-a-warning-for-the-world>.
- United States Patent and Trademark Office. New PE2E search tool using AI search features, 2022. URL <https://www.uspto.gov/web/offices/com/sol/og/2022/week02/TOC.htm#ref10>.
- United States Patent and Trademark Office. Artificial Intelligence strategy. Technical report, 2025. URL <https://cdn.patentlyo.com/media/2025/01/uspto-ai-strategy1.pdf>.
- United States Patent and Trademark Office. New artificial intelligence functionality in pe2e search, n.d. URL <https://www.uspto.gov/sites/default/files/documents/ai-sim-search.pdf>.
- U.S. Networking and Information Technology Research and Development Program. Artificial intelligence R&D investments: Fiscal year 2019–fiscal year 2025, n.d. URL <https://www.nitrd.gov/ai-rd-investments/>. Accessed January 5, 2026.
- Michelle Vaccaro, Abdullah Almaatouq, and Thomas Malone. When combinations of humans and AI are useful: A systematic review and meta-analysis. *Nature Human Behaviour*, 8(12):2293–2303, 2024. doi: 10.1038/s41562-024-02024-1.
- Sean Wang and Samuel Young. Unionization, employer opposition, and establishment closure. Technical report, 2024. URL https://sammygyoung.github.io/syoung_site/unionsClosure.pdf.
- Karen Weise. Amazon plans to replace more than half a million jobs with robots. *The New York Times*, October 2025. URL <https://www.nytimes.com/2025/10/21/technology/inside-amazons-plans-to-replace-workers-with-robots.html>.
- Wikipedia. Delta techops, 2025. URL https://en.wikipedia.org/w/index.php?title=Delta_TechOps&oldid=1298998991.

Accessed July 6, 2025.

William Julius Wilson. *When work disappears: The world of the new urban poor*. Vintage, 2011.