

NBER WORKING PAPER SERIES

APPLYING AI TO REBUILD MIDDLE CLASS JOBS

David Autor

Working Paper 32140

<http://www.nber.org/papers/w32140>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

February 2024

This article was drafted for NOEMA Magazine and will appear as “How AI Could Help Rebuild The Middle Class.” It draws on collaborative work with Daron Acemoglu, Caroline Chin, Joe Hazell, Simon Johnson, Pascual Restrepo, Anna Salomons, and Bryan Seegmiller, as well as my solo work. I also thank my co-panelists on the National Academy of Science and Engineering’s panel “Automation and the U.S. Workforce: An Update,” Erik Brynjolfsson, John Haltiwanger, Lawrence Katz, Tom Mitchell, Michael Strain, and Laura Tyson, who influenced my thinking on this subject. I am grateful to Tami Abdollah, Gavin Alcott, Sherry Autor, Lauren Fahey, Amy Finkelstein, Kathleen Miles, Anna Salomons, and Marika Tatsutani for comments that improved the article. I acknowledge support from Google, the William and Flora Hewlett Foundation, the NOMIS Foundation, and the Smith Richardson Foundation. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by David Autor. 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.

Applying AI to Rebuild Middle Class Jobs  
David Autor  
NBER Working Paper No. 32140  
February 2024  
JEL No. J01,J2,N30,O14

### **ABSTRACT**

While the utopian vision of the current Information Age was that computerization would flatten economic hierarchies by democratizing information, the opposite has occurred. Information, it turns out, is merely an input into a more consequential economic function, decision-making, which is the province of elite experts. The unique opportunity that AI offers to the labor market is to extend the relevance, reach, and value of human expertise. Because of AI's capacity to weave information and rules with acquired experience to support decision-making, it can be applied to enable a larger set of workers possessing complementary knowledge to perform some of the higher-stakes decision-making tasks that are currently arrogated to elite experts, e.g., medical care to doctors, document production to lawyers, software coding to computer engineers, and undergraduate education to professors. My thesis is not a forecast but an argument about what is possible: AI, if used well, can assist with restoring the middle-skill, middle-class heart of the US labor market that has been hollowed out by automation and globalization.

David Autor  
Department of Economics, E52-438  
Massachusetts Institute of Technology  
77 Massachusetts Avenue  
Cambridge, MA 02139  
and NBER  
dautor@mit.edu

NOEMA Magazine is available at  
<https://noemamag.com/how-ai-could-help-rebuild-the-middle-class/>

This article was drafted for NOEMA Magazine and will appear as  
“How AI Could Help Rebuild The Middle Class”

## Introduction

In a recent interview with U.K. Prime Minister Rishi Sunak, Elon Musk proclaimed artificial intelligence to be “the most disruptive force in history,” and noted that “there will come a point where no job is needed” (Henshall, 2023). Last year, AI godfather Geoffrey Hinton advised people to “get a job in plumbing” (Boran, 2023).

The message seems clear: The future of work, for many of us, is imperiled. A recent Gallup poll found that 75% of U.S. adults believe AI will lead to fewer jobs (Marken, 2023).

But this fear is misplaced.

The industrialized world is awash in jobs, and it’s going to stay that way. Four years after the Covid pandemic’s onset, the U.S. unemployment rate has fallen back to its pre-Covid nadir while total employment has risen to nearly three million above its pre-Covid peak (US Bureau of Labor Statistics, 2023). Due to plummeting birth rates and a cratering labor force, a comparable labor shortage is unfolding across the industrialized world (including in China) (Jones, 2020).

This is not a prediction, it’s a demographic fact. All the people who will turn thirty in the year 2053 have already been born and we cannot make more of them. Barring a massive change in immigration policy, the U.S. and other rich countries will run out of workers before we run out of jobs (Varian, 2020).

AI will change the labor market, but not in the way Musk and Hinton believe. Instead, it will reshape the value and nature of human expertise. Defining terms, expertise refers to the knowledge or competency required to accomplish a particular task like taking vital signs, coding an app or catering a meal. Expertise commands a market premium if it is both necessary for accomplishing an objective and relatively scarce. To paraphrase the character Syndrome in the movie “The Incredibles,” if everyone is expert, no one is expert.

Expertise is the primary source of labor’s value in the US and other industrialized countries. Jobs that require little training or certification, such as restaurant servers, janitors, manual laborers and (even) childcare workers, are typically found at the bottom of the wage ladder.

Consider the occupations of air traffic controller and crossing guard. In broad strokes, these are the same job: making rapid-fire, life-or-death decisions to avert collisions between passengers in vehicles and bystanders. But air traffic controllers were paid a median annual salary of \$132,250 in 2022, or nearly four times the \$33,380 median annual pay of crossing guards.<sup>1</sup>

---

<sup>1</sup>On crossing guards see Bureau of Labor Statistics and US Department of Labor (2023b). On air traffic controllers

The reason is expertise. Becoming an air traffic controller requires years of education and on-the-job apprenticeship — it is a scarce skill. Conversely, in most U.S. states, working as a crossing guard requires no formal training, specialized expertise or certification. An urgent need for more crossing guards could be filled by most air traffic controllers but the reverse would not be true.

Expertise is in constant flux. Forms that once commanded a substantial market premium — farriery, typesetting, fur-trapping, spell-checking — are all now either antiquated or automated. Simultaneously, many of the most highly paid jobs in industrialized economies — oncologists, software engineers, patent lawyers, therapists, movie stars — did not exist until specific technological or social innovations created a need for them. But the areas of expertise that will be eclipsed or gain currency shifts across technological epochs. The era of Artificial Intelligence heralds another such transformation.

The utopian vision of our Information Age was that computers would flatten economic hierarchies by democratizing information. In 2005, Marc Andreessen, co-founder of Netscape, told the New York Times’s Thomas Friedman that “today, the most profound thing to me is the fact that a 14-year-old in Romania or Bangalore or the [former] Soviet Union or Vietnam has all the information, all the tools, all the software easily available to apply knowledge however they want” (Friedman, 2005).

But the opposite of this vision has transpired.

Information, it turns out, is merely an input for a more consequential economic function, decision-making, which is the province of elite experts — typically the minority of U.S. adults who hold college or graduate degrees. By making information and calculation cheap and abundant, computerization catalyzed an unprecedented concentration of decision-making power, and accompanying resources, among elite experts (Deming, 2021).

Simultaneously, it automated away a broad middle-skill stratum of jobs in administrative support, clerical and blue-collar production occupations. Meanwhile, lacking better opportunities, the 60% of adults without a bachelor’s degree have been relegated to non-expert, low-paid service jobs.

The unique opportunity that AI offers humanity is to push back against the process started by computerization — to extend the relevance, reach and value of human expertise for a larger set of workers. Because artificial intelligence can weave information and rules with acquired experience to support decision-making, it can enable a larger set of workers equipped with necessary foundational training to perform higher-stakes decision-making tasks currently arrogated to elite experts, such as doctors, lawyers, software engineers and college professors. In essence, AI used well can assist with restoring the middle-skill, middle-class heart of the U.S. labor market that has been hollowed out by automation and globalization.

---

see Bureau of Labor Statistics and US Department of Labor (2023a)

While one may worry that AI will simply render expertise redundant and experts superfluous, history and economic logic suggest otherwise. AI is a tool, like a calculator or a chainsaw, and tools generally aren't substitutes for expertise but rather levers for its application.

By shortening the distance from intention to result, tools enable workers with proper training and judgment to accomplish tasks that were previously time-consuming, failure-prone or infeasible. Conversely, tools are useless at best — and hazardous at worst — to those lacking relevant training and experience. A pneumatic nail gun is an indispensable time-saver for a roofer and a looming impalement hazard for a home hobbyist.

For workers with foundational training and experience, AI can help to leverage expertise so they can do higher-value work. AI will certainly also automate existing work, rendering certain existing areas of expertise irrelevant. It will further instantiate new human capabilities, new goods and services that create demand for expertise we have yet to foresee.

My thesis is not a forecast but a claim about what is attainable. AI will not decide how AI is used, and its constructive and destructive applications are boundless. The erroneous assumption that the future is determined by technological inevitabilities — what Shoshana Zuboff terms inevitabilism — deprives citizens of agency in making, or even recognizing, the collective decisions that will shape the future (Zuboff, 2019). As Simone de Beauvoir wrote, “Fate triumphs as soon as we believe in it.”<sup>2</sup> AI offers vast tools for augmenting workers and enhancing work. We must master those tools and make them work for us.<sup>3</sup>

## From Artisanal to Mass Expertise

Most “experts” of our era would be at a loss if teleported back to the 18th century. Prior to the Industrial Revolution, goods were handmade by skilled artisans: wagon wheels by wheelwrights; clothing by tailors; shoes by cobblers; timepieces by clockmakers; firearms by blacksmiths. Artisans spent years acquiring at least two broad forms of expertise: procedural expertise, meaning following highly practiced steps to produce an outcome; and expert judgment, meaning adapting those procedures to variable instances.

If a blacksmith set out to make two muskets from the same design, not a single part belonging to one would be interchangeable with the other. Each part would be filed, honed and polished to fit precisely in the musket for which it was intended. Few experts of our era could do such work and fewer still could do so using the primitive tools of that time.

Although artisanal expertise was revered, its value was ultimately decimated by the rise of mass production in the 18th and 19th centuries (Hounshell, 1984). Mass production meant breaking the

---

<sup>2</sup>The original is, “La fatalité triomphe dès que l’on croit en elle” (de Beauvoir, 1952)

<sup>3</sup>Though not covered in this essay, see Acemoglu and Johnson (2023) and Acemoglu, Autor and Johnson (2023) for practical ideas on how to target these goals.

complex work of artisans into discreet, self-contained and often quite simple steps that could be carried out mechanistically by a team of production workers, aided by machinery and overseen by managers with higher education levels.

Mass production was vastly more productive than artisanal work, but conditions for rank-and-file workers were typically hazardous and grueling, requiring no specialized expertise beyond a willingness to labor under punishing conditions for extremely low pay.

Whereas skilled artisans were almost necessarily adult men — reflecting the years of apprenticeship required to master their trades as well as restrictive gender norms — early factories made abundant use of children and unmarried women. The skilled British weavers and textile workers who rose up to protest mechanization in the 19th century — the eponymous Luddites — are frequently derided for their supposed naive fear of technology.

But this fear was justified. As economic historian Joel Mokyr and his colleagues wrote in 2015, “the handloom weavers and frame knitters with their little workshops were quite rapidly wiped out by factories after 1815” (Mokyr et al., 2015; Bythell, 1969). Even as industrial-era innovations spurred a surge in productivity, it was five decades before working-class living standards began to rise.

As the tools, processes and products of modern industry gained sophistication, demand for a new form of worker expertise — “mass expertise” — burgeoned (Goldin and Katz, 1998; Buyst et al., 2018). Workers operating and maintaining complex equipment required training and experience in machining, fitting, welding, processing chemicals, handling textiles, dyeing and calibrating precision instruments, etc. Away from the factory floor, telephone operators, typists, bookkeepers and inventory clerks, served as information conduits — the information technology of their era.

Much of this newly demanded expertise was novel. There had been no demand for electricians until electricity found industrial and consumer uses. There were no skilled machinists prior to the invention of the machines that they operated. And there were no telephone operators prior to the construction of the telephone network. To master the tools, rules and exacting requirements of this work, workers frequently needed literacy and numeracy.

Not by coincidence, a large and growing fraction of the U.S. workforce was newly endowed with a high school diploma, meaning that such skills were increasingly available and substantially rewarded. This virtuous combination of rising industrial productivity and rising demand for worker expertise helped build a new middle class in industrializing countries, one that could afford luxuries such as full wardrobes, factory-made household goods and new industrial products like electric toasters and irons.

Unlike the artisans who preceded them, however, expert judgment was not necessarily needed — or even tolerated — among the “mass expert” workers populating offices and assembly lines. As

the management guru of the mass prediction era, Fredrick Winslow Taylor, wrote in 1911, “The work of every workman is fully planned out by the management at least one day in advance, and each man receives in most cases complete written instructions, describing in detail the task which he is to accomplish, as well as the means to be used in doing the work” (Taylor, 1911).

As a result, the narrow procedural content of mass expert work, with its requirement that workers follow rules but exercise little discretion, was perhaps uniquely vulnerable to technological displacement in the era that followed.

## **From Mass Expertise to Elite Expertise in the Information Age**

Stemming from the innovations pioneered during World War II, the Computer Era (AKA the Information Age) ultimately extinguished much of the demand for mass expertise that the Industrial Revolution had fostered. The unique power of the digital computer, relative to all technologies that preceded it, was its ability to cheaply, reliably and rapidly execute cognitive and manual tasks encoded in explicit, deterministic rules, i.e., what economists called “routine tasks” and what software engineers call programs (Autor et al., 2003).

This description might seem prosaic: don’t all machines simply follow deterministic rules? At one level, yes: Machines do what they’re built to do unless they are malfunctioning. But at another level, no. Distinct from prior mechanical devices that merely performed specific physical tasks, computers are symbolic processors that access, analyze and act upon abstract information. As Alan Turing proved in 1937, such machines can perform an infinite variety of tasks, provided that the task can be codified as a series of steps (more formally known as an algorithm) (Turing, 1937).

Prior to the computer era, there was essentially only one tool for symbolic processing: the human mind. The computer provided a second such tool, one with extraordinary capabilities and profound limitations. Prior to the computer era, workers who specialized in skilled office and production tasks were the embodiment of mass expertise.

As computing advanced, digital machines proved more proficient and much less expensive than workers in mastering tools and following rules. This eroded the value of mass expertise, just as the technologies of the Industrial Revolution previously eroded the value of artisanal expertise.

But not all tasks follow well-understood rules. As the philosopher Michael Polanyi observed in 1966, “We can know more than we can tell,” meaning that our tacit knowledge often exceeds our explicit formal understanding (Polanyi, 1966).

Making a persuasive argument, telling a joke, riding a bicycle or recognizing an adult’s face in a baby photograph, are subtle and complex undertakings that people accomplish without formal understanding of what they are doing — and often with little effort.

Mastery of these so-called “non-routine” tasks is not attained by learning the rules but instead through learning by doing. A child doesn’t need to study the physics of gyroscopes to learn to ride a bicycle — simple trial and error will suffice. But prior to AI, a programmer would have to specify all relevant steps, branches and exceptions for a robot to ride a bicycle. This observation — that there exist many tasks that human beings intuitively understand how to perform but whose rules and procedures they cannot verbalize — is often referred to as Polanyi’s Paradox .

Because many high-paid jobs are intensive in non-routine tasks, Polanyi’s Paradox proved a major constraint on what work traditional computers could do (Autor, 2014). Managers, professionals and technical workers are regularly called upon to exercise judgment (not rules) on one-off, high-stakes cases: choosing a treatment plan for an oncology patient, crafting a legal brief, leading a team or organization, designing a building, engineering a software product or safely landing a plane in dangerous conditions. Knowledge of rules is necessary but not sufficient for such cases.

Akin to the artisans of the pre-Industrial era, modern elite experts such as doctors, architects, pilots, electricians and educators, combine procedural knowledge with expert judgment and, frequently, creativity, to tackle specific, high stakes and often uncertain cases. Also akin to artisans, professionals develop expert judgment by spending years in supervised practice like apprenticeships, though the term is rarely applied to white-collar professionals.

Even as computerization eroded the value of mass expertise, it bordered on a divine gift to workers engaged in elite expert work. Computers enabled professionals to spend less time acquiring and organizing information, and more time interpreting and applying that information — that is, engaged in actual decision-making. This augmented the accuracy, productivity and thoroughness of expert professional judgment, thus magnifying its value.

As computerization advanced, the earnings of workers with four-year college and especially graduate degrees like those in law, medicine and science and engineering, rose steeply. This was a double-edged sword, however: computers automated away the mass expertise of the non-elite workers on whom professionals used to rely.

Ironically, computerization proved just as consequential for those employed in non-expert work. Many of the lowest-paid jobs in industrialized countries are found in hands-on service occupations: food service, cleaning and janitorial services, security, and personal care.

Although these jobs demand dexterity, sightedness, simple communication skills and common sense —and are therefore non-routine tasks, ill-suited to computerization— they are nevertheless low paid because they require little expertise: Most able-bodied adults can do them with minimal training and certification.

Computers could not “do” this work, but they did increase the supply of workers competing for it. Workers who in earlier decades would have qualified for mass expertise jobs in clerical, ad-



ministrative and production occupations were instead shunted into non-routine, hands-on service occupations. This placed downward pressure on already low wages for this work (Acemoglu and Restrepo, 2022).

Thus, rather than catalyzing a new era of mass expertise as did the Industrial Revolution, computerization fed a four-decade-long trend of rising inequality.

## **Complementing Judgement: Expertise in the Artificial Intelligence Era**

Like the Industrial and Computer revolutions before it, Artificial Intelligence marks an inflection point in the economic value of human expertise. To appreciate why, consider what distinguishes AI from the computing era that we're now leaving behind. Pre-AI, computing's core capability was its faultless and nearly costless execution of routine, procedural tasks. Its Achilles' heel was its inability to master non-routine tasks requiring tacit knowledge. Artificial Intelligence's capabilities are precisely the inverse.

In a case of cosmic irony, AI is not trustworthy with facts and numbers — it does not respect rules. AI is, however, remarkably effective at acquiring tacit knowledge. Rather than relying on hard-coded procedures, AI learns by example, gains mastery without explicit instruction and acquires capabilities that it was not explicitly engineered to possess.

If a traditional computer program is akin to a classical performer playing only the notes on the sheet music, AI is more like a jazz musician — riffing on existing melodies, taking improvisational solos and humming new tunes. Like a human expert, AI can weave formal knowledge (rules) with acquired experience to make — or support — one-off, high-stakes decisions.

AI's capacity to depart from script, to improvise based on training and experience, enables it to engage in expert judgment — a capability that, until now, has fallen within the province of elite experts.

Though only in its infancy, this is a superpower. As AI's facility in expert judgment becomes more reliable, incisive and accessible in the years ahead, it will emerge as a near-ubiquitous presence in our working lives. Its primary role will be to advise, coach and alert decision-makers as they apply expert judgment. If that sounds far-fetched, observe that AI's decision-making prowess has already bled into the edges of our everyday lives.

When your email application proposes to complete your sentence, your smartwatch asks if you've taken a fall or your car nudges your steering wheel to re-center your vehicle in a lane, AI is supplying expert judgment to interpret your intentions and guide your actions.

The stakes of most of these decisions are inconsequential at present, unless you're sleeping at the

wheel of your Tesla, but they will rise as AI advances and takes on higher value assignments in our lives.

What does this epochal advance in machine capability imply for the future of human expertise? Despite its novelty, I believe that the implications of AI have a relevant parallel in economic history, but the parallel runs counter to the present.

Recall that the advent of pre-AI computing made the expert judgment of professional decision-makers more consequential and more valuable by speeding the task of acquiring and organizing information. Simultaneously, computerization devalued and displaced the procedural expertise that was the stock-in-trade of many middle-skill workers.

But imagine a technology that could invert this process: what would it look like? It would support and supplement judgment, thus enabling a larger set of non-elite workers to engage in high-stakes decision-making. It would simultaneously temper the monopoly power that doctors hold over medical care, lawyers over document production, software engineers over computer code, professors over undergraduate education, etc.

Artificial Intelligence is this inversion technology. By providing decision support in the form of real-time guidance and guardrails, AI could enable a larger set of workers possessing complementary knowledge to perform some of the higher-stakes decision-making tasks currently arrogated to elite experts like doctors, lawyers, coders and educators.<sup>4</sup> This would improve the quality of jobs for workers without college degrees, moderate earnings inequality, and — akin to what the Industrial Revolution did for consumer goods — lower the cost of key services such as healthcare, education and legal expertise.

Most people understand that mass production lowered the cost of consumer goods. The contemporary challenge is the high and rising price of essential services like healthcare, higher education and law, that are monopolized by guilds of highly educated experts.

Federal Reserve Bank economists Emily Dohrman and Bruce Fallick estimate that over the last four decades, the prices of healthcare and education rose by around 200 and 600%, respectively, relative to U.S. household incomes (Dohrman and Fallick, 2020). Contributing to these increases is the escalating cost of employing elite decision-makers. Such gains are arguably justified: expertise commands a substantial premium when it is both necessary and scarce.

But AI has the potential to bring these costs down by reducing scarcity — that is, by empowering more workers to do this expert work.

To make this argument concrete, consider an example that is not from the AI realm: the job of Nurse Practitioner (NP). NPs are Registered Nurses (RNs) who hold an additional master's degree

---

<sup>4</sup>A related argument is offered in Agrawal et al. (2023)

that certifies them to administer and interpret diagnostic tests, assess and diagnose patients, and prescribe medications — services that were once exclusively within the domain of physicians.

Between 2011 and 2022, NP employment nearly tripled to approximately 224,000, with employment projected to grow by around 40% over the next decade, well above the national average. In 2022, the median Nurse Practitioner earned an annual salary of \$125,900.<sup>5</sup>

NPs are elite decision-makers. Their work combines procedural expertise with expert judgment so they can confront one-off patient cases where the stakes for judicious decision-making are extraordinarily high.

What makes the NP occupation relevant here is that it offers an uncommonly large-scale case where high-stakes professional tasks — diagnosing, treating and prescribing — have been reallocated (or co-assigned) from the most elite professional workers (MDs) to another set of professionals (NPs) with somewhat less elite (though still substantial) formal expertise and training.

What made this sharing of elite decision-making rights possible? The primary answer is institutional. In the early 1960s, a group of nurses and doctors recognized a growing shortage of primary care physicians, perceived that the skills of registered nurses were underused and pioneered a new medical occupation to address both shortfalls (Asubonteng et al., 1995).

This required launching new training programs, developing a certification regime and winning a change in the scope of medical practice regulations after a hard-fought (and perhaps never-ending) battle with the physicians’ primary lobbying arm, the American Medical Association.

A complementary answer to this origin story is that information technology, combined with improved training, facilitated this new division of labor. Speaking to the critical role that information and computing technology (ICT) play in the NP occupation, a 2012 study reported, “ICT supported the advanced practice dimension of the NP role in two ways: availability and completeness of electronic patient information enhanced timeliness and quality of diagnostic and therapeutic decision-making, expediting patient access to appropriate care . . . [and] patient data sourced from a central database supported and improved quality of communication between health professionals” (Li et al., 2012).

To put it more simply: Electronic medical records and improved communication tools enabled NPs to make better decisions.

Moving forward, AI could ultimately supplement the expert judgment of NPs engaging in a broader scope of medical care tasks. And this point applies much more broadly. From contract law to

---

<sup>5</sup>It is challenging to find consistent time-series data on NP employment in the U.S. BLS reports an increase from 81,000 NPs in 2011 to 224,000 in 2022 (Bureau of Labor Statistics, 2022). The BLS Occupational outlook Handbook reports 300,000 Nurse Anesthetists, Nurse Midwives, and Nurse Practitioners as of 2023 (Occupational Outlook Handbook, 2023).

calculus instruction to catheterization, AI could potentially enable a larger set of workers to perform high-stakes expert tasks. It can do this by complementing their skills and supplementing their judgment.

Is there evidence to support this hypothesis? Three recent studies provide “proof-of-concept” examples. In a 2023 paper, research economist Sida Peng of Microsoft Research and his coauthors from GitHub Inc. and the MIT Sloan School of Management demonstrate that GitHub Copilot, a generative AI-based programming aid, can significantly increase programmer productivity (Peng et al., 2023). In a controlled experiment, the treatment group that was given access to this tool completed the required programming task about 56% faster than the control group without access to Copilot.

In a 2023 paper in *Science*, Massachusetts Institute of Technology graduate students Shakked Noy and Whitney Zhang conducted an online experiment focused on writing tasks. Among the marketers, grant writers, consultants, managers and other diverse professionals recruited to the study, half were randomly given access to (and encouraged to use) ChatGPT for writing tasks. The other half used conventional non-AI tools like word processors and search engines to do their tasks.

Noy and Zhang found significant improvements in the speed and quality of writing output among those assigned to the ChatGPT group. Time spent on the task fell across the board by 40%. Remarkably, the biggest quality improvements were concentrated at the bottom. The least effective writers in the ChatGPT group were about as effective as the median writer without ChatGPT — a huge quality jump (Noy and Zhang, 2023).

ChatGPT did not eliminate the role of expertise. While the best writers remained at the top of the heap using either set of tools, ChatGPT enabled the most capable to write faster and the less capable to write both faster and better — so the productivity gap between adequate and excellent writers shrank.

In a third recent paper, researchers Erik Brynjolfsson of Stanford University, and Danielle Li and Lindsey Raymond of MIT evaluated the use of generative AI tools that provide suggested responses to customer service agents (Brynjolfsson et al., 2023). They also estimated a significant, roughly 14% improvement in productivity, and akin to Noy and Zhang’s study, these gains were most pronounced among novice workers.

AI tools helped novice workers attain the capabilities of experienced agents in three months rather than 10. As an unexpected side-effect, quit rates among new agents also fell substantially, likely due to less customer rage directed at them in chat windows. Buffered by the AI tool, support workers experienced significantly less hostility both from their clients and also toward their clients.

In all three instances, AI tools supplemented expertise rather than displaced experts. This occurred

through a combination of automation and augmentation. The benefit of automation was paid in time savings. AI wrote the first draft of computer code, advertising copy and customer support responses. The benefit of augmentation accrued in quality.

Using AI, less skilled workers produced work closer in quality to that of their more experienced and skilled peers. But quality improved not simply because workers were asleep at the wheel while AI did the driving. Workers were called upon to apply their expertise and judgment to produce the final product of code, text or customer support, while also harnessing AI's suggestions.

A fourth recent paper offers an exception that arguably proves the rule. In an experiment that randomly allocated the availability of AI assistance to professional radiologists, Nikhil Agarwal of MIT and his coauthors have found that AI did not improve the quality of radiologists' diagnoses — even though AI's predictions were at least as accurate as nearly two-thirds of the doctors studied (Agarwal et al., 2023).

That's because doctors did not understand how to use the AI tool effectively; when AI offered confident predictions, doctors frequently overrode those predictions with their own. When AI offered uncertain predictions, doctors frequently overrode their own better predictions with those supplied by the machine.

The lesson is not that AI is contraindicated for doctors but rather that its broad, complementary use will require training and the acquisition of additional expertise. With this expertise in hand, radiologists will become both faster and more accurate at diagnosis, and hence more valuable as experts.

If AI unleashes a surge of productivity in radiology, customer service, software coding, copywriting and many other domains, won't that mean that we'll be left with fewer workers doing the jobs previously done by many? In some arenas, the opposite may well be true.

Demand for healthcare, education and computer code appears almost limitless — and will rise further if as expected AI brings down the costs of these services. But in other domains, yes, rapid productivity growth will erode employment (Bessen, 2019). In 1900, about 35% of U.S. employment was in agriculture (Katz and Margo, 2014). After a century of sustained productivity growth, that share in 2022 was around 1% — and not because we're eating less (US Department of Agriculture, 2024).

But what's true about employment in a specific product or service has never been true of the economy writ large. When nearly 40% of U.S. workers were on farms, the fields of health and medical care, finance and insurance, and software and computing had barely germinated.

The majority of contemporary jobs are not remnants of historical occupations that have so far escaped automation (Autor et al., forthcoming). Instead, they are new job specialties that are

inextricably linked to specific technological innovations; they demand novel expertise that was unavailable or unimagined in earlier eras.

There were no air traffic controllers, electricians or gene editors until supporting innovations gave rise to the need for these specialized skill sets. Nor is technology the full story. Many expert personal service occupations, such as vegan chefs, college admissions consultants and personal trainers, owe their livelihoods not to specific innovations but to rising incomes, fluctuating fashions and shifting economic incentives. Innovation contributes to this by expanding the economic pie, thus allowing societies to call for richer slices.

Facing decades ahead of stagnating population growth and a rising share of citizens who are long past retirement, the challenge for the U.S. and the entire industrialized world is not a shortfall of work but a shortfall of workers. In the rapidly aging country of Japan, for example, the Financial Times reports that “Japanese retailers have shortened operating hours, installed avatars and hired foreign students to cope with the labour shortage” (Inagaki et al., 2024).

It would serve the U.S. and other industrialized countries well if AI-enabled more workers to use their expertise more effectively, thus boosting the share of high-productivity jobs while mitigating demographic labor market pressures.

## **Substitution versus Complementarity**

If AI can supply cheap expertise by the bucketful, won’t the remaining thimblefuls of human expertise be superfluous? I’ll answer with an analogy: YouTube. If you are handy at home repair or work in skilled trades, you probably spend time watching YouTube “how to” videos: how to replace a light switch, how to find a gas leak, how to tune up a snowblower, etc. Fifty-one percent of adult YouTube users report that the site is “very important” for “figuring out how to do things they haven’t done before,” according to a 2018 Pew Research study (Smith et al., 2018).

But for whom are these how-to videos useful? Not experts. They make the videos. What about amateurs? Let’s say that I wanted to replace the fuse box in my 19th-century home with a 20th-century circuit breaker panel. Assume, hypothetically, that I’ve never touched a pair of electrical pliers and I don’t own insulated gloves. But I have a free Saturday and there’s a Home Depot around the corner. Confidence high, I fire up one of the dozens to hundreds of YouTube how-to videos on this subject and get to work. Inevitably, but not immediately, I realize that my 19th-century fuse box is not quite like the one in the video. Whether I choose to reverse course or brazenly carry on, I face a palpable risk of a nasty shock or electrical fire.

That YouTube video was clearly not intended for me. To harness the free expertise on tap, I needed foundational expertise: procedural knowledge for handling high voltage circuits, expert judgment for problem-solving when the job went off script. With that expertise in hand, YouTube might

have been exactly what I needed.

My point: rather than making expertise unnecessary, tools often make it more valuable by extending its efficacy and scope. And the more powerful the tool, the higher the stakes. As Alexander Pope noted, “a little learning is a dang’rous thing” (Pope, 1711).

While AI is more than simply YouTube for white-collar professionals, its role in extending the capabilities of experts will be paramount. Most medical procedures, for example, follow a well-specified set of steps. But executing these steps requires hands-on practice and the tacitly acquired expert judgment that comes with it.

Plausibly, an experienced medical worker could, guided by AI, master a new medical device like the use of a new type of catheter, or carry out an unfamiliar procedure in an emergency. An untrained adult might also succeed in catheterizing a patient (or themselves) by following a “how to” video on YouTube. But when that procedure inevitably goes off script, someone with expert medical judgment better be on hand.

Artificial Intelligence will not in general enable untrained, non-expert workers to carry out high-stakes tasks, like catheterization. But it can enable workers with an appropriate foundation of expertise to level up. AI can extend the reach of expertise by building stories atop a good foundation and sound structure. Absent this footing, it is a structural hazard.

Am I overlooking the likelihood that AI-powered robots will soon be doing these jobs solo — with no human experts required? I don’t think so. AI will speed progress in robotics, but the era in which it is feasible and cost-effective to deploy robots to perform physically demanding tasks in unpredictable real-world environments, rather than in tightly controlled factory settings, remains distant.

If that sounds unduly pessimistic, consider how, despite tremendous investment and widespread pronouncements of imminent success, the leading technology companies of our era have faltered in delivering autonomous driving (Chafkin, 2022). Why? It’s not because operating a steering wheel, accelerator and brake pedal is a stretch for robots. It’s trivial. What remains profoundly challenging is interpreting and responding appropriately to a world of unpredictable pedestrians, ever-changing road hazards and inclement weather (Cano, 2023). Seen in this light, the cognitive and physical dexterity required to install a breaker box, prepare a meal or catharize a patient, appear awesome.

## **The Twilight of Expertise?**

One might object that I am merely describing the serene twilight of human expertise. Won’t AI do for human expertise what tractors did for ditch digging, assembly lines for artisanal expertise and calculators for long division — that is, automate them?

Although I doubt that most readers would prefer a world in which humanity must still pound tools out of wrought iron or do long division with pencil and paper, I recognize the concern. A future in which human labor has no economic value is, in my view, an ungovernable nightmare — though some guaranteed income aficionados will disagree. Regardless, the conclusion does not follow from the premise.

Innovation invariably provides new tools, and tools are often implements of automation. London cab drivers, for example, train for years to memorize all the streets of London (Brown, 2011; Lydall, 2023)— but smartphone navigation apps have made this hard-earned expertise technologically obsolete and economically superfluous (Berger et al., 2018).

Tools can and do encroach on their users' expertise. But the opposite is just as often true. Recall the air traffic controllers from earlier. Absent radar, GPS and two-way radios, these highly trained experts could do little more than squint at the sky. Similarly, the expertise of most longstanding occupations like doctors, builders and musicians, would be far less useful, and in many cases irrelevant, if they were deprived of the tools necessary to apply that expertise.

In economic parlance, navigation apps automated the expertise of London cabbies. But radar, GPS and two-way radios did the opposite for air traffic controllers. Innovation in this case did not automate, it created a new type of expert work.

If innovations were used exclusively for automation, we would have run out of work long ago. Instead, the industrialized world appears poised to run out of workers before it runs out of jobs. The likely reason is that the most important innovations have never been about automation. Automation did not, for example, give rise to airplanes, indoor plumbing, penicillin, CRISPR or television.

Rather than automating existing tasks, these innovations opened fundamentally new vistas of human possibility. Simultaneously, they generated new employment and new demands for expertise. There were no aircraft crews, household plumbers, geneticists or television actors until supporting innovations gave rise to the need for these specialized skill sets.

AI will automate the core tasks of some occupations, eliminate others and dramatically reshape some of those that remain. It will simultaneously instantiate new goods and services, generate new demands for expertise and open new possibilities for human advancement — though it is always difficult to predict what those will be.

These countervailing effects will create winners and losers, and the adjustment may be wrenching. There is no economic law that says that the forces of automation and new work creation are equal and offsetting; in fact, recent evidence indicates that the former is outpacing the latter (Acemoglu and Restrepo, 2022; Autor et al., forthcoming). Even if these opposing forces were to arm-wrestle one another to a draw, however, it is unlikely that workers whose expertise will be displaced by AI will be the same ones whose expertise is made newly valuable.



## A Scenario, Not a Forecast

History and scholarship demonstrate that the technologies that societies develop and the way that they use them — for exploitation or emancipation, for broadening prosperity or concentrating wealth — are determined foremost by the institutions in which they were created and the incentives under which they are deployed.

Scientific mastery of controlled nuclear fission in the 1940s enabled nations to produce both massively destructive weapons and near-carbon-free electricity generation plants. Eight decades later, countries have prioritized these technologies differently. North Korea possesses a phalanx of nuclear weapons but no civilian nuclear power plants. Japan, the only country against which offensive nuclear weapons have been used, possesses no nuclear weapons and dozens of civilian nuclear power plants.

Artificial Intelligence is far more malleable and broadly applicable than nuclear technology, and hence the range of both constructive and destructive uses is far wider. How AI is deployed, and who gains and loses out in the process, will depend upon the collective (and conflicting) choices of industry, governments, foreign nations, non-governmental organizations, universities, worker organizations and individuals.

The stakes are staggering, affecting not only economic efficiency but also income distribution, political power and civil rights. Some nations already use AI to heavily surveil their populations, squelch viewpoints that depart from official narratives and identify (and subsequently punish) dissidents — and they are rapidly exporting these capabilities to like-minded autocracies (Beraja et al., 2023; Mozur et al., 2022; Kroupe, 2022). In other settings, the same underlying AI technologies are used to advance medical drug discovery (including the development of Covid vaccines) (Qureshi et al., 2023; Dolgin, 2023), enable real-time translation of spoken languages and provide free, customized tutoring to struggling learners and precocious autodidacts.

AI poses a real risk to labor markets, but not that of a technologically jobless future. The risk is the devaluation of expertise. A future where humans supply only generic, undifferentiated labor is one where no one is an expert because everyone is an expert. In this world, labor is disposable and most wealth would accrue to owners of Artificial Intelligence patents. The political contours of such a world would be a hellscape: “WALL-E” meets “Mad Max.”

Remarkably, it is also the economic future that many AI visionaries seem to have in mind. For example, the charter of OpenAI, developer of ChatGPT and DALL-E, defines Artificial General Intelligence (AGI) as “highly autonomous systems that outperform humans at most economically valuable work” (OpenAI, 2018). In a 2023 New York Times bestselling book, the AI pioneer Mustafa Suleyman writes, “If the coming wave really is as general and wide-ranging as it appears, how will humans compete?” (Suleyman and Bhaskar, 2023).

The most charitable thing I can say about these ominous statements is that they are likely wrong — a flattening of the complexity of innovation into a single dimension of automation. Do these technological visionaries believe that Black and Decker tools make contractors' skills less valuable and that airplanes outperform their passengers? The latter question is of course nonsensical. Airplanes are not our competitors; we simply couldn't fly without them.

Replicating our existing capabilities, simply at greater speed and lower cost, is a minor achievement. The most valuable tools complement human capabilities and open new frontiers of possibility. The most pedestrian ones incrementally outperform existing tools.

My Maytag washing machine harnesses more computer processing power than the first Apollo mission so that I can start its spin cycle from anywhere on the planet. But that washing machine is never going to reach the moon. If AGI simply gives humanity a better washing machine rather than a new moonshot, AGI has not failed us, we have failed ourselves.

Amid a deluge of press reports on the impending AI robocalypse, one could easily fail to notice that the industrialized world is long on jobs and short on workers. The question is not whether we will have jobs — we will — but whether these will be the jobs we want.

For the lucky among us, work provides purpose, community and veneration. But the quality, dignity and respect of a substantial minority of jobs has eroded over the past four decades as computerization has marched onward and inequality has grown more prevalent.

The unique opportunity that AI offers humanity is to turn back this tide — to extend the relevance, reach and value of human expertise for a larger set of workers. Not only could this dampen earnings inequality and lower the costs of key services like healthcare and education, but it could also help restore the quality, stature and agency that has been lost to too many workers and jobs.

This alternative path is not an inevitable or intrinsic consequence of AI development. It is, however, technologically plausible, economically coherent and morally compelling. Recognizing this potential, we should ask not what AI will do to us, but what we want it to do for us.

## References

- Acemoglu, Daron and Pascual Restrepo**, “Tasks, Automation, and the Rise in U.S. Wage Inequality,” *Econometrica*, 2022, 90 (5), 1973–2016.
- **and Simon Johnson**, *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*, PublicAffairs, 2023.
- **, David Autor, and Simon Johnson**, “Can we Have Pro-Worker AI? Choosing a path of machines in service of minds,” Technical Report 123, CEPR, Paris & London October 2023.
- Agarwal, Nikhil, Alex Moehring, Pranav Rajpurkar, and Tobias Salz**, “Combining Human Expertise with Artificial Intelligence: Experimental Evidence from Radiology,” Technical Report w31422, National Bureau of Economic Research, Cambridge, MA July 2023.
- Agrawal, Ajay, Joshua S. Gans, and Avi Goldfarb**, “Do we want less automation?,” *Science*, July 2023, 381 (6654), 155–158.
- Asubonteng, Patrick, Karl J. McCleary, and George Munchus**, “Nurse practitioners in the USA – their past, present and future: some implications for the health care management delivery system,” *Health Manpower Management*, June 1995, 21 (3), 3–10.
- Autor, D. H., F. Levy, and R. J. Murnane**, “The Skill Content of Recent Technological Change: An Empirical Exploration,” *The Quarterly Journal of Economics*, November 2003, 118 (4), 1279–1333.
- Autor, David**, “Polanyi’s Paradox and the Shape of Employment Growth,” Technical Report w20485, National Bureau of Economic Research, Cambridge, MA August 2014.
- **, Caroline Chin, Anna Salomons, and Bryan Seegmiller**, “New frontiers: The origin and content of new work, 1940 – 2018,” *The Quarterly Journal of Economics*, forthcoming, forthcoming.
- Beraja, Martin, Andrew Kao, David Y Yang, and Noam Yuchtman**, “AI-tocracy,” *The Quarterly Journal of Economics*, 03 2023, 138 (3), 1349–1402.
- Berger, Thor, Chinchih Chen, and Carl Benedikt Frey**, “Drivers of disruption? Estimating the Uber effect,” *European Economic Review*, November 2018, 110, 197–210.
- Bessen, James**, “Automation and jobs: when technology boosts employment,” *Economic Policy*, October 2019, 34 (100), 589–626.
- Boran, Marie**, “Godfather of AI: Prepare for Artificial Intelligence before it’s smarter than us,” July 2023.
- Brown, Mark**, “How Driving a Taxi Changes London Cabbies’ Brains,” *Wired UK*, December 2011.
- Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond**, “Generative AI at Work,” Technical Report w31161, National Bureau of Economic Research, Cambridge, MA April 2023.
- Bureau of Labor Statistics**, “Employed full time: Wage and salary workers: Nurse practitioners

- occupations: 16 years and over,” <https://fred.stlouisfed.org/graph/?g=15br9> 2022.
- Bureau of Labor Statistics and US Department of Labor**, “Air Traffic Controllers,” <https://www.bls.gov/ooh/transportation-and-material-moving/air-traffic-controllers.htm> September 2023.
- , “Data for Occupations Not Covered in Detail,” <https://www.bls.gov/ooh/about/data-for-occupations-not-covered-in-detail.htm> September 2023.
- Buyst, Erik, Maarten Goos, and Anna Salomons**, “Job polarization: an historical perspective,” *Oxford Review of Economic Policy*, 2018, 34 (3), 461–474.
- Bythell, Duncan**, *The Handloom Weavers: A Study in the English Cotton Industry During the Industrial Revolution*, Cambridge University Press, 1969.
- Cano, Ricardo**, “Self-driving cars are causing mayhem on S.F. streets, officials say. Will their expansion be restricted?,” *San Francisco Chronicle*, January 2023.
- Chafkin, Max**, “Even After \$100 Billion, Self-Driving Cars Are Going Nowhere,” *Bloomberg*, 2022.
- de Beauvoir, Simone**, *America Day by Day*, London: G. Duckworth, 1952.
- Deming, David**, “The Growing Importance of Decision-Making on the Job,” Technical Report w28733, National Bureau of Economic Research, Cambridge, MA April 2021.
- Dohrman, Emily and Bruce C. Fallick**, “Is the Middle Class Worse Off Than It Used to Be?,” *Economic Commentary (Federal Reserve Bank of Cleveland)*, February 2020, pp. 1–6.
- Dolgin, Elie**, “‘Remarkable’ AI tool designs mRNA vaccines that are more potent and stable,” *Nature*, May 2023.
- Friedman, Thomas**, “It’s a Flat World, After All,” *The New York Times*, April 2005, p. 33.
- Goldin, Claudia and Lawrence F. Katz**, “The Origins of Technology-Skill Complementarity,” *The Quarterly Journal of Economics*, August 1998, 113 (3), 693–732.
- Henshall, Will**, “Elon Musk Tells Rishi Sunak AI Will Eliminate the Need for Jobs,” *Time Magazine*, November 2023.
- Hounshell, David A.**, *From the American system to mass production, 1800 - 1932: the development of manufacturing technology in the United States* number 4. In ‘Studies in industry and society.’, Baltimore: Johns Hopkins University Press, 1984.
- Inagaki, Kana, Leo Lewis, and David Keohane**, “Avatars, robots and AI: Japan turns to innovation to tackle labour crisis,” *The Financial Times*, January 2024.
- Jones, Kate**, “These countries are aging the fastest - here’s what it will mean,” *World Economic Forum*, February 2020.
- Katz, Lawrence F. and Robert Margo**, “Technical Change and the Relative Demand for Skilled Labor: The United States in Historical Perspective,” in Leah Bouston, Carola Frydman,

- and Robert Margo, eds., *Human Capital in History*, University of Chicago Press, 2014, pp. 15–57.
- Kroupe, Anastasiia**, “Russia Uses Facial Recognition to Hunt Down Draft Evaders,” *Human Rights Watch*, October 2022.
- Li, Julie, Johanna Westbrook, Joanne Callen, and Andrew Georgiou**, “The role of ICT in supporting disruptive innovation: a multi-site qualitative study of nurse practitioners in emergency departments,” *BMC Medical Informatics and Decision Making*, December 2012, *12* (1), 27.
- Lydall, Ross**, “Hundreds of black taxi driver sign-ups in London after ‘The Knowledge’ test made free,” *The Standard*, December 2023.
- Marken, Stephanie**, “Three in Four Americans Believe AI Will Reduce Jobs,” *Gallup*, September 2023.
- Mokyr, Joel, 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*, August 2015, *29* (3), 31–50.
- Mozur, Paul, Muye Xiao, and John Liu**, “‘An Invisible Cage’: How China Is Policing the Future,” *New York Times*, 2022.
- Noy, Shakked and Whitney Zhang**, “Experimental evidence on the productivity effects of generative artificial intelligence,” *Science*, July 2023, *381* (6654), 187–192.
- Occupational Outlook Handbook**, “Nurse Anesthetists, Nurse Midwives, and Nurse Practitioners,” <https://www.bls.gov/ooh/healthcare/nurse-anesthetists-nurse-midwives-and-nurse-practitioners.htm> September 2023.
- OpenAI**, “OpenAI Charter,” April 2018.
- Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer**, “The Impact of AI on Developer Productivity: Evidence from GitHub Copilot,” *arXiv*, 2023. Publisher: arXiv Version Number: 1.
- Polanyi, Michael**, *The Tacit Dimension*, Doubleday & Co, 1966.
- Pope, Alexander**, “An Essay on Criticism,” 1711.
- Qureshi, Rizwan, Muhammad Irfan, Taimoor Muzaffar Gondal, Sheheryar Khan, Jia Wu, Muhammad Usman Hadi, John Heymach, Xiuning Le, Hong Yan, and Tanvir Alam**, “AI in drug discovery and its clinical relevance,” *Heliyon*, July 2023, *9* (7), e17575.
- Smith, Aaron, Skye Toor, and Patric Van Kessel**, “Many Turn to YouTube for Children’s Content, News, How-To Lessons,” Technical Report, Pew Research Center November 2018.
- Suleyman, Mustafa and Michael Bhaskar**, *The Coming Wave: Technology, Power, and the Twenty-first Century’s Greatest Dilemma*, Crown, 2023.
- Taylor, Frederick Winslow**, *The Principles of Scientific Management*, Harper & Brothers, 1911.

**Turing, A. M.**, “Computability and -definability,” *Journal of Symbolic Logic*, December 1937, 2 (4), 153–163.

**US Bureau of Labor Statistics**, “Employment Level,” Retrived from <https://fred.stlouisfed.org/graph/?g=1eHul> 2023.

**US Department of Agriculture**, “Ag and Food Sectors and the Economy,” <https://www.ers.usda.gov/data-products/ag-and-food-statistics-charting-the-essentials/ag-and-food-sectors-and-the-economy/#:~:text=Agriculture%20and%20its%20related%20industries,percent%20of%20total%20U.S.%20employment> January 2024.

**Varian, Hal**, “Automation versus procreation (aka bots versus tots),” *VoxEU Columns*, March 2020.

**Zuboff, Shoshana**, *The age of surveillance capitalism: the fight for a human future at the new frontier of power*, PublicAffairs, 2019.