



The labor market impacts of technological change: from unbridled enthusiasm to qualified optimism to vast uncertainty

David Autor

May 2022

NBER WORKING PAPER SERIES

THE LABOR MARKET IMPACTS OF TECHNOLOGICAL CHANGE:
FROM UNBRIDLED ENTHUSIASM TO QUALIFIED OPTIMISM TO VAST UNCERTAINTY

David Autor

Blueprint Labs Discussion Paper #2022.08

NBER Working Paper 30074

<http://www.nber.org/papers/w30074>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

May 2022

This paper was prepared for publication by the Brookings Institution under its Global Forum on Democracy and Technology project. I thank Daron Acemoglu, Lauren Fahey, and Zia Qureshi for thoughtful comments that improved the paper. I acknowledge financial support from Google.org, the Hewlett Foundation, the Smith Richardson Foundation, and the Washington Center for Equitable Growth. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 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.

The Labor Market Impacts of Technological Change: From Unbridled Enthusiasm to Qualified Optimism to Vast Uncertainty

David Autor

NBER Working Paper No. 30074

May 2022

JEL No. J23,J24,O33

ABSTRACT

This review considers the evolution of economic thinking on the relationship between digital technology and inequality across four decades, encompassing four related but intellectually distinct paradigms, which I refer to as the education race, the task polarization model, the automation-reinstatement race, and the era of Artificial Intelligence uncertainty. The nuance of economic understanding has improved across these epochs. Yet, traditional economic optimism about the beneficent effects of technology for productivity and welfare has eroded as understanding has advanced. Given this intellectual trajectory, it would be natural to forecast an even darker horizon ahead. I refrain from doing so because forecasting the “consequences” of technological change treats the future as a fate to be divined rather than an expedition to be undertaken. I conclude by discussing opportunities and challenges that we collectively face in shaping this future.

David Autor

Department of Economics, E52-438

Massachusetts Institute of Technology

77 Massachusetts Avenue

Cambridge, MA 02139

and NBER

dautor@mit.edu

Introduction

Citizens in industrialized countries believe that digital technology is fostering inequality and that this problem is likely to grow worse in the decades ahead (Smith and Anderson, 2017; Wike and Stokes, 2018). Although public and expert opinions often diverge on economic questions, survey data confirm that academic economists share this worry. A 2017 Chicago Booth poll found that 35 to 40 percent of leading U.S. economists believe that robots and artificial intelligence are likely to substantially increase long-term unemployment rates.¹ What is the economic basis for this concern? In this review, I consider the evolution of economic thinking on the relationship between digital technology and inequality across four decades, encompassing four intellectually related but distinct paradigms.

I start from the premise that what workers earn in a market economy depends substantially, though not exclusively, on their productivity—that is, the value they produce through their labor. Their productivity depends in turn on two things: first, their capabilities (concretely, the tasks they can accomplish); and second, their scarcity. The fewer workers that are available to accomplish a given task and the more that employers need that task accomplished by workers (rather than by, for example, machines or algorithms), the higher is the workers' economic value and thus their potential earnings. In conventional terms, the skill premium depends upon the supply of skills and the demand for skills.

Stated in these terms, what is the role of technology—digital or otherwise—in determining wages and shaping wage inequality? The answer is not obvious, and the successive evolution of thinking on this topic reflects the subtlety of the question. I present four answers below, corresponding to four strands of thinking on this topic, and discuss the distinct implications of each. I refer to these four paradigms as the education race, the task polarization model, the automation-reinstatement race, and the era of Artificial Intelligence uncertainty. The nuance of economic understanding has improved across each of these epochs. Yet, traditional economic optimism about the beneficent effects of technology for productivity and welfare has eroded as understanding has advanced. Given this intellectual trajectory, it would be natural to forecast an even darker horizon ahead. I refrain from doing that, however, because forecasting the “consequences” of technological change treats the future as a fate to be divined rather than an expedition to be undertaken. I conclude by discussing the opportunities and the challenges we collectively face in shaping this future.

¹ See <https://www.igmchicago.org/surveys/robots-and-artificial-intelligence-2/>. European economists are somewhat less pessimistic, however. See <https://www.igmchicago.org/surveys/robots-and-artificial-intelligence/>

1 The Education Race

Perhaps the most influential conceptual frame for understanding how technology shapes wage inequality originates with a short article published in 1974 by Dutch economist and Nobel Laureate, Jan Tinbergen (Tinbergen, 1974).² Tinbergen was intrigued by the observation that the wages of Dutch workers with post-high school education (which he called ‘third-level’ education) had been rising over the course of many decades despite vast increases in their supply. This pattern is hard to rationalize in a standard competitive setting since it seemingly implies that the demand curve for skilled labor is *upward* sloping.

To interpret these facts, Tinbergen offered a simple but remarkably powerful analogy. Modern economies face an ongoing race between the demand for and supply of skill, with technological change propelling the demand curve outward and the educational system racing to push the supply curve outward to match it.³ In this telling, when the demand curve pulls ahead in the race, inequality between more and less-educated workers—college and non-college workers in the contemporary setting—rises, since more-educated workers are becoming relatively scarce. Conversely, when the supply of college-educated workers surges, as occurred during the 1970s, for example, when American men could defer the Vietnam draft by enrolling in college (Card and Lemieux, 2001b), earnings inequality between college and non-college workers falls. Notably, there is no “equilibrium” quantity of education that holds inequality constant in this framework. Rather, technologically advancing countries must keep raising educational attainment cohort by cohort to keep pace with the moving target of rising skill demands. Or, quoting Lewis Carroll’s Red Queen, “it takes all the running you can do, to keep in the same place.”

Tinbergen’s metaphor of a race between education and technology, now formalized mathematically, has proved remarkably powerful. A series of papers and books, commencing with Katz and Murphy (1992), demonstrates that the evolution of inequality between education groups (generally, college-educated versus non-college) in many advanced countries is remarkably well explained by two forces: steadily rising demand for college workers, who are needed to perform increasingly sophisticated and skill-intensive jobs (presumably, the

² Seminal work by Robert Solow in the 1950s demonstrated that technological progress was the central force behind rising aggregate productivity. But Solow did not consider inequality. ‘Labor’ is an undifferentiated commodity in the Solow-Swan model, meaning that wage inequality was not a meaningful construct in this model.

³ In Tinbergen’s words, there is a “‘race’ between the demand for skill—that is, demand for third-level manpower—driven by technological development and supply of it due to increased schooling.”

technological developments that Tinbergen had in mind); and booms and busts in the rate of college attendance among young adults that affect supply.⁴

Figure 1, reproduced from Autor (2014), illustrates the capacity of this simple model to rationalize the evolution of the U.S. college / high-school earnings premium over the nearly five decades between 1963–2012. The model can explain both why the college premium fell during the 1970s as the rate of college attainment was rising rapidly, and further, why the college premium surged in the 1980s when college attainment of younger cohorts of U.S. adults plateaued. In fact, this model can in broad brushstrokes explain the evolution of inequality between college and non-college workers in the United States over the course of nearly two centuries (cf. Autor et al. (2020b)).

Of course, the college vs. non-college earnings premium is only one component of wage inequality; most earnings inequality occurs among workers of the same education levels. The data show, however, that the growth of educational earnings gaps is the predominant contributor to rising earnings inequality over the last four decades. Specifically, Autor et al. (2020b) estimate that the growth of education-earnings differentials explains approximately 60% of the growth of overall earnings inequality between 1980 and 2017 and 40% of the growth between 2000 and 2017. Hence, if we can understand the causes of rising educational earnings inequality, we understand a lot about the sources of the overall rise in earnings inequality.

The empirical success of the education-race model raises a foundational question: what is it about technology that raises the demand for better-educated workers? The model does not directly address this question. Taken (too) literally, it portrays technological progress as an autonomous force that intrinsically makes highly educated workers more productive and hence more in demand. To be sure, researchers have added considerable nuance to this framework as they have applied it. For example, Goldin and Katz (1998, 2008) offer theory and detailed historical evidence that early industrial-age factories primarily demanded *less-skilled* workers. But as factories adopted continuous-process methods requiring sophisticated machinery, they increasingly demanded more educated workers with the expertise needed to operate these sophisticated factories.⁵ The education race model's simplicity is both a strength and a limitation. The model can explain much with little—specifically, the evolution of two centuries of educational inequality as a function of only two factors: changes in educational supply and an ongoing (though not directly measured) technologically-propelled increase in educational demand. The limitation is that the model lacks an underlying notion of *why* technology affects

⁴ This model is further developed, elaborated, and applied in Autor et al. (2020b, 2008, 1998); Card and Lemieux (2001a); Goldin and Katz (2008); Goldin et al. (2007); Goldin and Margo (1992); Katz and Autor (1999)

⁵ In a related vein, Krusell et al. (2000) argue that technological change became more skill-demanding when improvements in the quality-adjusted price of industrial equipment accelerated in the 1970s.

skill demand. Specifying this notion is left to successor models that build on Tinbergen’s foundation.

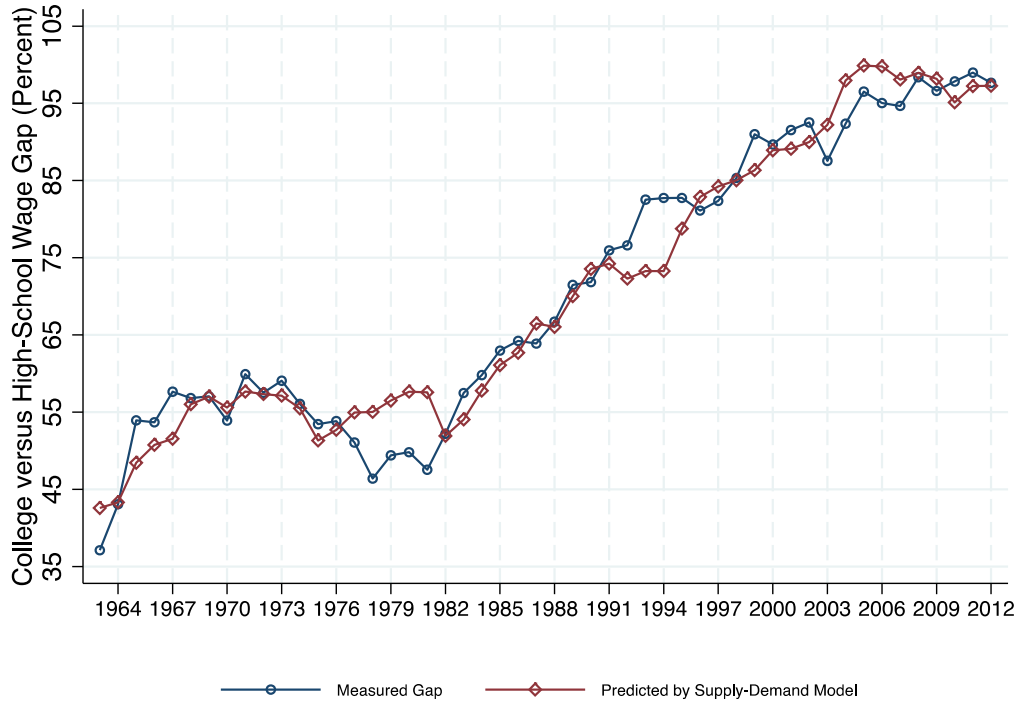


Figure 1: The supply of college graduates and the U.S. college/high school premium, 1963–2012

Source: Autor (2014).

Note: Figure uses March CPS data for earnings years 1963 to 2012. The series labeled “Measured Gap” is constructed by calculating the mean of the natural logarithm of weekly wages for college graduates and non-college graduates and plotting the (exponentiated) ratio of these means for each year. This calculation holds constant the labor market experience and gender composition within each education group. The series labeled “Predicted by Supply-Demand Model” plots the (exponentiated) predicted values from a regression of the log college/non-college wage gap on a quadratic polynomial in calendar years and the natural log of college/non-college relative supply. See Autor (2014) for details.

Beyond its simplicity, another feature of the education-race model has proven conceptually appealing but less empirically relevant. Technological change in the education-race model, *as conventionally applied*, affects labor demand only by *raising* (i.e., augmenting) the productivity of specific skill groups (e.g., college or non-college workers). In economic terms, this means that technological change in the simplest education-race model is *factor-augmenting*—it makes at least some workers better at the work that they do. The labor market impacts of factor-augmenting technological change are somewhere between benign and benevolent: benign

because no worker is made directly worse off (setting aside envy or other social externalities); and benevolent because, under conventional assumptions, *all* workers benefit from technological progress, at least to some degree.⁶ Thus, although technological change can raise inequality in the education-race framework (i.e., when demand surges ahead of supply), it does so by augmenting some workers more than others—which is not a terrible problem to have.

This implication of the model—that technological change at least weakly augments every worker’s productivity—is not well supported by the data. Figure 2, reproduced from Autor (2019), depicts the steep rise of earnings inequality by education group. Between 1979 and 2017, the real weekly earnings of full-time, full-year working men with a post-baccalaureate degree rose by 43 percent, and earnings for men with a 4-year degree but no graduate study rose by 12 percent. Conversely, real earnings *fell* substantially among men without a four-year degree: by 10 percent among men with some-college; by 21 percent among men with exactly a high school diploma; and by 25 percent among men without a high school diploma. (Real earnings rose among women of all educational levels, though the increases were very modest among the least educated women.) *If* the supply of non-college men and women had increased steeply in this period, these earnings declines could be consistent with the education-race model. But in reality, the share of working-age adults possessing less than a four-year degree dropped sharply.⁷ All else equal, this should have raised the relative wage of non-college workers, yet the opposite occurred.

Though not the standard approach, it is entirely possible to generalize the education-race model so that technological change can either augment or replace factors. Specifically, one can introduce factor-replacing technological change that reduces the real wages of non-college workers by reallocating tasks from non-college to college workers (or vice versa). The task polarization model, outlined below, provides a foundation for understanding when and why such task reallocation might occur.

⁶ Formally, all workers necessarily benefit so long as capital is elastically supplied and college and non-college workers are (at least) weakly substitutable for one another, meaning that when a skill group becomes more productive (e.g., due to technological augmentation), employers demand more of that group. Considerable evidence supports the assumption that college and non-college workers are substitutable in this sense. (Katz and Autor, 1999).

⁷ The share of labor hours supplied by workers with high school or lower education fell from more than 75 percent in 1963 to less than 40 percent in 2017. Conversely, the share of labor hours supplied by workers with a bachelor’s or post-college degree rose from less than 15 percent to more than 35 percent (Autor, 2019)).

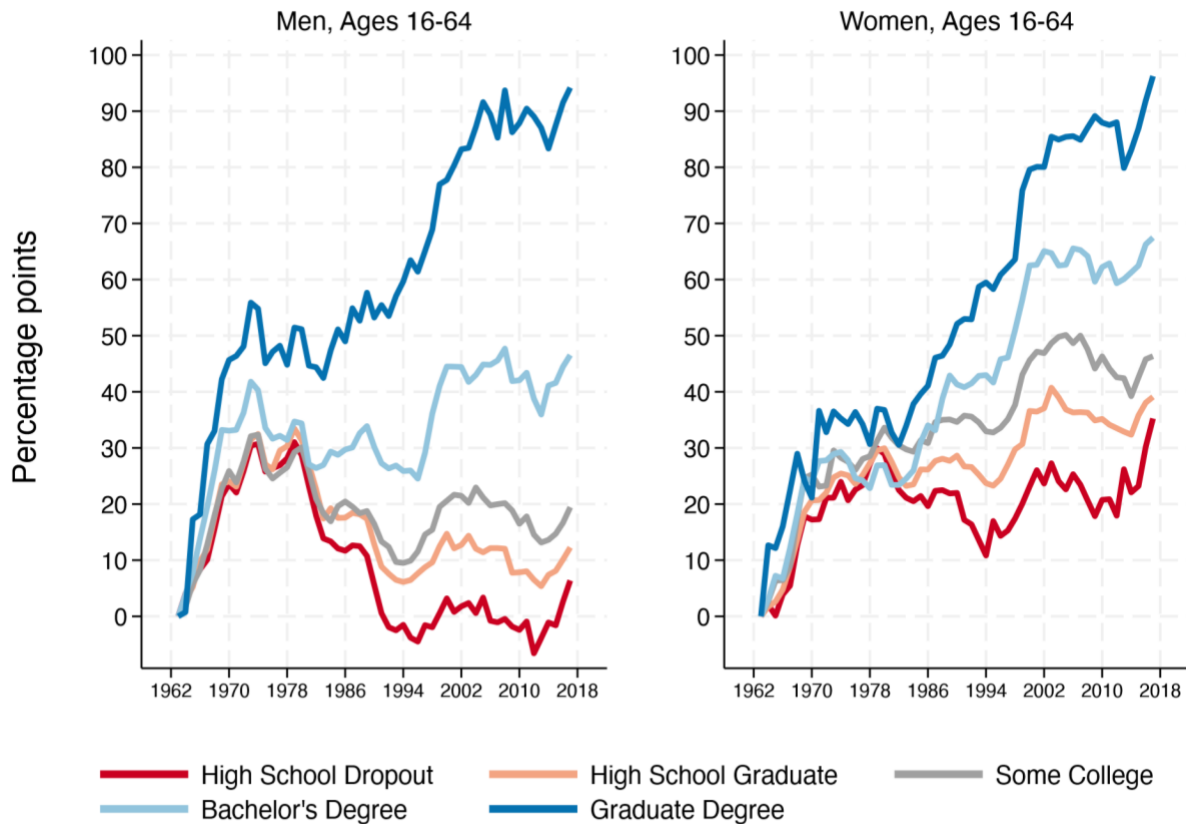


Figure 2: Cumulative percentage point changes in real mean weekly earnings of full-time, full-year workers ages 18–64, United States, 1963–2017

Source: *Autor (2019)*.

Note: Figure uses March Current Population Survey Annual Social and Economic Supplement data for earnings years 1963 to 2017. Series correspond to percentage point changes since 1963 in (composition-adjusted) exponentiated mean real (constant \$2019) log wages for each group, using data on full-time, full-year workers ages 16 to 64. Data are sorted into sex-education-experience groups of two sexes, five education categories (high school dropout, high school graduate, some college, college graduate, and post-college degree), and four potential experience categories (0–9, 10–19, 20–29, and 30–39 years). Educational categories are harmonized following the procedures in *Autor et al. (2008)*. Log weekly wages of full-time, full-year workers are regressed in each year separately by sex on dummy variables for four education categories, a quartic in experience, three region dummies, black and other race dummies, and interactions of the experience quartic with three broad education categories (high school graduate, some college, and college plus). The (composition-adjusted) mean log wage for each of the 40 groups in a given year is the predicted log wage from these regressions evaluated for whites, living in the mean geographic region, at the relevant experience level (5, 15, 25, or 35 years depending on the experience group). Mean log wages for broader groups in each year represent weighted averages of the relevant (composition-adjusted) cell means using a fixed set of weights, equal to the mean share of total hours worked by each group over 1963–2005. All earnings numbers are deflated by the chain-weighted (implicit) Personal Consumption Expenditure price deflator and exponentiated for plotting purposes. Earnings of less than 67 per week in 1982 dollars are dropped. Allocated earnings observations are excluded in earnings years 1967 forward using either family earnings allocation flags (1967–1974) or individual earnings allocation flags (1975 earnings year forward).

In short, while the rising wages of college-educated workers in the face of rising relative supply is consistent with the education-race model—corresponding to a case in which technology pushes demand outward faster than supply is rising—the substantial, sustained fall in the real earnings of non-college workers is less consistent with this model. Other factors aside from technology may be at play, of course, such as declining unionization, falling real minimum wages, or accelerating globalization. Nevertheless, I will argue below that technological change is at least partly responsible, but not in a form that is easily captured in the canonical education-race model.

2 The Task-Polarization Model

Building on this conceptual foundation, a subsequent literature takes up a central question that the education-race model leaves unanswered: *why* do recent waves of technology appear to complement more-educated workers (Acemoglu and Autor, 2011; Autor et al., 2003)? In answering this question, this research helps to explain why the real earnings of some skill groups have fallen, even while technological change has augmented the productivity and earnings of other skill groups. In short, this framework offers a more nuanced but less benign view on the effects of technological change on earnings levels, inequality, and the value of skills.

The starting point of the task-polarization model (‘task model’) is to conceptualize the process of accomplishing a job as performing a series of tasks. For example, the tasks that go into writing a research paper might include managing a research team, collecting data, developing and testing hypotheses, performing calculations, crafting a report, proofreading that report, and distributing it to recipients. The second step is to ask which tasks will be carried out by machines and which by workers. In the pre-digital era, most research and writing tasks would have been accomplished more or less manually with human labor plus books, calculators, typewriters, and postal mail. Human expertise would also be heavily applied to leading and managing teams, interpreting data, forming and testing hypotheses, and writing the report.

Computerization changes this picture by reallocating many of these tasks from human tasks to machine tasks—for example, collecting (machine-readable) data, performing calculations, proofreading, and distributing the report. Notice that in this new division of labor, computers accomplish a distinctive subset of tasks, those involving routine codifiable activities that can be fully described by a set of rules and procedures, encoded in software, and carried out by non-sentient machines. Tasks such as data gathering (from machine-readable sources), calculation, and certain types of error-checking are well-suited for computerization because they follow deterministic scripts. Conversely, it has proved far more challenging to program computers to lead teams, develop and test novel hypotheses, draw robust conclusions, and write compelling reports conveying the findings (though this is changing, more on this below). The simple reason is that these tasks are not well described by a tightly specified scripts that machines can faithfully

execute to achieve successful results—at least, not without substantial reliance on human expertise and judgment. Accordingly, such ‘non-routine’ tasks have primarily been performed by workers rather than machines—until recently (more below). Paired with computers, workers can focus their efforts primarily on the tasks that machines cannot accomplish, which opens the possibility for faster work, better work, or both.

This simple framing offers two refinements relative to the education-race model. First, it embraces the reality that automation directly *replaces* human labor in accomplishing a subset of tasks—something that does not happen in the canonical education-race model. An immediate implication is that workers whose most valuable skills are collecting data, performing calculations, proofreading documents, etc. are potentially made worse off because computers directly substitute for their skills. Concretely, because the real cost of symbolic processing (i.e., what computers do) has been falling by double-digits annually for decades (Nordhaus, 2007), what workers can now earn by carrying out these once well-remunerated but now fully automated information processing tasks is essentially zero.⁸

A second strength of the task framework is that it offers a plausible explanation for *why* computerization seems to complement more educated workers. Observe that in the paper-writing example above, many of the tasks that are *not* computerized would be considered high-skill tasks: leading a team, forming a hypothesis, crafting a paper, etc. These “non-routine cognitive” abstract-reasoning (expert judgment, creativity) and interpersonal (leadership, management) tasks have proven hard to automate because, simply put, we don’t know “the rules.” As the philosopher Michael Polanyi observed, “We know more than we can tell” (Polanyi and Sen, 1966), meaning that there are many things that we regularly accomplish—riding a bicycle, making a compelling argument, recognizing a current friend’s face in their baby photograph—that we understand tacitly but not explicitly how to do. People can achieve mastery through tacit knowledge because they learn by doing. A child doesn’t need to read up on the physics of gyroscopes to learn how to ride a bicycle—simple trial and error will do it. For a computer program to successfully accomplish a task, however, the computer programmer must usually specify *all* the relevant steps, branches, and exceptions in advance. For this reason, “non-routine” abstract-reasoning and interpersonal communication tasks have remained largely out of reach for machines (again, until recently).

The argument goes one step further: not only are abstract-reasoning and communication tasks not substituted by computers, they are generally *complemented*. The productivity and

⁸ Concretely, there is no positive price at which an employer would hire someone to add columns of numbers, route telephone calls between exchanges, or look up the current trading price of a group of stocks—yet, these tasks used to comprise many full-time jobs (see Feigenbaum and Gross (2020) on the elimination of telephone operators by mechanical and then digital switching).

earnings power of workers who specialize in abstract reasoning, expert judgment, and interpersonal interactions and leadership rises as the inputs into their work—information access, analysis, and communication—becomes less expensive and more productive. Thus, computerization increases the productivity of better-educated workers whose jobs rely on information, calculation, problem-solving, and communication, e.g., doctors, architects, researchers, and stock analysts. But this is a double-edged sword: computerization increases the productivity of highly educated workers by displacing the tasks of the middle-skill workers who in many cases previously provided these information-gathering, organizational, and calculation tasks (e.g., sales workers, office workers, administrative support workers, and assembly line production workers).

However, not all tasks that are hard to automate would be classified as high-skill tasks. Tasks such as waiting tables, cleaning rooms, picking and boxing items, or assisting elderly people to perform acts of daily living, require dexterity, sightedness, simple communications, and common sense, all of which draw on substantial reservoirs of tacit knowledge.⁹ Such tasks are commonly found in personal services jobs, e.g., food service, cleaning, security, entertainment, recreation, and personal care. Computerization has generally not substituted for workers in performing such jobs. But neither has it strongly complemented them. Rather, it leaves this work largely untouched, neither automating the central tasks of this job nor augmenting the workers doing it. Moreover, because a large fraction of adults can, with modest training, perform the core tasks of many non-routine manual jobs, such jobs will generally not pay high wages even when demand is rising, except when the labor market is very tight (as is currently the case).

There is now a vast literature testing the task framework empirically, extending it theoretically, and of course, critiquing it vigorously (Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2021; Autor and Dorn, 2013; Autor et al., 2006; Deming, 2017; Goos and Manning, 2007; Goos et al., 2009, 2014; Gregory et al., 2021; Harrigan et al., 2021; Levy and Murnane, 2004; and Michaels et al., 2014). A central implication of this framework, one that receives ample empirical support, is that across firms, industries, and countries, computerization spurs a ‘polarization’ of job growth into traditionally high-wage and traditionally low-wage occupations at the expense of the middle tier. We see this clearly in the U.S. data: at the high end of the labor market, a growing cadre of high-education, high-wage occupations offer strong career prospects, rising lifetime earnings, and significant employment security. At the other end, low-education, low-wage occupations, often in personal services, provide little economic security and limited career earnings growth. Traditional middle-tier jobs in production, operative, clerical and administrative support, and sales are in decline. Figure 3, reproduced from Autor (2019),

⁹ A hotel employee cleaning a guest room must determine which items are personal and which are trash. A soda can found on the floor is likely trash; a similarly situated perfume bottle likely fell there by accident.

documents this pattern for the United States. Figure 4 shows an analogous pattern in European data over a shorter time period (Goos et al., 2014).¹⁰

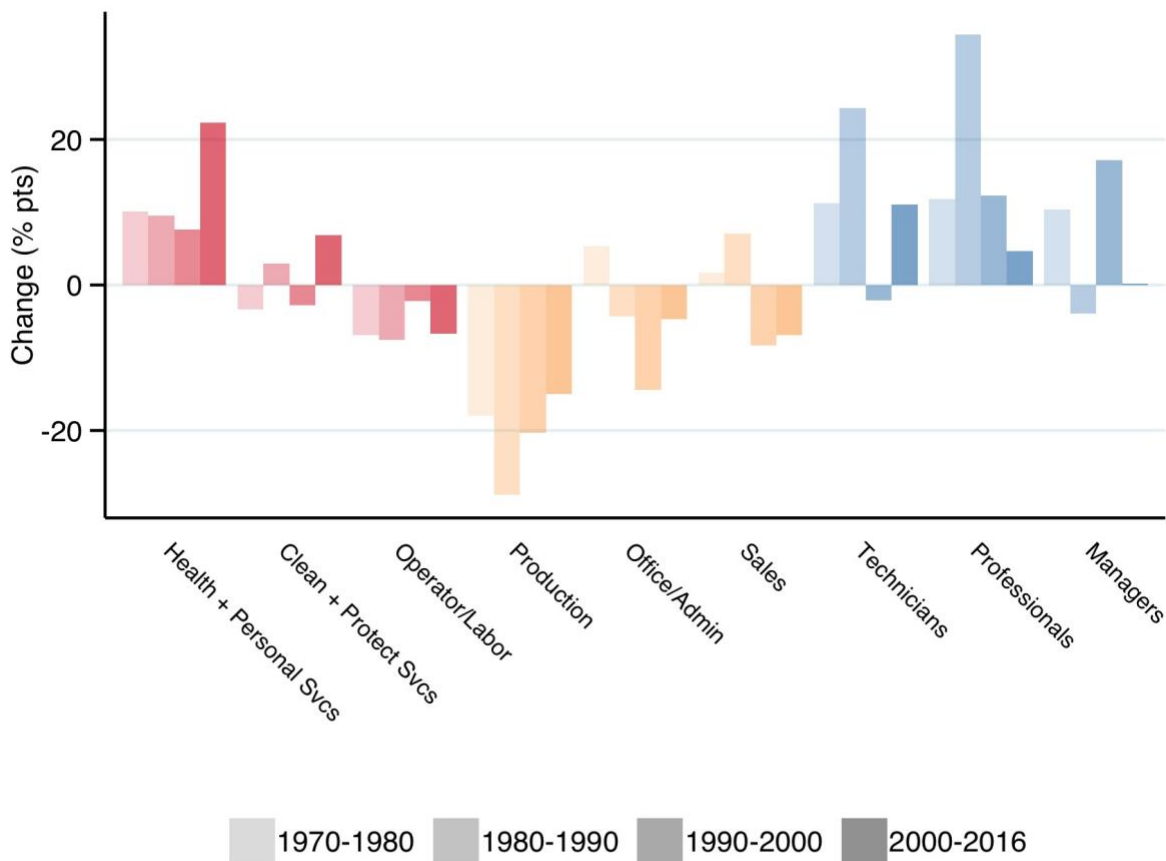


Figure 3: Percent changes in occupational employment shares among working-age adults, United States, 1970–2016

Source: Autor (2019)

Note: Data source is as in Figure 2. Sample consists of all persons aged 16 to 64 who reported having worked at least one week in the earnings years, excluding those in the military. For each individual, hours worked are the product of usual hours worked per week and the number of weeks worked last year. Individual hours worked are aggregated using CPS sampling weights. Occupational classifications are harmonized following Dorn (2009), and updated through 2017.

¹⁰ Polarization does not, however, describe the experience of developing countries, where the skill levels associated with different tasks are quite different and computerization is less pervasive and still relatively expensive in comparison with human labor (Maloney and Molina, 2016).

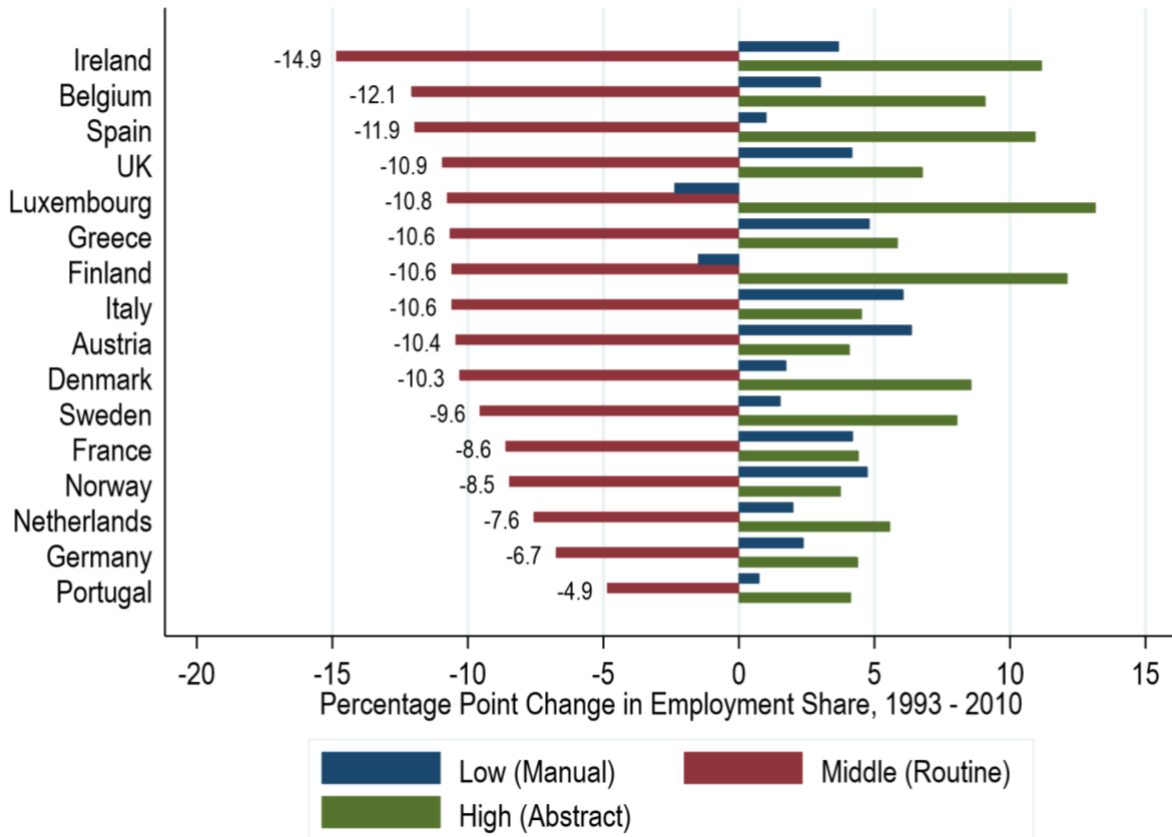


Figure 4: Change in occupational employment shares in European Union Countries, 1993–2010

Source: [Goos et al. \(2014\)](#).

The evidence on occupational change is clear. The implications of the task framework for wages are, however, more nuanced. For highly educated workers, those performing non-routine analytic and interpersonal tasks, the task framework unambiguously predicts higher earnings. By the same logic, one might surmise that wages in middle-skill routine-task-intensive occupations should fall while wages in lower-skill service occupations should remain unaffected. This can occur, but the prediction is ambiguous. The reason why is that when wages in middle-skill occupations fall, workers who would otherwise do those jobs will tend to enter previously lower-paid service occupations, thus placing downward pressure on wages in those occupations as well.¹¹ Thus, while the task model unambiguously predicts the U-shaped pattern of occupational growth seen in Figures 3 and 4, it is formally ambiguous as to whether this also leads to a U-shaped pattern of wage growth (Autor and Dorn, 2013; Böhm, 2020; Böhm et al., 2019).

¹¹ Some workers will also transition into higher-paid occupations. However, degree and credential requirements for these occupations (e.g., law degree, medical degree, engineering certification) will constrain rapid entry.

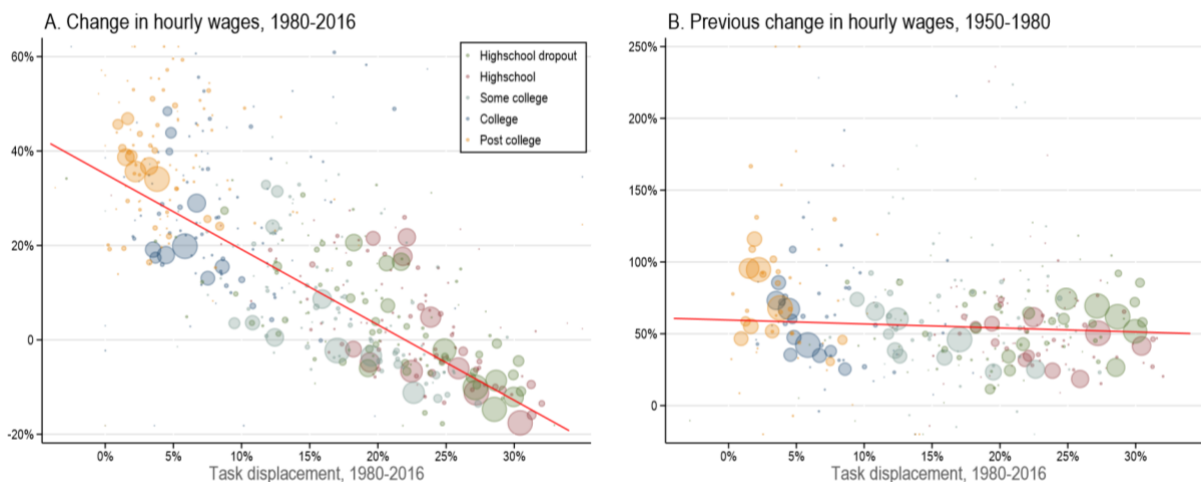


Figure 5: Exposure to task displacement and changes in real wages by demographic group, United States, 1980–2016 and 1950–1980

Source: [Acemoglu and Restrepo \(2021\)](#).

Note: Each marker corresponds to one of 500 demographic groups, defined by gender, age, education, race, and native/immigrant status. Marker sizes indicate the share of hours worked by each group and different colors indicate education levels.

Recent work by [Acemoglu and Restrepo \(2021\)](#) makes progress on this empirical challenge by taking a fresh approach to measuring wage impacts. Rather than studying wage changes in the occupations that workers do *at present*, they instead study the exposure of different demographic groups to displacement of routine tasks according to the industries and occupations in which these groups worked in 1980, before polarization got underway. The simple idea is that if workers of given education, age, gender, and race/ethnicity tended to work in routine-task-intensive jobs in 1980 (e.g., production occupations, clerical occupations), and the industries that employ them apply computers to automate those tasks, then the onset and evolution of widespread computerization over the ensuing decades would be expected to place downward pressure on their earnings. Evidence for this mechanism is seen in Figure 5, which reports a striking downward sloping relationship between exposure to routine task replacement in 1980 and changes in wages by demographic group between 1980 and 2016 (panel A). Equally striking is that this downward-sloping relationship is *not* present in the three prior decades, as shown in Panel B. This adds to the case that the negative relationship in Panel A reflects the adverse effect of routine-task displacement on the earnings of workers who, in earlier decades, tended to specialize in routine-task-intensive jobs.

Notice, however, that this evidence does not imply that *most* workers are harmed by computerization. For example, in panel A of Figure 5, only a subset of workers—those most exposed to task displacement—appear to have lost ground (in real earnings terms) between 1980 and 2016. This subset is almost entirely made up of workers with high school or lower education,

consistent with the evidence in Figure 2 that real wages of non-college workers have stagnated or fallen over the last four decades.¹² For the majority of workers, however, real earnings growth was positive in these decades, reflecting in part the productivity gains emanating from computerization (though many other factors are at play).

The task model thus underscores that technological change, like most economic transformations, creates both winners and losers. Akin to the education-race model, the task model also implies that computerization has contributed to rising inequality. Unlike the education-race framework, however, the task framework further implies that a substantial component of this effect stems from the adverse impacts of technological change on the earnings of less-educated workers rather than (exclusively) the positive effect of factor-augmentation on the earnings of high-skill workers. How large is that contribution? Acemoglu and Restrepo (2021) estimate that 50% to 70% of the increase in earnings inequality between education, sex, race, and age groups during 1980 through 2016—and the entirety of the fall in real wages of men without high school—are due to the adverse effects of automation on worker groups that were initially more specialized in routine task-intensive work.

3 New Work and Task Reinstatement

An important limitation of the task framework in its basic form is that it conceptualizes the set of tasks as static—meaning that none are added or subtracted, it’s only their allocation between workers and machines that shifts as technology and education evolve. This is a convenient fiction, but it has significant downsides. First, casual empiricism suggests that work is continually evolving, with demands for new skills and expertise that were previously unimagined (e.g., drone pilots, artificial intelligence programmers, vegan chefs, and executive coaches). Second, if the set of tasks were truly static, then it seems likely that advancing automation would inexorably crowd humans into an ever-diminishing subset of tasks, perhaps finally making human labor altogether obsolete, as envisioned by Susskind (2020). While one should not categorically exclude the possibility that this could occur, it does not accurately reflect the last century of technological change, during which the world of work has grown more complex, varied, and intellectually interesting (Autor, 2015).

An ingenious 2011 paper by Jeffrey Lin brings concrete evidence to these informal observations (Lin, 2011). Using historical Census documents from 1965 through 2000, Lin shows that the Census Bureau regularly captures novel job titles based on the occupational descriptions

¹² In a similar vein, Autor (2019) shows that the polarization of occupational structure primarily reflects the movement of non-college workers out of middle-skill occupations and into traditionally low-paid services. College-educated workers remain highly concentrated in professional, technical, and managerial occupations.

that survey respondents supply on their Census forms. While many of these novel write-ins are, of course, simply idiosyncratic descriptions or misspellings, the Census Bureau filters out the chaff to identify bona fide new job titles, reported by a significant number of Census respondents. Lin's work makes two contributions: first, it provides representative evidence on the appearance of 'new work'; second, it offers a methodology for systematically capturing new work hiding in plain sight in the Census Bureau's existing data infrastructure.

What precisely is new work? Table 1, drawn from Autor et al. (2021b), list examples of new titles added to the Census Bureau's internal occupational classification manual in each decade between 1940 and 2018.¹³ The left-hand column reveals, as intuition would suggest, that many new titles—such as textile chemists (added in 1960) or controllers of remotely piloted vehicles (added in 1980)—involve operating, installing, maintaining, integrating, or selling new technologies. While technology-related new titles are commonplace, just as prevalent are new titles that do *not* relate to a technological innovation but instead reflect changing tastes, incomes, and demographics (right-hand column). For example, beauticians (added in 1950), hypnotherapists (added in 1980), and sommeliers (added in 2000) provide specialized services. Surely, many new 'gig' titles will soon enter this list of titles: on-demand personal driver (Uber and Lyft); warehouse pick worker (Amazon); and on-demand shopper (Instacart), among others.

How does new work relate to the task-polarization framework elaborated above? Building on Lin's observations, Acemoglu and Restrepo (2018b) fuse the notion of new work (more precisely, new tasks) into the canonical task model. In their extended framework, automation displaces workers from existing job tasks as before; but now, new task creation potentially 'reinstates' demand for workers by generating new tasks that require human expertise.¹⁴ Thus, akin to the education-race model, the competing forces of *task automation* and *task displacement* determine the net effect of technological change on labor demand: if automation outpaces reinstatement, labor demand falls; and conversely, if reinstatement outpaces automation, labor demand rises.¹⁵

¹³ New titles introduced in a given decade, say 1940, correspond to those captured by the Census Bureau in the preceding decade, i.e., between 1931 and 1940.

¹⁴ While in the long run, these tasks may also be automated, it appears plausible that many novel activities are first accomplished and perfected by workers before they are subsequently routinized and automated.

¹⁵ This explanation oversimplifies for brevity. The net effect of automation and reinstatement depends not only on the relative speed of these forces but also their impact on aggregate labor demand through the productivity growth channel. Automation can raise labor demand even while displacing worker tasks if the resulting productivity boost raises demand sufficiently to offset employment losses due to task displacement. Logically, the impacts of automation and reinstatement may differ by skill group, as explored in Acemoglu and Restrepo (2021) and Autor et al. (2021b).

YEAR	EXAMPLE OF TITLES ADDED	
1940	Automatic welding machine operator	Gambling dealer
1950	Airplane designer	Beautician
1960	Textile chemist	Pageants director
1970	Engineer computer application	Mental-health counselor
1980	Controller, remotely piloted vehicle	Hypnotherapist
1990	Certified medical technician	Conference planner
2000	Artificial intelligence specialist	Chat room host/monitor
2010	Wind turbine technician	Sommelier
2018	Pediatric vascular surgeon	Drama therapist

Table 1: Examples of new occupational titles added to the U.S. Census Bureau’s Classified Index of Occupations between 1940 and 2018 Source:

Autor et al. (2021b).

Of course, knowing that old work is being automated and new work is being created does not tell us which effect dominates in net, which occupations or skill groups are most positively or negatively affected, and what underlying forces guide this process. Evidence is now emerging on these questions, though much more is needed. Employing an indirect measure of task change based on changes in labor’s share of income by industry, [Acemoglu and Restrepo \(2019\)](#) conduct a macroeconomic analysis of task displacement and task reinstatement for two long time intervals, 1950–1987 and 1987–2017. Their analysis suggests that these two forces—automation and task reinstatement—were roughly in balance in the first time interval of 1950–1987, but that automation subsequently outpaced task reinstatement in the second time interval of 1987–2017, which is consistent with labor’s falling share of national income occurring simultaneously ([Autor et al., 2020a](#); [Karabarbounis and Neiman, 2014](#)).¹⁶

To analyze representative evidence over a substantial time horizon, [Autor et al. \(2021b\)](#) build on the approach pioneered by [Lin \(2011\)](#) to analyze data on new work creation in eight decades of U.S. data from 1940 through 2018. These data suggest that new work is quantitatively

¹⁶ [Acemoglu and Restrepo \(2018b\)](#) offer a general theory of new work creation based on changes in the relative price of capital and labor, where declines in the price of labor spur labor-using innovations (and hence task reinstatement) and, conversely, declines in the price of capital spur capital-using automation innovations (and hence task displacement).

important. Autor et al. (2021b) estimate that more than 60% of employment in 2018 was found in job titles that did not exist in 1940, as shown in Figure 6.¹⁷ The introduction of new work, however, is *not* uniform across skill groups. Between 1940 and 1980, most new work that employed non-college workers was found in construction, transportation, production, clerical, and sales jobs—which are squarely middle-skill occupations. In the subsequent four decades (1980–2018), however, the locus of new work creation for non-college workers shifted away from these middle-tier occupations and towards traditionally lower-paid personal services. Conversely, new work creation employing college-educated workers became increasingly concentrated in professional, technical, and managerial occupations. In combination, these patterns indicate that new work creation has polarized, mirroring (and in part driving) the aggregate polarization of employment seen in Figure 3.

What explains the shifting locus of new work creation across occupations and skill groups during these decades? Autor et al. (2021b) document three critical forces. One is the introduction of automation innovations, which they find erode employment in occupations that are most exposed to them. But not all technological innovations are directed at automation. Using U.S. utility patent data, Autor et al. (2021b) develop a method to distinguish among innovations that *automate* the tasks that workers supply versus those that *augment* the outputs or services that their work generates. For example, the introduction of photocopying would constitute an automation innovation since it replaces the labor inputs of workers who previously duplicated documents using more cumbersome means (e.g., carbon paper). Conversely, the introduction of an electronic workbook for performing calculations (i.e., a spreadsheet) would constitute an augmentation innovation since it enhances the services provided by financial analysts, allowing them to conduct faster and deeper analyses.¹⁸ In contrast to the role of automation technologies, Autor et al. (2021b) document that *augmentation* innovations spur employment growth in the occupations most exposed to them. Given that many occupations are simultaneously exposed to both augmentation and automation innovations, this finding is particularly striking.

¹⁷ Using data from Lin (2011), Acemoglu and Restrepo (2018b) estimate a similar fraction of employment in new work for the shorter time interval of 1980 to 2015.

¹⁸ This example also highlights that many innovations contain elements of both automation and augmentation. The spreadsheet was surely an augmentation innovation for analysts, but it might also have been an automation innovation for routine bookkeepers.

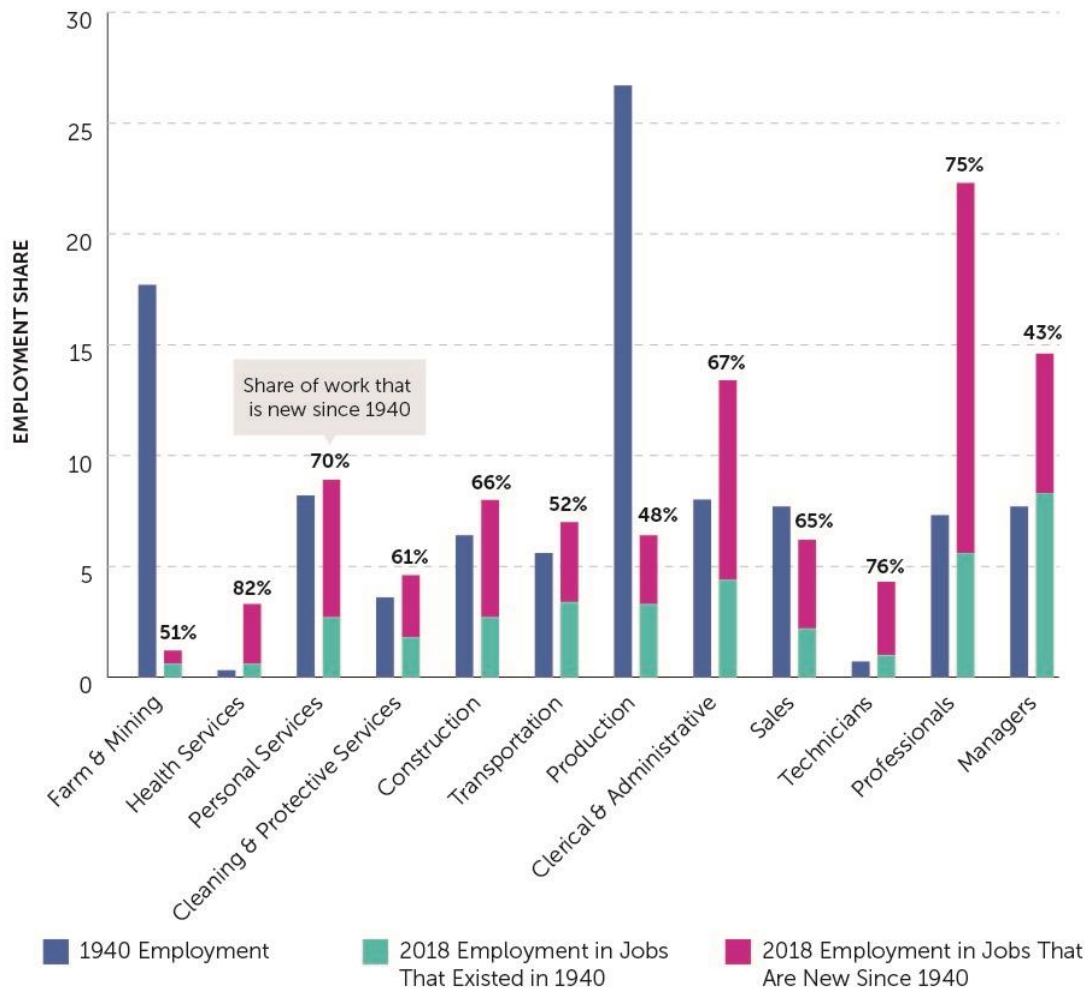


Figure 6: More than 60% of jobs done in the United States in 2018 had not yet been “invented” in 1940

Source: Autor et al. (2022).

Note: This figure compares the distribution of employment in 1940 and 2018 across all major occupations, distinguishing between employment in new job titles added between 1940 and 2018 versus job titles that were present in 1940.

Alongside these two faces of innovation—automation and augmentation—Autor et al. (2021b) analyze a third factor affecting new work creation: demand and supply forces that directly shape when and where new work emerges. When occupations are exposed to adverse demand shocks, for example, the contraction of manufacturing employment in the United States in the face of the China trade shock of the 1990s and 2000s (Autor et al., 2021a), not only does employment contract but the pace of introduction of new occupational titles slows. Conversely, when demand for an occupation expands, as for example has occurred in many personal care and healthcare occupations in the face of population aging, employment rises *and* the pace of new work introduction accelerates.

While much remains to be understood about the potential of new work creation to temper the task-eroding consequences of automation, it is clear that new work plays a critical role in shaping and tempering the long-run consequences of technological change for labor demand.

4 The Present Era of Artificial Intelligence Uncertainty

The task framework outlined above is well-suited to understanding the economic consequences of the last four decades of advancing digital computing. But how well does it fit the current era of Artificial Intelligence (AI)? Does AI fundamentally change the relationship between technological change, labor demand, and inequality—and if so, how do we characterize these changes analytically? The task framework provides a natural starting point both for considering what AI may do, and for understanding how AI differs from the technologies that preceded it.

The task framework encompasses two conceptual pieces. One is the notion of ‘tasks’ as units of work that can be accomplished by workers, machines, or potentially by service providers in other countries (see [Grossman and Rossi-Hansberg \(2008\)](#)). The second is a specific characterization of what tasks computers can accomplish—in particular, routine tasks in the terminology of [Autor et al. \(2003\)](#). What makes a task routine is that it follows an *explicit*, fully specified set of rules and procedures. Tasks fitting this description can in many cases be codified in computer software and executed by machines. Conversely, tasks that rely on what [Polanyi and Sen \(1966\)](#) called ‘tacit’ knowledge (e.g., riding a bicycle, telling a clever joke) have historically been challenging to program because the explicit steps for accomplishing these tasks are often not formally known.

Artificial intelligence overturns the second piece of the task framework—specifically, the stipulation that computers can accomplish only explicitly understood (i.e., ‘routine’) tasks. AI tools surmount this longstanding constraint because they can be used to infer tacit relationships that are not fully specified by underlying software. For example, it is extraordinarily challenging to explicitly define what makes a chair a chair: must it have legs, and if so, how many; must it have a back; what range of heights is acceptable; must it be comfortable; and what makes a chair comfortable, anyway? Writing the rules for this problem is maddening. If written too narrowly, they will exclude stools and rocking chairs. If written too broadly, they will include tables and countertops. In a well-known paper, ([Grabner et al., 2011](#)) argue that the fundamental problem is that what makes a chair a chair is its suitability for sitting upon. But what makes something “suitable” for sitting is as elusive as the original problem. Given this morass, this chair classification task would be categorized as ‘non-routine’ for purposes of conventional computing—a human task rather than a machine task.

Fast forward to the present and AI can now ‘solve’ this classification problem. It does not solve it by following explicit rules, however. Instead, it learns the solution inductively by training on examples. Given a suitable database of tagged images and sufficient processing power, AI can infer what image attributes are statistically associated with the label “chair” and can then use that information to classify untagged images of chairs with a high degree of accuracy (Brynjolfsson and Mitchell, 2017; Brynjolfsson et al., 2018). What rules does AI use for this classification? In general, we do not know because the rules remain tacit. Nowhere in the learning process does AI formally codify or reveal the underlying features (i.e., rules) that constitute “chair-ness”. Rather, the classification decision emerges from layers of learned statistical associations with no human-interpretable window into that decision-making process.¹⁹ And herein lies an irony: Polanyi’s paradox survives the paradigm shift in computing, but with a twist. In the pre-AI era, programmers struggled to imbue computers with the tacit knowledge needed for accomplishing non-routine tasks; in the present AI-era, computers can readily acquire this tacit knowledge, but they cannot (in almost all cases) communicate that knowledge explicitly to people. That is, computers now know more than they can tell us.²⁰

Returning to the task model, how does the relaxation of the tacit knowledge constraint affect our predictions of what machines and people will do in the future? One potential answer is that the task model is now irrelevant given that machines are increasingly capable of accomplishing non-routine tasks.²¹ An alternative answer is that the task model remains conceptually and empirically valuable because it provides an analytic tool for rigorously studying the interactions between human and machine capabilities in accomplishing work (Acemoglu and Restrepo, 2018a; Autor, 2013)—though it makes fewer crisp predictions about what tasks are likely to be automated in the years ahead. I see three questions as particularly relevant:

- Looking through the lens of the task framework, what work tasks will AI prove capable of accomplishing in the years (and decades) ahead? AI’s applicability is in my assessment sufficiently vast that I find it harder to say what AI *cannot* do than what it *can and will* do.²² It is commonly argued, for example, that because AI is blissfully unaware of the rich context of many real-world problems, it cannot accomplish the high-stakes, multifaceted decision tasks that humans regularly undertake in their work. This argument would be convincing if

¹⁹ Schematically, AI learns by adjusting connection weights among layers of (virtual) nodes on an information network. The representation of the decision-making process in this network has essentially no relationship to the formal structure of the problem as a human would understand it.

²⁰ The field of explainable AI seeks to make the tacit knowledge acquired by AI explicit. See, for example, https://en.wikipedia.org/wiki/Explainable_artificial_intelligence.

²¹ See Bresnahan (2021) for a strident argument that the task model is irrelevant in the AI era, and perhaps was irrelevant in all prior eras.

²² See Marcus and Davis (2019) for a counterargument.

humans were highly effective and reliable at making such decisions. But the evidence strongly suggests that they are not (Kahneman et al., 2021).

- Second, what new demands for human skills and capabilities will emerge as AI displaces a growing set of traditional human work tasks? As per the discussion above, I am certain that such new work tasks will emerge, and that many forms of human capability and expertise will become newly valuable. Because technological advances have always generated new demands for human specialization, as have rising societal wealth and ongoing changes in norms, tastes, and institutions, I do not foresee a moment when labor scarcity (and hence, labor income) is eliminated. Simultaneously, many currently valuable human capabilities will eventually be rendered obsolete. This will be costly for many and disruptive for society in general. These disruptions are *also* characteristic of technological upheavals, but because of the rapidity with which AI is evolving, they may be particularly acute.
- Third, while the task framework offers a useful starting point for analyzing the impact of AI on labor markets and inequality, it is unlikely to be encompassing enough to reflect all relevant labor market impacts of AI. And it is certainly insufficient to capture many of the broader societal impacts. How do we get a fuller analytical grasp on the terrain ahead? Works by Agrawal et al. (2018), Bresnahan (2021), and Korinek and Stiglitz (2018) offer different lenses on these questions that bring different issues into focus. We are only at the start of the intellectual journey to understand AI's implications for work and inequality, so it would be premature to proclaim that we've already found the most promising route to that destination.

A small but rapidly growing literature that includes Babina et al. (2020), Brynjolfsson and Mitchell (2017), Brynjolfsson et al. (2018), Felten et al. (2018, 2019), and Webb (2020) applies a task approach to analyze the labor market impacts of AI adoption.²³ These recent papers make an important break with prior task-based studies. Earlier incarnations of this literature often focused specifically on whether routine tasks were substituted by computers and non-routine tasks were complemented. Thus, they applied both the general task framework *and* the specific characterization of the intrinsic capabilities and limitations of procedural computing supplied by Autor et al. (2003). In contrast, recent works studying the labor market impacts of AI apply the task framework generally but do not for the most part characterize analytically precisely what AI can do—which makes sense because such a characterization remains elusive.²⁴ Instead, these

²³ Though not specifically focused on AI, recent papers by Atalay et al. (2020) and Deming and Noray (2020) present novel, closely related analyses.

²⁴ Qualifying these generalizations, Brynjolfsson and Mitchell (2017) and Brynjolfsson et al. (2018) offer a rubric for assessing the suitability of job tasks to machine learning (their SML index), while Agrawal et al. (2018) offer a formal

papers develop or apply expert or crowd-sourced assessments of the tasks for which AI is currently suitable to determine which tasks, occupations, firms, and industries are most AI-exposed.

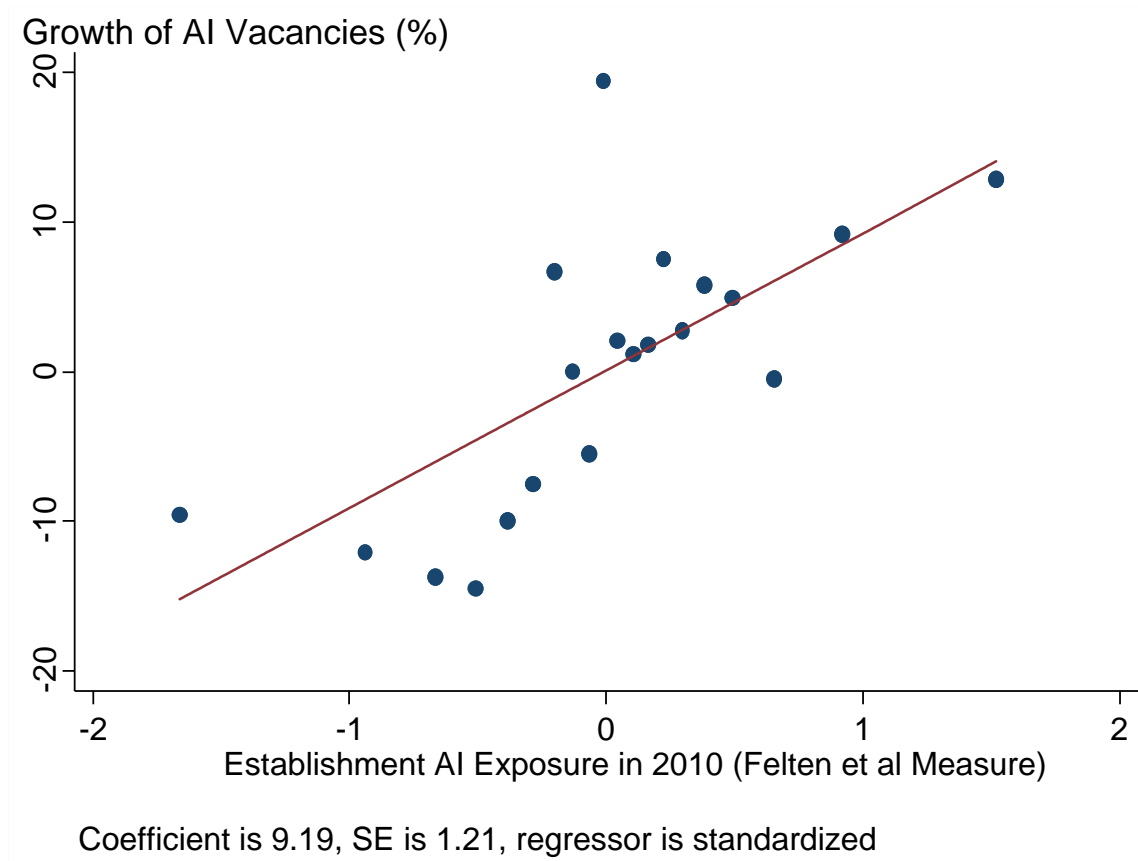


Figure 7: U.S. establishments whose task structures in 2010 were more suitable for AI subsequently posted relatively more AI vacancies between 2010 and 2018

Source: [Acemoglu et al. \(2022a\)](#).

Note: This figure shows the relationship between establishment-level AI exposure in 2010, computed using the [Felten et al. \(2018\)](#) method, and establishments’ subsequent increase in posting of job positions requiring AI skills between 2010 and 2018. The solid line corresponds to a regression with 2010 establishment vacancies as the weight. Each bin represents about 50,000 establishments. Point estimates and standard errors are reported at the bottom of the figure. This analysis excludes vacancies in AI-producing sectors of the economy, specifically Information (NAICS code 51) and Business Services (NAICS code 54).

[Acemoglu et al. \(2022a\)](#) offer one recent example. This study uses establishment-level job vacancy postings from the online job-posting aggregator Burning Glass Technologies to assess the impacts of recent AI adoption on the demand for workers who perform non-AI jobs. For this analysis, the study defines “AI jobs” as those that advertise specific expertise requirements in

characterization of what tasks AI accomplishes. Specifically, they argue that AI is essentially a prediction machine—a tool that forecasts the immediate (or long-term) future based on past inputs.

contemporary AI tools.²⁵ Conversely, non-AI jobs are the (vast) remainder that do not demand AI-specific skills but nevertheless may be affected by it. This could include a whole range of jobs, from financial analysts, to pharmacists, to pilots, to warehouse managers. Drawing on AI-suitability indexes developed by Brynjolfsson et al. (2018), Felten et al. (2018), and Webb (2020), the paper first predicts which establishments are likely to adopt AI as a function of the suitability of their job task structures (visible in job postings) in the pre-AI era. Consistent with this prediction, the paper documents that establishments whose occupational structures in 2010 made them suitable for AI *did* in fact differentially increase posting of vacancies for workers with AI skills as AI took off between 2010 and 2018, as is shown in Figure 7.

With these predictions in hand, Acemoglu et al. (2022a) explore whether AI adoption (spurred by AI-suitability) is affecting hiring in non-AI jobs. The answer is a qualified yes. They find that as AI-exposed establishments adopted AI between 2010 and 2018 (particularly after 2014), they differentially changed their mix of job skill requirements in *non-AI* positions—suggesting that non-AI job tasks were affected—and modestly reduced hiring in non-AI positions simultaneously. This evidence confirms that AI’s imprint can already be seen at the firms and establishments whose preexisting task structures make them more suitable for using AI. Yet, Acemoglu et al. (2022a) find that AI is not so far having detectable labor market impacts at the aggregate occupation or industry level, though such affects appear likely in the future. In net, these conclusions are evocative but not dramatic; they hint at potential *aggregate* effects of AI but do not so far confirm them.

“Aggregate effects” is a pregnant phrase: what might those effects be? Here, I speculate:

1. One such aggregate effect is that further improvements in AI’s capabilities may accelerate the process of task automation relative to task augmentation. Broadly, this will mean that labor’s share of national income will decline further, beyond what has already occurred over the last two decades as documented in Autor et al. 2020a, and, concomitantly, the share of national income paid to owners of capital (i.e., machines, robots, algorithms, etc.) will grow. Ironically, this process of aggregate labor displacement can occur without any reduction in wage inequality among workers—or with wage inequality rising even further. Specifically, all workers could get a smaller slice of the aggregate economic pie while the proportional difference among those slices remained just as pronounced. This fall in labor’s share of national income does *not*, however, necessarily imply that employment will fall. So as long as

²⁵ Examples include Machine Learning, Computer Vision, Machine Vision, Deep Learning, Virtual Agents, Image Recognition, Natural Language Processing, Speech Recognition, Pattern Recognition, Object Recognition, and Neural Networks, among many others.

people need to work for a living, falling wages do not preclude stable or rising employment.²⁶ Additionally, a fall in the labor share does not necessarily mean a decline in wages; the same capabilities that make AI labor-displacing could in theory generate sufficient productivity growth that average wages would rise even as labor's share of national income falls. In this case, the size of the pie grows faster than labor's share of that pie shrinks. Nevertheless, a fall in labor's share of national income is problematic, the simple reason being that the ownership of capital is far more concentrated than the ownership of labor (i.e., absent coercion, each person owns only their own labor). Thus, a substantial fall in labor's share of national income implies a dramatic rise in *income* inequality—that is, wage plus non-wage income—even absent a change in *wage* inequality.

2. A second scenario—which could in theory co-occur with the one above—is that, spurred by advancing AI, the twin forces of task-automation and task-augmentation reshape the set of tasks (and associated worker skills) that are complemented and substituted by technologies. While the last four decades of conventional computing capabilities have fomented occupational polarization and rising wage inequality, this need not be true going forward, or at least not to the same degree. It is a near-certainty that AI will increasingly be deployed to accomplish mid- and high-level decision-making tasks that have historically been performed by managers and professionals. This is already occurring in finance and investing, inventory management, credit issuance, fraud detection, and even some fields of design. An expanding set of these expert and semi-expert tasks will almost surely become technologically equivalent to the “routine tasks” of earlier years: equally well accomplished by machines, and with greater rapidity and at lower cost. Accordingly, it is possible that even those with moderately high levels of educational attainment—those, for example, with a BA but no post-graduate degree—will find that their primary work tasks are increasingly substituted by AI.

That some of their tasks are substituted does mean that these workers' skills are necessarily devalued. It is in part by displacing a subset of human tasks that, in many cases, automation makes the remaining set of worker tasks more valuable. (Imagine the value of a statistician stripped of her computer or a construction crew denied use of power tools.) Whether workers' skills are complemented or substituted by new technologies depends in part on their ability to adapt to changing task demands. Economists have long understood that education makes people better at adapting to, and capitalizing on, novel circumstances (Schultz, 1975).

²⁶ For those skeptical of this point, note that employment rates are generally higher in poor than rich countries and that hours worked per capita tend to fall among both men and women as countries become wealthier (Bick et al. 2018).

But this resilience is not guaranteed. At the turn of the twentieth century, high school graduates were an elite education group who commanded substantial premia as bookkeepers and clerks. In essence, they were the leading “information technology” of big business in that era (Goldin and Katz, 2008). Today, however, there is little difference in the wages paid to high school graduates and those without a high school degree. Thus, the high school credential has lost much of its market value, except as a waypoint on the road to higher education (Goldin and Katz, 2008; Card, 2009).

Still, there is an upper limit to this substitution process at present. While there is no consensus on the topic, many experts do not expect artificial general intelligence (AGI) to emerge for some decades, if at all (Fijelland, 2020). Assuming this expectation is correct (which I believe it is), humans will continue to have comparative advantage in creativity, judgment, hypothesis formation, contextual thinking, causal analysis, communication, emotional intelligence, and many more arenas, the importance of which we likely do not fully appreciate and the difficulty of which we surely vastly underestimate. I feel confident that the *most* skilled workers will likely continue to be complemented by advances in computing and AI—such as workers who invent, design, research, lead, entertain, and educate. But this observation is not limited to those with elite educations. People effortlessly do extraordinary things on an ongoing basis, such as applying common sense to tease apart otherwise intractable problems; drawing generalizable inferences from *small* data; and using abductive reasoning to form plausible interpretations of a sparse set of observations. Such seemingly quotidian tasks are currently beyond the frontier of the most advanced AI, and yet children accomplish them effortlessly (Davis and Marcus, 2019). Recalling Polanyi’s observations that we know more than we can tell, I would add that we do not fully appreciate how much of this tacit knowledge we possess.

3. I similarly do not expect AI to rapidly reach deep into the ranks of low-paid service occupations, those comprising the left-hand side of the occupational polarization plot shown in Figure 3. There are three reasons why not. First, most service occupations demand dexterous, fluid, adaptive interactions with people and the environment, whether in care jobs, services, entertainment, etc. Automating these activities will require substantial advances in low-cost robots that can navigate in the highly variable human environment rather than in the predictable engineered environment of a factory floor. These advances will take place much more slowly than advances in AI, which depend primarily on more of the same ingredients—more data, greater computing power. Second, while machines will surely slowly gain many of these human-like capabilities, their cost may remain high relative to the low cost of labor performing those same activities. This cost comparison makes the

economics of automating many service tasks less attractive.²⁷ Finally, many low-paid service tasks—such as caregiving, coaching, and advising and selling—are unattractive targets for automation not only because the technical challenge is steep but because personal attention from another human being is intrinsically part of the service.

4. Finally, while it is easy to imagine which tasks and what jobs will succumb to automation, it is far harder to forecast what and where new work will emerge. Millions of workers are currently employed in order-fulfillment and ride-hailing jobs that were in effect created by e-commerce and mobile telephony. Similarly, the U.S. Bureau of Labor Statistics maintains information on “green jobs” associated with the transformation of the power sector.²⁸ Many of these occupations are relatively new or rapidly growing, such as solar plumbers, solar site assessors, and specialized plumbers, pipefitters, and steamfitters. Artificial intelligence itself has created a host of new skill demands and occupational specialties, as documented in Acemoglu et al. (2022a). As discussed in Autor et al. (2021b), new innovations almost always generate new work as people deploy, master, maintain, refine, and improve new technologies, tools, and services. Nor does new work generation depend exclusively on innovation. Autor et al. (2021b) further demonstrate that changes in demographics, tastes, and income levels also drive the generation of new work.

What these observations imply is that the work of the future is not an empty set—not even remotely. In Autor et al. (2022), we write that “No compelling historical or contemporary evidence suggests that technological advances are driving us toward a jobless future. On the contrary, we anticipate that in the next two decades, industrialized countries will have more job openings than workers to fill them, and that robotics and automation will play an increasingly crucial role in closing these gaps. Nevertheless, the impact of robotics and automation on workers will not be benign. These technologies, in concert with economic incentives, policy choices, and institutional forces, will alter the set of jobs available and the skills they demand.” It is that adaptation that creates both challenge and opportunity. The problem that industrialized countries face in the immediate decades ahead is not a shortfall in the quantity of jobs. It is rather that many of the

²⁷ An exception to this dictum is that service tasks that are done at large scale may attract automation. For example, Amazon, which employs hundreds of thousands of warehouse workers, has invested heavily in robotics to automate part of the product fulfillment process. Similarly, White Castle restaurants have deployed fryolator-operating robots in some of their many stores. This same economic logic may drive robotics in table-waiting, hotel room cleaning, shelf-stocking, and checkout operations, even though all are low-paid tasks that require substantial human flexibility. The attractiveness of automation will increase if the cost of human labor in these tasks rises—a healthy economic process. The scenario to be concerned about is one where automation makes formerly scarce labor broadly abundant (and hence cheap); not one in which scarce labor makes automation more attractive at the margin.

²⁸ <https://www.bls.gov/green/overview.htm>

jobs may be of low quality, use only generic human capacities, and provide little opportunity for skills acquisition, specialization, and rising lifecycle productivity. This is not a new problem, however. It has been unfolding over four decades. And in general, the U.S. has adapted to it poorly.

5 Conclusions

I began by asking what the role of technology—digital or otherwise—is in determining wages and shaping wage inequality. I presented four answers corresponding to four strands of thinking on this topic: the education race, the task-polarization model, the automation-reinstatement race, and the era of AI uncertainty. The nuance of economic understanding has improved substantially across these epochs. Yet, traditional economic optimism about the beneficent effects of technology for productivity and welfare has eroded as understanding has advanced. Fundamentally, technological change expands the frontier of human possibilities, but one should expect it to create many winners and many losers, and to pose vast societal challenges and opportunities along the way.

What are the policy implications of these observations? The question is so broad that almost any answer is bound to appear vague and inadequate. One can reliably predict that technological innovations will foster new ways of accomplishing existing work, new business models, and entirely new industries, and these in turn will generate new jobs and spur some productivity gains. But absent complementary institutional investments, technological innovation alone will not generate broadly shared gains. [Autor et al. \(2022\)](#) sketch a long-form policy vision of what form these investments may take, focusing on three domains: education and training; labor market institutions; and innovation policy itself. I provide a capsule summary of those conclusions here, stopping short of policy recommendations.

As a starting point, educating and training the workforce to meet the demands of the moment can improve access to good jobs for workers who would otherwise face barriers and boost the quality of existing jobs by creating opportunities for career ladders. While boosting the supply of skills is never controversial, pure supply-side policies are inadequate to meet the gravity of the labor market challenge. In 1979, 60% of U.S. males at the median of the wage distribution possessed high school or lower education, whereas 21% had some college, and 20% held a bachelor's degree or above. By 2019, 31% of males at the median of the earnings distribution had some college, 36% had attained at least a four-year college degree — a 75% increase — and only

one-third had high school or less education.²⁹ Despite this dramatic skill upgrading, the real earnings of the median male rose by only 10% in these forty years.

More broadly, over more than four decades, the link between rising productivity and commensurate improvements in job opportunities and earnings has decoupled for the majority of U.S. workers. The poor quality of jobs available to workers lacking four-year college degrees or specialized credentials provides one of the starkest examples of this failure. Low-wage U.S. workers earn substantially less than low-wage workers in almost all other wealthy industrialized countries.³⁰

The divergence between the upward path of productivity growth and the near plateauing of median wage growth among U.S. workers was not an inevitable consequence of technology, globalization, or market forces (Autor et al. 2022). Rather, a set of U.S.-specific institutional and policy choices failed to blunt—and in some cases magnified—the consequences of technological and globalization pressures on the U.S. labor market. To contend effectively with these challenges would require institutional and policy reforms that realign labor market opportunities with the rising productivity and societal wealth that the U.S. has reaped from decades of innovation and investments in human and physical capital.

Thus, a second locus for investment would be to revitalize the governmental, nongovernmental, and private sector institutions that translate—or fail to translate—rising productivity into shared prosperity. Some such steps might include updating and more vigorously enforcing labor standards, recalibrating federal minimum wage policies, extending the scope and flexibility of the unemployment insurance system, and transforming the U.S. employer-based health insurance provision into a system with portable benefits. Autor et al. (2022) further call for reevaluating the U.S.’ uncritical embrace of pure shareholder capitalism. While shareholder capitalism can plausibly be credited with some of the productive dynamism of the U.S. economy, it has also arguably helped fuel the drive to curtail wages and benefits for low-wage workers and, more broadly (Acemoglu et al. 2022b).

A third potential domain for policy is to directly shape innovation itself to speed productivity growth and complement the skills of the labor force. It is well known that the U.S. has a strong national innovation system, fueled by federal R&D investments, to develop fundamental science and new technologies, that has led to scientific leadership and new industries (Gross and Sampat, 2020). Less recognized, however, is the crucial link between those new industries as complements to the inevitable loss of jobs that results from productivity-enhancing technologies. New industries grew out of a flourishing innovation ecosystem that created new companies and new

²⁹ These statistics from Donovan and Bradley (2020, Table 5) refer to male workers at the 45th-55th percentiles of the male wage distribution in the corresponding year.

³⁰ This paragraph and the next excerpt from Autor et al. (2022), p101-102.

applications, alongside older industries that increased mechanization and automation as they matured. Arguably, the U.S. has let those important R&D investments wither (Bloom et al. 2019). Through increased and targeted R&D investments supported by a reinvigorated federal R&D program, as well as a tax policy that keeps workers and social challenges at the forefront, the country's innovation system could be put to work for a broader number of people and regions than it has in recent decades (Gruber and Johnson, 2019).³¹

Although I would prefer to end this essay with optimistic assurances, I will instead end with one uncertainty, one certainty, and one admonition. The *uncertainty* is that we have less clarity about our technological future than we did two decades ago. AI has extended the frontier of technological possibility towards boundaries that are barely visible at present. The tasks that machines will be able to accomplish, the rate at which new innovations may emerge, and the speed with which socially impactful technological innovations may diffuse is unknown. But the range of possibilities has surely gotten broader, and our certainty about the boundaries has accordingly diminished.

Conversely, the *certainty* is that these technological advances will expand the set of desirable possibilities that are within the reach of humanity. Artificial Intelligence has the potential to help humanity tackle some of its most pressing challenges: climate change, disease, poverty, malnutrition, and inadequate education. But whether societies will successfully realize this potential, or instead squander it or, worse, disastrously misuse it, is highly uncertain and, I would argue, fundamentally indeterminate.

The *admonition* is this: given the potential applicability of AI to a vast set of purposes, we collectively (meaning individuals, organizations, and governments) should not simply be asking what AI will accomplish but what we want it to accomplish. How can we use AI most productively to complement workers, raise productivity, and more broadly, tackle humanity's most pressing challenges? Simultaneously, how can we blunt or reshape the commercial incentives to use AI for socially counterproductive objectives such as displacing workers, preying upon people's cognitive and emotional frailties, or consolidating the power of governments or corporations to exercise social control (Acemoglu and Restrepo, 2020)? As we ponder our uncertain AI future, our goal should not merely be to predict that future but to create it.

References

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022a). "AI and Jobs: Evidence from Online Vacancies." *Journal of Labor Economics*, 40(S1):S293-S340.

³¹ This paragraph excerpts from Autor et al. (2022), p121.

- Acemoglu, D., He, Alex X., and Le Maire, D. (2022b). "Eclipse of Rent-Sharing: The Effects of Managers' Business Education on Wages and the Labor Share in the US and Denmark." NBER Working Paper No. 29874, March.
- Acemoglu, D. and Autor, D. H. (2011). "Skills, Tasks and Technologies: Implications for Employment and Earnings." *Handbook of Labor Economics*, 4(11):1043–1171.
- Acemoglu, D. and Restrepo, P. (2018a). "Modeling Automation." *American Economic Review: Papers and Proceedings*, 108:48–53.
- Acemoglu, D. and Restrepo, P. (2018b). "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review*, 108(6):1488–1542.
- Acemoglu, D. and Restrepo, P. (2019). "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives*, 33(2):3–30.
- Acemoglu, D. and Restrepo, P. (2020). "The Wrong Kind of AI? Artificial Intelligence and the Future of Labor Demand." *Cambridge Journal of Regions, Economy and Society*, 13(1):25–35.
- Acemoglu, D. and Restrepo, P. (2021). "Tasks, Automation, and the Rise in U.S. Wage Inequality." NBER Working Paper No. 28920, June.
- Agrawal, A., Gans, J., and Goldfarb, A. (2018). *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Press.
- Atalay, E., Phongthientham, P., Sotelo, S., and Tannenbaum, D. (2020). "The Evolution of Work in the United States." *American Economic Journal: Applied Economics*, 12(2):1–34.
- Autor, D. (2013). "The 'Task Approach' to Labor Markets: An Overview." *Journal for Labour Market Research*, 46(3):185–199.
- Autor, D., Dorn, D., and Hanson, G. H. (2021a). "On the Persistence of the China Shock." *Brookings Papers on Economic Activity*, forthcoming.
- Autor, D., Dorn, D., Katz, L. F., Patterson, C., and Van Reenen, J. (2020a). "The Fall of the Labor Share and the Rise of Superstar Firms." *The Quarterly Journal of Economics*, 135(2):645–709.
- Autor, D., Goldin, C., and Katz, L. F. (2020b). "Extending the Race between Education and Technology." In *AEA Papers and Proceedings*, volume 110, pages 347–51.
- Autor, D., Mindell, D., and Reynolds, E. (2022). *The Work of the Future: Building Better Jobs in an Age of Intelligent Machines*. Cambridge: MIT Press.
- Autor, D., Salomons, A., and Seegmiller, B. (2021b). "New Frontiers: The Origin and Content of New Work, 1940 – 2018." MIT Working Paper, July.
- Autor, D. H. (2014). "Skills, Education, and the Rise of Earnings Inequality among the 'Other 99 Percent'." *Science*, 344(6186):843–851.

- Autor, D. H. (2015). "Why are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives*, 29(3):3–30.
- Autor, D. H. (2019). "Work of the Past, Work of the Future." *AEA Papers and Proceedings*, volume 109, pages 1–32.
- Autor, D. H. and Dorn, D. (2013). "The Growth of Low Skill Service Jobs and the Polarization of the U.S. Labor Market." *American Economic Review*, 103(5):1553–1597.
- Autor, D. H., Katz, L. F., and Kearney, M. S. (2006). "The Polarization of the U.S. Labor Market." *American Economic Review*, 96(2):189–194.
- Autor, D. H., Katz, L. F., and Kearney, M. S. (2008). "Trends in U.S. Wage Inequality: Revising the Revisionists." *The Review of Economics and Statistics*, 90(2):300–323.
- Autor, D. H., Katz, L. F., and Krueger, A. B. (1998). "Computing Inequality: Have Computers Changed the Labor Market?" *The Quarterly Journal of Economics*, 113(4):1169–1213.
- Autor, D. H., Levy, F., and Murnane, R. J. (2003). "The Skill Content of Recent Technological Change: An Empirical Exploration." *The Quarterly Journal of Economics*, 118(4):1279–1333.
- Babina, T., Fedyk, A., He, A., and Hodson, J. (2020). "Artificial Intelligence, Firm Growth, and Industry Concentration." Available at SSRN: <https://ssrn.com/abstract=3651052>.
- Bick, A., Fuchs-Schündeln, N., & Lagakos, D. (2018). "How Do Hours Worked Vary with Income? Cross-Country Evidence and Implications." *American Economic Review*, 108(1), 170-99.
- Bloom, N., Van Reenen, J. and Williams, H. (2019). "A Toolkit of Policies to Promote Innovation." *Journal of Economic Perspectives* 33, no. 3: 163–184.
- Böhm, M. J. (2020). "The Price of Polarization: Estimating Task Prices under Routine-Biased Technical Change." *Quantitative Economics*, 11(2):761–799.
- Böhm, M. J., von Gaudecker, H.-M., and Schran, F. (2019). "Occupation Growth, Skill Prices, and Wage Inequality." CESifo Working Paper No. 7877.
- Bresnahan, T. (2021). "Artificial Intelligence Technologies and Aggregate Growth Prospects" in Diamond, John, and George Zodrow (eds.), *Prospects for Economic Growth in the United States*, Cambridge University Press.
- Brynjolfsson, E. and Mitchell, T. (2017). "What Can Machine Learning Do? Workforce Implications." *Science*, 358(6370):1530–1534.
- Brynjolfsson, E., Mitchell, T., and Rock, D. (2018). "What Can Machines Learn, and What Does it Mean for Occupations and the Economy?" *AEA Papers and Proceedings*, 108:43–47.
- Card, D. and Lemieux, T. (2001a). "Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis." *The Quarterly Journal of Economics*, 116(2):705–746.

- Card, D. and Lemieux, T. (2001b). "Going to College to Avoid the Draft: The Unintended Legacy of the Vietnam War." *American Economic Review*, 91(2):97–102.
- Card, D. (2009). "Immigration and Inequality." *American Economic Review: Papers & Proceedings*, 99(2), 1-21.
- Deming, D. J. (2017). "The Growing Importance of Social Skills in the Labor Market." *The Quarterly Journal of Economics*, 132(4):1593–1640.
- Deming, D. J. and Noray, K. (2020). "Earnings Dynamics, Changing Job Skills, and STEM Careers." *The Quarterly Journal of Economics*, 135(4):1965–2005.
- Donovan, Sarah A, and David H Bradley. (2020). "Real Wage Trends, 1979 to 2019." Congressional Research Service, Report R45090, December.
- Dorn, D. (2009). "Essays on Inequality, Spatial Interaction, and the Demand for Skills" (Doctoral dissertation, University of St. Gallen).
- Feigenbaum, J. and Gross, D. P. (2020). "Answering the Call of Automation: How the Labor Market Adjusted to the Mechanization of Telephone Operation." NBER Working Paper No. 28061. November.
- Felten, E., Manav, R., and Seamans, R. (2018). "A Method to Link Advances in Artificial Intelligence to Occupational Abilities." *American Economic Association Papers and Proceedings*, 108(54):54–57.
- Felten, E., Raj, M., and Seamans, R. C. (2019). "The Effect of Artificial Intelligence on Human Labor: An Ability-Based Approach." In *Academy of Management Proceedings*, volume 2019, page 15784. Academy of Management Briarcliff Manor, NY 10510.
- Fjelland, R. (2020). "Why General Artificial Intelligence Will Not Be Realized." *Nature: Humanities and Social Sciences Communications*, 7(1):1-9.
- Goldin, C. and Katz, L. F. (1998). "The Origins of Technology-Skill Complementarity." *The Quarterly Journal of Economics*, 113(3):693–732.
- Goldin, C. and Katz, L. F. (2008). *The Race Between Education and Technology*. Harvard University Press.
- Goldin, C., Katz, L. F., et al. (2007). "Long-Run Changes in the Wage Structure: Narrowing, Widening, Polarizing." *Brookings Papers on Economic Activity*, 38(2):135–168.
- Goldin, C. and Margo, R. A. (1992). "The Great Compression: The Wage Structure in the United States at Mid-Century." *The Quarterly Journal of Economics*, 107(1):1–34.
- Goos, M. and Manning, A. (2007). "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain." *Review of Economics and Statistics*, 89(1):118–133.
- Goos, M., Manning, A., and Salomons, A. (2009). "Job Polarization in Europe." *American Economic Review: Papers & Proceedings*, 99(2):58–63.

- Goos, M., Manning, A., and Salomons, A. (2014). "Explaining Job Polarization: Routine Biased Technological Change and Offshoring." *American Economic Review*, 104(8):2509– 2526.
- Grabner, H., Gall, J., and Van Gool, L. (2011). "What Makes a Chair a Chair?" In *CVPR 2011*, pages 1529–1536. IEEE.
- Gregory, T., Salomons, A., and Zierahn, U. (2022). "Racing with or Against the Machine? Evidence on the Role of Trade in Europe." *Journal of the European Economic Association*, 20(2):869-906.
- Gross, Daniel P., and Bhaven N. Sampat. (2020). "Inventing the Endless Frontier: The Effects of the World War II Research Effort on Post-War Innovation." NBER Working Paper No. 27375, June.
- Grossman, G. M. and Rossi-Hansberg, E. (2008). "Trading Tasks: A Simple Theory of Offshoring." *American Economic Review*, 98(5):1978–97.
- Gruber, J. and Johnson, S. (2019). *Jump-Starting America: How Breakthrough Science Can Revive Economic Growth and the American Dream*. Public Affairs.
- Harrigan, J., Reshef, A., and Toubal, F. (2021). "The March of the Techies: Technology, Trade, and Job Polarization in France, 1994-2007." *Research Policy*, 50(7).
- Kahneman, D., Sibony, O., and Sunstein, C. R. (2021). *Noise: A Flaw in Human Judgment*. Little, Brown.
- Karabarbounis, L. and Neiman, B. (2014). "The Global Decline of the Labor Share." *The Quarterly Journal of Economics*, 129(1):61–103.
- Katz, L. F. and Autor, D. H. (1999). "Changes in the Wage Structure and Earnings Inequality." *Handbook of Labor Economics*, 3:1463–1555.
- Katz, L. F. and Murphy, K. M. (1992). "Changes in Relative Wages, 1963–1987: Supply and Demand Factors." *The Quarterly Journal of Economics*, 107(1):35–78.
- Korinek, A. and Stiglitz, J. E. (2018). "Artificial Intelligence and its Implications for Income Distribution and Unemployment." In *The Economics of Artificial Intelligence: An Agenda*, pages 349–390. University of Chicago Press.
- Krusell, P., Ohanian, L. E., Ríos-Rull, J.-V., and Violante, G. L. (2000) "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis." *Econometrica*, 68(5): 1029-1053.
- Levy, F. and Murnane, R. J. (2004). *The New Division of Labor: How Computers are Creating the Next Job Market*. Princeton University Press; Russell Sage Foundation.
- Lin, J. (2011). "Technological Adaptation, Cities, and New Work." *The Review of Economics and Statistics*, 93(2):554–574.

- Marcus, G. and Davis, E. (2019). *Rebooting AI: Building Artificial Intelligence We Can Trust*. New York: Vintage Press.
- Maloney, William F., and Carlos Molina. (2016). "Are Automation and Trade Polarizing Developing Country Labor Markets, Too?" World Bank Policy Research Working Paper 7922.
- Michaels, G., Natraj, A., and Van Reenen, J. (2014). "Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years." *Review of Economics and Statistics*, 96(1):60–77.
- Nordhaus, W. D. (2007). "Two Centuries of Productivity Growth in Computing." *The Journal of Economic History*, 67(1):128–159.
- Polanyi, M. (1966). *The Tacit Dimension*. Doubleday.
- Schultz, T. W. (1975). "The Value of the Ability to Deal with Disequilibria." *Journal of Economic Literature*, 13(3):827-846.
- Smith, A. and Anderson, M. (2017). *Automation in Everyday Life*. Pew Research Center.
- Susskind, D. (2020). *A World Without Work: Technology, Automation, and How We Should Respond*. Metropolitan Books, New York, N.Y.
- Tinbergen, J. (1974). "Substitution of Graduate Labor by Other." *Kyklos*, 27(2):217–226.
- Webb, M. (2020). "The Impact of Artificial Intelligence on the Labor Market." Working paper.
- Wike, R. and Stokes, B. (2018). "In Advanced and Emerging Economies Alike, Worries about Job Automation." Pew Research Center, Global Attitudes & Trends.