

Why Are There Still So Many Jobs? The History and Future of Workplace Automation[†]

David H. Autor

There have been periodic warnings in the last two centuries that automation and new technology were going to wipe out large numbers of middle class jobs. The best-known early example is the Luddite movement of the early 19th century, in which a group of English textile artisans protested the automation of textile production by seeking to destroy some of the machines. A lesser-known but more recent example is the concern over “The Automation Jobless,” as they were called in the title of a *TIME* magazine story of February 24, 1961:

The number of jobs lost to more efficient machines is only part of the problem. What worries many job experts more is that automation may prevent the economy from creating enough new jobs. . . . Throughout industry, the trend has been to bigger production with a smaller work force. . . . Many of the losses in factory jobs have been countered by an increase in the service industries or in office jobs. But automation is beginning to move in and eliminate office jobs too. . . . In the past, new industries hired far more people than those they put out of business. But this is not true of many of today’s new industries. . . . Today’s new industries have comparatively few jobs for the unskilled or semiskilled, just the class of workers whose jobs are being eliminated by automation.

Concerns over automation and joblessness during the 1950s and early 1960s were strong enough that in 1964, President Lyndon B. Johnson empaneled a

■ *David H. Autor is Professor of Economics, Massachusetts Institute of Technology, Cambridge, Massachusetts. From 2009 to 2014, he was Editor of the Journal of Economic Perspectives.*

[†]To access the Data Appendix and disclosure statement, visit
<http://dx.doi.org/10.1257/jep.29.3.3>

“Blue-Ribbon National Commission on Technology, Automation, and Economic Progress” to confront the productivity problem of that period—specifically, the problem that productivity was rising so fast it might outstrip demand for labor. The commission ultimately concluded that automation did not threaten employment: “Thus technological change (along with other forms of economic change) is an important determinant of the precise places, industries, and people affected by unemployment. But the general level of demand for goods and services is by far the most important factor determining how many are affected, how long they stay unemployed, and how hard it is for new entrants to the labor market to find jobs. The basic fact is that technology eliminates jobs, not work” (Bowen 1966, p. 9). However, the Commission took the reality of technological disruption as severe enough that it recommended, as one newspaper (*The Herald Post* 1966) reported, “a guaranteed minimum income for each family; using the government as the employer of last resort for the hard core jobless; two years of free education in either community or vocational colleges; a fully administered federal employment service, and individual Federal Reserve Bank sponsorship in area economic development free from the Fed’s national headquarters.”

Such concerns have recently regained prominence. In their widely discussed book *The Second Machine Age*, MIT scholars Erik Brynjolfsson and Andrew McAfee (2014, p. 11) offer an unsettling picture of the likely effects of automation on employment:

Rapid and accelerating digitization is likely to bring economic rather than environmental disruption, stemming from the fact that as computers get more powerful, companies have less need for some kinds of workers. Technological progress is going to leave behind some people, perhaps even a lot of people, as it races ahead. As we’ll demonstrate, there’s never been a better time to be a worker with special skills or the right education, because these people can use technology to create and capture value. However, there’s never been a worse time to be a worker with only ‘ordinary’ skills and abilities to offer, because computers, robots, and other digital technologies are acquiring these skills and abilities at an extraordinary rate.

Clearly, the past two centuries of automation and technological progress have not made human labor obsolete: the employment-to-population ratio rose during the 20th century even as women moved from home to market; and although the unemployment rate fluctuates cyclically, there is no apparent long-run increase. But those concerned about automation and employment are quick to point out that past interactions between automation and employment cannot settle arguments about how these elements might interact in the future: in particular, the emergence of greatly improved computing power, artificial intelligence, and robotics raises the possibility of replacing labor on a scale not previously observed. There is no fundamental economic law that guarantees every adult will be able to earn a living solely on the basis of sound mind and good character. Whatever the future holds, the present clearly offers a resurgence of automation anxiety (Akst 2013).

In this essay, I begin by identifying the reasons that automation has not wiped out a majority of jobs over the decades and centuries. Automation does indeed substitute for labor—as it is typically intended to do. However, automation also complements labor, raises output in ways that lead to higher demand for labor, and interacts with adjustments in labor supply. Indeed, a key observation of the paper is that journalists and even expert commentators tend to overstate the extent of machine substitution for human labor and ignore the strong complementarities between automation and labor that increase productivity, raise earnings, and augment demand for labor.

Changes in technology do alter the types of jobs available and what those jobs pay. In the last few decades, one noticeable change has been “polarization” of the labor market, in which wage gains went disproportionately to those at the top and at the bottom of the income and skill distribution, not to those in the middle. I will offer some evidence on this phenomenon. However, I will also argue that this polarization is unlikely to continue very far into the foreseeable future.

The final section of this paper reflects on how recent and future advances in artificial intelligence and robotics should shape our thinking about the likely trajectory of occupational change and employment growth. I argue that the interplay between machine and human comparative advantage allows computers to substitute for workers in performing routine, codifiable tasks while amplifying the comparative advantage of workers in supplying problem-solving skills, adaptability, and creativity. The frontier of automation is rapidly advancing, and the challenges to substituting machines for workers in tasks requiring flexibility, judgment, and common sense remain immense. In many cases, machines both substitute for and complement human labor. Focusing only on what is lost misses a central economic mechanism by which automation affects the demand for labor: raising the value of the tasks that workers uniquely supply.

How Automation and Employment Interact

In 1900, 41 percent of the US workforce was employed in agriculture; by 2000, that share had fallen to 2 percent (Autor 2014), mostly due to a wide range of technologies including automated machinery. The mass-produced automobile drastically reduced demand for many equestrian occupations, including blacksmiths and stable hands. Successive waves of earth-moving equipment and powered tools displaced manual labor from construction. In more recent years, when a computer processes a company’s payroll, alphabetizes a list of names, or tabulates the age distribution of residents in each Census enumeration district, it is replacing a task that a human would have done in a previous era. Broadly speaking, many—perhaps most—workplace technologies are designed to save labor. Whether the technology is tractors, assembly lines, or spreadsheets, the first-order goal is to substitute mechanical power for human musculature, machine-consistency for human handiwork, and digital calculation for slow and error-prone “wetware.”

Given that these technologies demonstrably succeed in their labor saving objective and, moreover, that we invent many more labor-saving technologies all the time, should we not be somewhat surprised that technological change hasn't *already* wiped out employment for the vast majority of workers? Why doesn't automation *necessarily* reduce aggregate employment, even as it demonstrably reduces labor requirements per unit of output produced?

These questions underline an economic reality that is as fundamental as it is overlooked: tasks that cannot be substituted by automation are generally complemented by it. Most work processes draw upon a multifaceted set of inputs: labor and capital; brains and brawn; creativity and rote repetition; technical mastery and intuitive judgment; perspiration and inspiration; adherence to rules and judicious application of discretion. Typically, these inputs *each* play essential roles; that is, improvements in one do not obviate the need for the other. If so, productivity improvements in one set of tasks almost necessarily increase the economic value of the remaining tasks.

An iconic representation of this idea is found in the O-ring production function studied by Kremer (1993).¹ In the O-ring model, failure of any one step in the chain of production leads the entire production process to fail. Conversely, improvements in the reliability of any given link increase the value of improvements in all of the others. Intuitively, if $n - 1$ links in the chain are reasonably likely to fail, the fact that link n is somewhat unreliable is of little consequence. If the other $n - 1$ links are made reliable, then the value of making link n more reliable as well rises. Analogously, when automation or computerization makes some steps in a work process more reliable, cheaper, or faster, this increases the value of the remaining human links in the production chain.

As a contemporary example, consider the surprising complementarities between information technology and employment in banking, specifically the experience with automated teller machines (ATMs) and bank tellers documented by Bessen (2015). ATMs were introduced in the 1970s, and their numbers in the US economy quadrupled from approximately 100,000 to 400,000 between 1995 and 2010. One might naturally assume that these machines had all but eliminated bank tellers in that interval. But US bank teller employment actually rose modestly from 500,000 to approximately 550,000 over the 30-year period from 1980 to 2010 (although given the growth in the labor force in this time interval, these numbers do imply that bank tellers declined as a share of overall US employment). With the growth of ATMs, what are all of these tellers doing? Bessen observes that two forces worked in opposite directions. First, by reducing the cost of operating a bank branch, ATMs indirectly increased the demand for tellers: the number of tellers *per branch* fell by more than a third between 1988 and 2004, but the number of urban bank branches (also encouraged by a wave of

¹ The name of the O-ring production function refers to the 1986 accident of Space Shuttle Challenger, which exploded and crashed back to earth less than two minutes after takeoff, killing its seven crew members. The proximate cause of the Challenger crash was an inexpensive and seemingly inconsequential rubber O-ring seal in one of its booster rockets that failed after hardening and cracking during the icy Florida weather on the night before takeoff.

bank deregulation allowing more branches) rose by more than 40 percent. Second, as the routine cash-handling tasks of bank tellers receded, information technology also enabled a broader range of bank personnel to become involved in “relationship banking.” Increasingly, banks recognized the value of tellers enabled by information technology, not primarily as checkout clerks, but as salespersons, forging relationships with customers and introducing them to additional bank services like credit cards, loans, and investment products.

This example should not be taken as paradigmatic; technological change is not necessarily employment-increasing or Pareto-improving. Three main factors can mitigate or augment its impacts. First, workers are more likely to benefit directly from automation if they supply tasks that are complemented by automation, but not if they primarily (or exclusively) supply tasks that are substituted. A construction worker who is expert with a shovel but cannot drive an excavator will generally experience falling wages as automation advances. Similarly, a bank teller who can tally currency but cannot provide “relationship banking” is unlikely to fare well at a modern bank.

Second, the elasticity of labor supply can mitigate wage gains. If the complementary tasks that construction workers or relationship bankers supply are abundantly available elsewhere in the economy, then it is plausible that a flood of new workers will temper any wage gains that would emanate from complementarities between automation and human labor input. While these kinds of supply effects will probably not offset productivity-driven wage gains fully, one can find extreme examples: Hsieh and Moretti (2003) document that new entry into the real estate broker occupation in response to rising house prices fully offsets average wage gains that would otherwise have occurred.

Third, the output elasticity of demand combined with income elasticity of demand can either dampen or amplify the gains from automation. In the case of agricultural products over the long run, spectacular productivity improvements have been accompanied by declines in the share of household income spent on food. In other cases, such as the health care sector, improvements in technology have led to ever-larger shares of income being spent on health. Even if the elasticity of final demand for a given sector is below unity—meaning that the sector shrinks as productivity rises—this does not imply that aggregate demand falls as technology advances; clearly, the surplus income can be spent elsewhere. As passenger cars displaced equestrian travel and the myriad occupations that supported it in the 1920s, the roadside motel and fast food industries rose up to serve the “motoring public” (Jackson 1993). Rising income may also spur demand for activities that have nothing to do with the technological vanguard. Production of restaurant meals, cleaning services, haircare, and personal fitness is neither strongly complemented nor substituted by current technologies; these sectors are “technologically lagging” in Baumol’s (1967) phrase. But demand for these goods appears strongly income-elastic, so that rising productivity in technologically leading sectors may boost employment nevertheless in these activities. Ultimately, this outcome requires that the elasticity of substitution between leading and lagging sectors is less than or equal to unity (Autor and Dorn 2013).

Over the very long run, gains in productivity have not led to a shortfall of demand for goods and services: instead, household consumption has largely kept pace with household incomes. We know this because the share of the population engaged in paid employment has generally risen over (at least) the past century despite vast improvements in material standards of living. An average US worker in 2015 wishing to live at the income level of an average worker in 1915 could roughly achieve this goal by working about 17 weeks per year.² Most citizens would not consider this tradeoff between hours and income desirable, however, suggesting that consumption demands have risen along with productivity. Of course, citizens in high-income countries work fewer annual hours, take more vacations, and retire earlier (relative to death) than a century ago—implying that they choose to spend part of their rising incomes on increased leisure. This is clearly good news on many fronts, but does it also imply that consumption demands are approaching satiation? I think not. In high-income countries, consumption and leisure appear to be complements; citizens spend much of their leisure time consuming—shopping, traveling, dining, and, less pleasantly, obtaining medical care.³


What about the Marxian concern that automation will immiserate workers by obviating the demand for labor? In simple economic models, this outcome cannot really occur because capital is owned by the economic agents who are presumably also the workers; but, alternatively, the returns could accrue to a narrow subset of agents. Sachs and Kotlikoff (2012) and Sachs, Benzell, and LaGarda (2015) explore multigenerational economic environments in which a burst of robotic productivity can enrich one generation of capital owners at the expense of future generations. These later generations suffer because the fruits of the productivity surge are consumed by the old, while the young face diminished demand for their labor and, in some cases, also experience credit constraints that inhibit their human capital investments. In these models, the fundamental threat is not technology per se but misgovernance; an appropriate capital tax will render the technological advance broadly welfare-improving, as these papers stress. Thus, a key takeaway is that rapid automation may create distributional challenges that invite a broad policy response, a point to which I will return.

² Douglas (1930; reproduced in US Bureau of the Census 1949) reports average annual earnings across all sectors in 1915 at \$633. Inflating this to 2015 dollars using the US Bureau of Labor Statistics historical Consumer Price Index calculator yields a current dollar equivalent of \$14,711. The BLS employment report from April 2015 reports mean weekly private nonfarm earnings of \$858. Thus, it would take 17 weeks of work at the average US weekly wage to earn a full-time annual 1915 income.

³ This outcome is a modern version of the “coal paradox” posed by William Stanley Jevons in his 1865 book *The Coal Question*. Jevons argued that as we became more efficient in mining coal, we would use more of it, not less. Modern environmental economists term this idea the “rebound effect.” In this discussion, the broad parallel is that greater efficiency of production of all goods and services means that we consume more of them, not the same or less.

Figure 1

Average Change per Decade in US Occupational Employment Shares for Two Periods: 1940–1980 and 1980–2010



Source: Based on Katz and Margo (2014), table 1.6, panel A, which is based upon the 1920 through 2000 Census of population IPUMS and 2010 American Community Survey.

Notes: Observed long changes in US occupational employment shares over 1940–1980 and 1980–2010 are scaled by the number of intervening decades to yield average change per decade. Occupations are classified into occupational groups based on 1950 occupation codes using the consistent coding of occupations in all years into 1950 codes (the OCC1950 variable) in the IPUMS. Additional details are found in Katz and Margo (2014, p. 46).

Polarization in the US Labor Market

Even if automation does not reduce the quantity of jobs, it may greatly affect the qualities of jobs available. For the three decades or so from the end of World War II and up through the late 1970s, the US experienced rapid automation and technological change—inspiring, for example, the *TIME* magazine story in 1961 and Lyndon Johnson’s 1964 National Commission mentioned earlier. While it’s difficult to paint an accurate picture of occupational change over a large time interval, Figure 1, which draws from Katz and Margo (2014), provides a high-level overview by depicting the average change per decade in employment for seven broad occupational categories, ranked from lowest to highest paid, for two periods: 1940–1980 and 1980–2010. In the first four decades after World War II, the thrust of occupational change skewed strongly away from physically demanding, dangerous, and

menial work and towards skilled blue- and white-collar work. Agricultural employment declined by almost 4 percentage points per decade. Professional, technical, and managerial employment—the highest skill categories—grew by 3 percentage points per decade (2.5 for the professionals and technicians plus 0.5 for the managers). And among the vast middle group of workers between agriculture (at the bottom) and professional, technical, and managerial (the three groups at the top), service and skilled blue-collar occupations were stable, clerical/sales occupations rose, and operative and laborer occupations fell sharply.

Thus, physically demanding, repetitive, dangerous, and cognitively monotonous work was receding, ushered out by extraordinary productivity gains in agriculture. Rising consumer affluence spurred demand for manufactured goods and leisure complements. Growth of technologically intensive corporations, health care services, and higher education created employment for credentialed professionals and a cadre of supporting clerical, administrative, and sales workers. Though automation was clearly reducing labor demand across a large swath of occupations, it is easy to see why overall job prospects appeared broadly favorable during this period.

But after the late 1970s, these favorable winds slowed and in some cases reversed. While jobs at the top of the skill ladder—professional, technical, and managerial occupations—grew even more rapidly between 1980 and 2010 than in the four decades prior, positive occupational shifts outside of these categories mostly halted. Skilled blue-collar occupations shrank rapidly and clerical and sales occupations—the vulnerable “production jobs” of the information age—sharply reversed course. While physically demanding operative and laborer jobs continued to atrophy, low-paid personal services began absorbing an increasing share of noncollege labor. By this time, the vast movement away from agricultural work had already played out.

Many forces distinguish the labor markets of these two epochs of 1940–1980 and 1980–2010: a partial list would include changes in the relative supply of college and noncollege labor, rising trade penetration, offshoring, and globalization of production chains, declines in labor union penetration, the changing “bite” of the minimum wage, and certain shifts in tax policy. Of course, many of these factors combine and interact as well such that attributing changes to a single cause would be foolish. However, my focus here is on the effects of technological change, and especially information technology, on employment and occupations (and later wages). To understand the role that information technology has played (and may play), it is useful to start from first principles: What do computers do? And how does their widespread adoption change what workers do?

Fundamentally, computers follow procedures meticulously laid out by programmers. The typical pattern has been that for a computer to accomplish a task, a programmer must first fully understand the sequence of steps required to perform that task, and then must write a program that, in effect, causes the machine to simulate these steps precisely. (The field of machine learning, discussed below, provides an interesting exception to this process.) When a computer processes a company’s

payroll, alphabetizes a list of names, or tabulates the age distribution of residents in each Census enumeration district, it is “simulating” a work process that would, in a previous era, have been done by humans using nearly identical procedures. The principle of computer simulation of workplace tasks has not fundamentally changed since the dawn of the computer era—but its cost has. An ingenious 2007 paper by William Nordhaus estimates that the cost of performing a standardized set of computations has fallen by at least 1.7 trillion-fold since the manual computing era, with most of that decline occurring since 1980. Thus, firms have strong economic incentives to substitute ever-cheaper computing power for relatively expensive human labor. What are the effects?

One first-order effect is, of course, substitution. As the price of computing power has fallen, computers and their robot cousins have increasingly displaced workers in accomplishing explicit, codifiable tasks. In Autor, Levy, and Murnane (2003), my coauthors and I label these activities as “routine tasks,” not because they are mundane, but because they can be fully codified and hence automated (see Levy and Murnane 2004 for many examples). Routine tasks are characteristic of many middle-skilled cognitive and manual activities: for example, the mathematical calculations involved in simple bookkeeping; the retrieving, sorting, and storing of structured information typical of clerical work; and the precise executing of a repetitive physical operation in an unchanging environment as in repetitive production tasks. Because core tasks of these occupations follow precise, well-understood procedures, they are increasingly codified in computer software and performed by machines. This force has led to a substantial decline in employment in clerical, administrative support, and to a lesser degree, in production and operative employment.

But the scope for this kind of substitution is bounded because there are many tasks that people understand tacitly and accomplish effortlessly but for which neither computer programmers nor anyone else can enunciate the explicit “rules” or procedures. I have referred to this constraint as Polanyi’s paradox, named after the economist, philosopher, and chemist who observed in 1966, “We know more than we can tell” (Polanyi 1966; Autor 2015). When we break an egg over the edge of a mixing bowl, identify a distinct species of birds based on a fleeting glimpse, write a persuasive paragraph, or develop a hypothesis to explain a poorly understood phenomenon, we are engaging in tasks that we only tacitly understand how to perform. Following Polanyi’s observation, the tasks that have proved most vexing to automate are those demanding flexibility, judgment, and common sense—skills that we understand only tacitly.⁴

Polanyi’s paradox also suggests *why* high-level reasoning is straightforward to computerize and certain sensorimotor skills are not. High-level reasoning uses a set

⁴ Computer scientists often refer to this phenomenon as Moravec’s paradox, after Moravec (1988) who wrote, “[I]t is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility.”

of formal logical tools that were developed specifically to address formal problems: for example, counting, mathematics, logical deduction, and encoding quantitative relationships. In contrast, sensorimotor skills, physical flexibility, common sense, judgment, intuition, creativity, and spoken language are capabilities that the human species evolved, rather than developed. Formalizing these skills requires reverse-engineering a set of activities that we normally accomplish using only tacit understanding. Hoffman and Furcht (2014) discuss the challenge that Polanyi's paradox poses for scientific innovation more broadly.

If computers largely substitute for routine tasks, how do we characterize the nonroutine tasks for which they do not substitute? In Autor, Levy, and Murnane (2003), we distinguish two broad sets of tasks that have proven stubbornly challenging to computerize. One category includes tasks that require problem-solving capabilities, intuition, creativity, and persuasion. These tasks, which we term "abstract," are characteristic of professional, technical, and managerial occupations. They employ workers with high levels of education and analytical capability, and they place a premium on inductive reasoning, communications ability, and expert mastery. The second broad category includes tasks requiring situational adaptability, visual and language recognition, and in-person interactions—which we call "manual" tasks. Manual tasks are characteristic of food preparation and serving jobs, cleaning and janitorial work, grounds cleaning and maintenance, in-person health assistance by home health aides, and numerous jobs in security and protective services. These jobs tend to employ workers who are physically adept and, in some cases, able to communicate fluently in spoken language. While these activities are not highly skilled by the standards of the US labor market, they present daunting challenges for automation. Equally noteworthy, many outputs of these manual task jobs (haircuts, fresh meals, housecleaning) must be produced and performed largely on-site or in person (at least for now), and hence these tasks are not subject to outsourcing. The potential supply of workers who can perform these jobs is very large.


Because jobs that are intensive in either abstract or manual tasks are generally found at opposite ends of the occupational skill spectrum—in professional, managerial, and technical occupations on the one hand, and in service and laborer occupations on the other—this reasoning implies that computerization of "routine" job tasks may lead to the simultaneous growth of high-education, high-wage jobs at one end and low-education, low-wage jobs at the other end, both at the expense of middle-wage, middle education jobs—a phenomenon that Goos and Manning (2003) called "job polarization." A large body of US and international evidence confirms the presence of employment polarization at the level of industries, localities, and national labor markets (Autor, Katz, and Kearney 2006, 2008; Goos and Manning 2007; Autor and Dorn 2013; Michaels, Natraj, and Van Reenen 2014; Goos, Manning, and Salomons 2014; Graetz and Michaels 2015; Autor, Dorn, and Hanson 2015).⁵

⁵ Mishel, Shierholz, and Schmitt (2013) offer an extended, and for the most part extremely careful, critique of the literature on technological change, employment, and wage inequality. Their paper argues

Figure 2

Change in Employment by Major Occupational Category, 1979–2012

(the y-axis plots 100 times log changes in employment, which is nearly equivalent to percentage points for small changes)



Sources: Author using data from the 1980, 1990, and 2000 Census IPUMS files, American Community Survey combined file 2006–2008, and American Community Survey 2012. The sample includes the working-age (16–64) civilian noninstitutionalized population. Employment is measured as full-time equivalent workers. *Notes:* Figure 2 plots percentage point changes in employment (more precisely, the figure plots 100 times log changes in employment, which is close to equivalent to percentage points for small changes) by decade for the years 1979–2012 for ten major occupational groups encompassing all of US nonagricultural employment. Agricultural occupations comprise no more than 2.2 percent of employment in this time interval, so this omission has a negligible effect.

Figure 2 illustrates this pattern for the United States by plotting percentage point changes in employment by decade for the years 1979–2012 for ten major occupational groups encompassing all of US nonagricultural employment. (More

that the growth of low-wage service employment does not commence in the United States until the 2000s, a finding that is at odds with all other work using contemporary occupation codes of which I am aware (including the Bureau of Labor Statistic's own tabulations of Occupational Employment Statistics data for this time period provided in Alpert and Auyer 2003, table 1). At a methodological level, work in this area always requires adjustments and judgment calls in comparing occupational data across Census years, but the adjustments that Mishel et al. apply to the data generate occupational patterns that appear anomalous. Substantively, I believe the main issue is not whether *employment polarization* has occurred—on this, the evidence appears unambiguous—but the extent to which these occupational employment shifts are helpful for understanding *wage polarization* or wage inequality more broadly.

precisely, the figure plots 100 times log changes in employment, which are close to equivalent to percentage points for small changes. Agricultural occupations comprise no more than 2.2 percent of employment in this time interval, so this omission has a negligible effect.) These ten occupations can be divided into three groups. On the right-hand side of the figure are managerial, professional, and technical occupations, which are highly educated and highly paid. Moving leftward, the next four columns display employment growth in middle-skill occupations, comprising sales; office and administrative support; production, craft and repair; and operator, fabricator, and laborer. The leftmost three columns of Figure 2 depict employment trends in service occupations, defined by the Census Bureau as jobs that involve helping, caring for, or assisting others. The majority of workers in service occupations have no post-secondary education, and average hourly wages in service occupations are in most cases below the other seven occupational categories.

As Figure 2 illustrates, the rapid employment growth in both high- and low-education jobs has substantially reduced the share of employment accounted for by “middle-skill” jobs. In 1979, the four middle-skill occupations (sales; office and administrative workers; production workers; and operatives) accounted for 60 percent of employment. In 2007, this number was 49 percent, and in 2012, it was 46 percent. The employment share of service occupations was essentially flat between 1959 and 1979, and so their rapid growth since 1980 marks a sharp trend reversal (Autor and Dorn 2013).


The polarization of employment across occupations is not unique to the United States. Figure 3 plots changes in the share of employment between 1993 and 2010 within three broad sets of occupations—low-, middle-, and high-wage—covering all nonagricultural employment in 16 European Union economies. In all countries, middle-wage occupations declined as a share of employment while both high-wage and low-wage occupations increased their shares of employment over this 17-year period. While the US and EU data are not precisely comparable, the US economy would fall roughly in the middle of the pack of this set of countries in terms of its employment polarization. The comparability of these occupational shifts across a large set of developed countries makes it likely that a common set of forces contributes to these shared labor-market developments. Simultaneously, the substantial differences among countries underscores that no single factor or common cause explains the diversity of experiences across the United States and the European Union.

Does Employment Polarization Lead to Wage Polarization?

From the barbell shape of occupational employment growth depicted in Figures 2 and 3, one might surmise that occupational polarization would also catalyze wage polarization—that is, rising relative wages in both high-education, abstract task-intensive jobs and in low-education, manual task-intensive jobs. However, this

Figure 3

Change in Occupational Employment Shares in Low, Middle, and High-Wage Occupations in 16 EU Countries, 1993–2010



Source: Goos, Manning, and Salomons (2014, table 2).

Notes: High-paying occupations are corporate managers; physical, mathematical, and engineering professionals; life science and health professionals; other professionals; managers of small enterprises; physical, mathematical, and engineering associate professionals; other associate professionals; life science and health associate professionals. Middle-paying occupations are stationary plant and related operators; metal, machinery, and related trade work; drivers and mobile plant operators; office clerks; precision, handicraft, craft printing, and related trade workers; extraction and building trades workers; customer service clerks; machine operators and assemblers; and other craft and related trade workers. Low-paying occupations are laborers in mining, construction, manufacturing, and transport; personal and protective service workers; models, salespersons, and demonstrators; and sales and service elementary occupations.

reasoning does not take into account the role played by the three mitigating forces discussed above: complementarity, demand elasticity, and labor supply.

Let's first consider the effect of computerization on wages in abstract task-intensive occupations such as managerial, professional, and technical occupations. These occupations all draw upon large bodies of constantly evolving expertise: for example, medical knowledge, legal precedents, sales data, financial analysis, programming languages, and economic statistics. Information technology and computerization should strongly complement workers performing abstract task-intensive jobs. By dramatically lowering the cost and increasing the scope of information and analysis available to them, computerization enables workers performing abstract tasks to further specialize in their area of comparative advantage, with less time spent on acquiring and crunching information, and more time spent on interpreting and applying it. By the same token, information technology substitutes for many of the

support occupations that these professions employ, including medical secretaries, paralegals, and research assistants. Similarly, computerization and information technology appears to allow “delaying” of management structures (Caroli and Van Reenen 2001). Arguably, many of the middle managers displaced by delaying performed routine information-processing tasks.

If demand for the output of abstract task-intensive activities is inelastic, these productivity gains might work to lower expenditure on these outputs, which could mitigate wage gains. However, all outward evidence suggests that as technology has boosted the output of the professions, demand for their services has more than kept pace. Health care is an obvious example, but one can readily make similar arguments about finance, law, engineering, research, and design.

What about reactions from labor supply? If workers could quickly move into the highly educated professions, such a shift would mute earnings gains. But of course, many professions require both college and graduate degrees, so the production pipeline for new entrants is at least five to ten years in length. Indeed, young US adults, particularly US males, have responded with remarkable sluggishness to the rising educational premium over the last 30 years (Autor 2014). For example, in 1975, approximately 40 percent of hours worked by males with fewer than ten years of experience (a group that has made the more recent choices about college) were supplied by those with a college education. Forty years later in 2005, this share was almost unchanged. For women workers with less than ten years of experience, the share of total hours worked by those with a college education was 42 percent in 1982 but had risen to 53 percent by 2005. In the last decade, the share of hours worked by those with less than ten years of experience and a college degree has increased for both men and women: in 2012, it was 52 percent of hours for men in this group and 62 percent of the hours for women. Thus, while the stock of workers with college and graduate degrees has certainly grown, the supply response has not been nearly large enough to swamp the contemporaneous movements in labor demand.

Workers in abstract task-intensive occupations therefore benefit from information technology via a virtuous combination of strong complementarities between routine and abstract tasks, elastic demand for services provided by abstract task-intensive occupations, and inelastic labor supply to these occupations over the short and medium term. In combination, these forces mean that information technology should raise earnings in occupations that make intensive use of abstract tasks and among workers who intensively supply them.

These same synergies do not apply to jobs that are intensive in manual tasks, such as janitors and cleaners, vehicle drivers, security guards, flight attendants, food service workers, and home health aides. Most manual task-intensive occupations are only minimally reliant on information or data processing for their core tasks, and involve only limited opportunities for either direct complementarity or substitution.⁶

⁶ There are partial exceptions to this generalization: global positioning system satellites and scheduling software allows truckers and delivery services to minimize wasted mileage; calendar, contact, and billing

Aggregate evidence suggests that final demand for manual task-intensive work—services in particular—is relatively *price* inelastic (Baumol 1967; Autor and Dorn 2013). If so, productivity gains in manual task-intensive occupations that tend to reduce their price per unit of service provided will not necessarily raise expenditure on their outputs. On the other hand, demand for manual task-intensive work appears to be relatively *income* elastic (Clark 1951; Mazzorali and Ragusa 2013), so that rising aggregate incomes will tend to increase demand for these activities. New technology and productivity growth in other areas may therefore *indirectly* raise demand for manual task-intensive occupations by increasing societal income.

Labor supply to manual task-intensive occupations is intrinsically elastic, due to their generally low education and training requirements. This insight does not preclude the possibility that wages in manual tasks will rise, at least to some extent. As Baumol (1967) observed, even absent productivity growth in technologically lagging occupations, wages in these occupations *must* rise over time with societal income to compensate workers for *not* entering other sectors (again, assuming that demand for these activities is relatively inelastic). But it does suggest that wage increases in these jobs will be restrained to some extent by the labor supply response, including from workers displaced in other sectors of the economy.

Overall, manual task-intensive activities are at best weakly complemented by computerization, do not benefit from elastic final demand, and face elastic labor supply that tempers demand-induced wage increases. Thus, while information technology has strongly contributed to *employment* polarization measured in quantity of jobs, we would not generally expect these employment changes to culminate in a corresponding *wage* polarization except perhaps at certain times or in certain labor markets. Indeed, in Autor and Dorn (2013), we present evidence that wages for manual-task occupations rose during the 1990s when labor markets were extremely tight, but after 2000, the expansion of manual task-intensive service occupations accelerated while wages in these occupations fell.

For insight about the evolution of wage patterns, consider Figure 4. The horizontal axis of this figure is based on a ranking of all 318 detailed occupations from lowest to highest by their initial skill level, as measured by its 1979 mean hourly occupational wage. These categories are weighted by their initial size, and then grouped into 100 bins of equal size. The vertical axis of the figure then shows the percentage change in wages over each of four periods across the skill distribution—with the line smoothed for clarity. (Again, more precisely, the figure plots 100 times log changes in employment, which is nearly equivalent to percentage points for small changes.)


The right-hand two-thirds of Figure 4 look like the plots of employment polarization. From 1979 through 2007, wages rose consistently across the high-skill portion

software assists home health workers to manage data more effectively; and computerized ordering systems enable food service workers to rapidly tally customer tabs. In a few years time, many retailers may employ RFID “chip” technology that will scan purchases without needing a human checkout cashier at all.

Figure 4

Changes in Mean Wages by Occupational Skill Percentile among Full-Time, Full-Year (FTFY) Workers, 1979–2012

(the y-axis plots 100 times log changes in employment, which is nearly equivalent to percentage points for small changes)



Sources: Author, calculated using 1980, 1990, and 2000 Census IPUMS files; American Community Survey combined file 2006–2008, American Community Survey 2012.

Notes: The figure plots changes in mean log wages over each period, by 1979 occupational skill percentile rank using a locally weighted smoothing regression (bandwidth 0.8 with 100 observations), where skill percentiles are measured as the employment-weighted percentile rank of an occupation's mean log wage in the Census IPUMS 1980 5 percent extract. The sample includes the working-age (1–64) civilian non-institutionalized population with 48+ annual weeks worked and 35+ usual weekly hours. Weekly wages are calculated as annual earnings divided by weeks worked.

of the figure, which is disproportionately made up of the abstract task-intensive categories of professional, technical, and managerial occupations. By contrast, wage growth in the middle-skill, typically routine task-intensive occupations was less rapid and generally decelerated over time. For the low-education, manual task-intensive occupations heavily represented on the left-hand side of Figure 4, in the 1980s, wage growth was a little *more* rapid than in the middle-skill occupations—and in the 1990s, it was much more rapid. However, that changed in the 2000s: while Figure 2 showed that employment growth in these occupations exceeded that in all other categories between 1999 and 2007, Figure 4 shows wage growth was generally negative in the low-skill percentiles, lower than in all other categories (Mishel, Shierholz, and Schmitt 2013). During this time period, my strong hunch is that the explanation is that declining employment in middle-skill routine task-intensive

jobs led middle-skill workers—including new entrants, those displaced from routine task-intensive jobs, and those who lost jobs during recession—to enter manual task-intensive occupations instead (Smith 2013; Cortes, Jaimovich, Nekarda, and Siu 2014; Foote and Ryan 2014).

A final set of facts illustrated by Figure 4 is that overall wage growth was anemic throughout the 2000s, even prior to the Great Recession. Between 1999 and 2007, real wage changes were negative below approximately the 15th percentile, and were below 5 percentage points up to the 70th percentile of the distribution. Indeed, wage growth was greater at all percentiles during both the 1980s and 1990s than in the *pre-recession* 2000s.⁷ Of course, wage growth was essentially zero at all percentiles from 2007 to 2012.

Why are the rapidly rising earnings of the top 1 percent (as discussed in Atkinson, Piketty, and Saez 2011, for example) not strongly evident in Figure 4? One reason reflects substance; another is an artifact of the data. Substantively, the plot depicts changes in earnings by *occupational* percentile rather than *wage* percentile. Wage growth by occupational percentile is less concentrated than wage growth across wage percentiles because the highest earners are found across a variety of occupations. In addition, the very highest percentiles of earnings are censored in public use Census and American Community Survey data files, which further masks earnings gains at extreme quantiles.


The Recent Slowdown in the Growth of High-Skill Occupations

The hypothesis that automation and information technology has led to occupational and, to a lesser degree, wage polarization in the US labor force can explain some key features of the US and the cross-national data. But reality invariably proves more complicated than any single theory anticipates.

For my thesis linking technological change to occupational change, one concern is the unexplained deceleration of employment growth in abstract task-intensive occupations after 2000 (Beaudry, Green, and Sand 2014, forthcoming; Mishel, Shierholz, and Schmitt 2013). Figure 5 follows the format of Figure 4 but instead of showing (approximate) percentage changes in wages on the vertical axis, it shows percentage changes in the employment share of the jobs ranked by their skill level in 1979. Since the sum of shares must equal one at any time period, the changes in these shares across the decades must total zero, and thus, the height at each skill percentile measures the growth in each occupation's employment relative to the whole.

⁷ Because the 2000–2007 interval is two years shorter than the 1979–1989 period, one should multiply the later changes by 1.25 to put them on the same temporal footing. But even after making such an adjustment, wage growth was still considerably weaker at all percentiles from 2000–2007 than in the earlier two decades.

Figure 5

Smoothed Employment Changes by Occupational Skill Percentile, 1979–2012


Sources: Author, calculated using 1980, 1990, and 2000 Census Integrated Public Use Microdata Series (IPUMS) files; American Community Survey combined file 2006–2008, American Community Survey 2012.

Notes: The figure plots changes in employment shares by 1980 occupational skill percentile rank using a locally weighted smoothing regression (bandwidth 0.8 with 100 observations), where skill percentiles are measured as the employment-weighted percentile rank of an occupation's mean log wage in the Census IPUMS 1980 5 percent extract. Employment in each occupation is calculated using workers' hours of annual labor supply times the Census sampling weights. Consistent occupation codes for Census years 1980, 1990, and 2000, and 2008 are from Autor and Dorn (2013).

Figure 5 contributes three nuances to the occupational polarization story above. First, the pace of employment gains in low-wage, manual task-intensive jobs has risen successively across periods, as shown at the left-hand side of the figure. Second, the occupations that are losing employment share appear to be increasingly drawn from higher ranks of the occupational distribution. For example, the highest ranked occupation to lose employment share during the 1980s lay at approximately the 45th percentile of the skill distribution. In the final two subperiods, this rank rose still further to above the 75th percentile—suggesting that the locus of displaced middle-skill employment is moving into higher-skilled territories. Third, growth of high-skill, high-wage occupations (those associated with abstract work) decelerated markedly in the 2000s, with no relative growth in the top two deciles of the occupational skill distribution during 1999 through 2007, and only a modest recovery between 2007 and 2012. Stated plainly, the growth of occupational employment across skill levels looks U-shaped earlier in the period, with gains at low-skill and high-skill levels. By the 2000s, the pattern of occupational employment across

Figure 6

Private Fixed Investment in Information Processing Equipment and Software as a Percentage of Gross Domestic Product, 1949–2014



Source: FRED, Federal Bank of St. Louis. <http://research.stlouisfed.org/fred2/graph/?g=GXc> (accessed 8/3/2014).

skill levels began to resemble a downward ramp. In Autor (2015), I present a more detailed breakdown of these patterns, and in particular suggest that the set of abstract task-intensive jobs is not growing as rapidly as the potential supply of highly educated workers.

What explains the slowing growth of abstract task-intensive employment? One interpretation is that automation, information technology, and technological progress in general are encroaching upward in the task domain and beginning to substitute strongly for the work done by professional, technical, and managerial occupations. While one should not dismiss this possibility out of hand, it doesn't fit well with the pattern of computer and software investment. If information technology is increasingly replacing workers high in the skill distribution, one would expect a surge of corporate investment in computer hardware and software. Instead, Figure 6 shows that in early 2014, information processing equipment and software investment was only 3.5 percent of GDP, a level last seen in 1995 at the outset of the "dot-com" era. To me, the evidence in Figure 6 suggests a temporary dislocation of demand for information technology capital during the latter half of the 1990s, followed by a sharp correction after 2000. I suspect that the huge falloff in

information investment may have dampened innovative activity and demand for high-skilled workers more broadly.

As noted earlier, technological change is far from the only factor affecting US labor markets in the last 15 years. For example, the deceleration of wage growth and changes in occupational patterns in the US labor market after 2000, and further after 2007, is surely associated to some extent with two types of macroeconomic events. First, there are the business cycle effects—the bursting of the “dot-com” bubble in 2000, and the collapse of the housing market and the ensuing financial crisis in 2007–2008—both of which curtailed investment and innovative activity. Second, there are the employment dislocations in the US labor market brought about by rapid globalization, particularly the sharp rise of import penetration from China following its accession to the World Trade Organization in 2001 (Autor, Dorn, and Hanson 2013; Pierce and Schott 2012; Acemoglu, Autor, Dorn, Hanson, and Price forthcoming). China’s rapid rise to a premier manufacturing exporter had far-reaching impacts on US workers, reducing employment in directly import-competing US manufacturing industries and depressing labor demand in both manufacturing and nonmanufacturing sectors that served as upstream suppliers to these industries.

Of course, these forces are in various ways linked with the spread of automation and technology. Advances in information and communications technologies have changed job demands in US workplaces directly and also indirectly, by making it increasingly feasible and cost-effective for firms to source, monitor, and coordinate complex production processes at disparate locations worldwide and altering competitive conditions for US manufacturers and workers. This multidimensional complementarity among causal factors makes it both conceptually and empirically difficult to isolate the “pure” effect of any one factor.

Polanyi’s Paradox: Will It Be Overcome?

Automation, complemented in recent decades by the exponentially increasing power of information technology, has driven changes in productivity that have disrupted labor markets. This essay has emphasized that jobs are made up of many tasks and that while automation and computerization can substitute for some of them, understanding the interaction between technology and employment requires thinking about more than just substitution. It requires thinking about the range of tasks involved in jobs, and how human labor can often complement new technology. It also requires thinking about price and income elasticities for different kinds of output, and about labor supply responses.

The tasks that have proved most vexing to automate are those demanding flexibility, judgment, and common sense—skills that we understand only tacitly. I referred to this constraint above as Polanyi’s paradox. In the past decade, computerization and robotics have progressed into spheres of human activity that were considered off limits only a few years earlier—driving vehicles, parsing legal documents, even

performing agricultural field labor. Is Polanyi's paradox soon to be at least mostly overcome, in the sense that the vast majority of tasks will soon be automated?⁸

My reading of the evidence suggests otherwise. Indeed, Polanyi's paradox helps to explain what has *not* yet been accomplished, and further illuminates the paths by which more *will* ultimately be accomplished. Specifically, I see two distinct paths that engineering and computer science can seek to traverse to automate tasks for which we "do not know the rules": environmental control and machine learning. The first path circumvents Polanyi's paradox by regularizing the environment, so that comparatively inflexible machines can function semi-autonomously. The second approach inverts Polanyi's paradox: rather than teach machines rules that we do not understand, engineers develop machines that attempt to infer tacit rules from context, abundant data, and applied statistics.

Environmental Control

Most automated systems lack flexibility—they are brittle. Modern automobile plants, for example, employ industrial robots to install windshields on new vehicles as they move through the assembly line. But aftermarket windshield replacement companies employ technicians, not robots, to install replacement windshields. Evidently, the tasks of removing a broken windshield, preparing the windshield frame to accept a replacement, and fitting a replacement into that frame demand more real-time adaptability than any contemporary robot can cost-effectively approach.

The distinction between assembly line production and the in-situ repair highlights the role of environmental control in enabling automation. Engineers can in some cases radically simplify the environment in which machines work to enable autonomous operation, as in the familiar example of a factory assembly line. Numerous examples of this approach to environmental regularization are so ingrained in daily technology that they escape notice, however. To enable the operation of present-day automobiles, for example, humanity has adapted the naturally occurring environment by leveling, re-grading, and covering with asphalt a nontrivial percentage of the earth's land surface.⁹

The ongoing automation of warehouses provides another example. Large online retailers, such as Amazon.com, Zappos.com, and Staples, operate systems of warehouses that have traditionally employed legions of dexterous, athletic "pickers," who run and climb through shelves of typically non-air-conditioned warehouses to locate, collect, box, label, and ship goods. There is at present no cost-effective robotic

⁸ For a glimpse of the view that just about anything can now be computerized, see the widely cited (albeit unpublished) article by the economists Carl Frey and Michael Osborne, who write (2013, p. 24) that, "recent developments in ML [machine learning] and MR [mobile robotics], building upon big data, allow for pattern recognition, and thus enable computer capital to rapidly substitute for labour across a wide range of non-routine tasks. Yet some inhibiting engineering bottlenecks to computerization persist. Beyond these bottlenecks, however, we argue that it is largely already technologically possible to automate almost any task, provided that sufficient amounts of data are gathered for pattern recognition."

⁹ According to Wikipedia, so-called impervious surfaces (mostly roads and parking lots) cover 43,000 square miles of land in the lower 48 United States—roughly equal to the land area of the state of Ohio (http://en.wikipedia.org/wiki/Impervious_surface, accessed 8/4/2014).

facsimile for these human pickers. The job's steep requirements for flexibility, object recognition, physical dexterity, and fine motor coordination are too formidable.

But large components of warehousing can be automated, as demonstrated by Kiva Systems, a robotic warehousing startup that was purchased by Amazon in 2012. The core of the Kiva system is a dispatch program that oversees the flow of all goods through the warehouse, coordinating the work of robots, which carry shelves, with the work of humans. As objects arrive at the facility for stocking, the dispatch software directs robots to transport and line up empty shelves to a loading area, where *human* stockers place merchandise on shelves. Robots then carry the loaded shelves back to a storage warehouse, where the dispatch software directs their placement to optimize product availability for expected product demand. As new orders arrive, the dispatch software sends robots to retrieve shelves and lines them up in a packing area. Then a human picker, directed by a laser pointer controlled by the dispatch software, takes objects from the assembled shelves, packs them in shipping boxes, applies a shipping label, and drops the package in a chute for delivery. As items are picked, the robots take the shelves away until needed again for packing or restocking. Thus, in a Kiva-operated warehouse, robots handle only the routine task of moving shelves across a level surface; workers handle merchandise; and the dispatch software coordinates the activity.

While Kiva Systems provides a particularly clear example of exploiting environmental control to extend the reach of automation, the same principle is often lurking behind more sophisticated packaging. Perhaps the least recognized—and most mythologized—is the self-driving Google Car. Computer scientists sometimes remark that the Google car does not drive on roads, but rather on maps. A Google car navigates through the road network primarily by comparing its real-time audio-visual sensor data against painstakingly hand-curated maps that specify the exact locations of all roads, signals, signage, and obstacles. The Google car adapts in real time to obstacles, such as cars, pedestrians, and road hazards, by braking, turning, and stopping. But if the car's software determines that the environment in which it is operating differs from the environment that has been preprocessed by its human engineers—when it encounters an unexpected detour or a crossing guard instead of a traffic signal—the car requires its human operator to take control. Thus, while the Google car appears outwardly to be adaptive and flexible, it is somewhat akin to a train running on invisible tracks.

These examples highlight both the limitations of current technology to accomplish nonroutine tasks, and the capacity of human ingenuity to surmount some of these obstacles by re-engineering the environment in which work tasks are performed.

Machine Learning

Polanyi's paradox—"we know more than we can tell"—presents a challenge for computerization because, if people understand how to perform a task only tacitly and cannot "tell" a computer how to perform the task, then seemingly programmers cannot automate the task—or so the thinking has gone. But this understanding

is shifting rapidly due to advances in machine learning. Machine learning applies statistics and inductive reasoning to supply best-guess answers where formal procedural rules are unknown. Where engineers are unable to program a machine to “simulate” a nonroutine task by following a scripted procedure, they may nevertheless be able to program a machine to master the task autonomously by studying successful examples of the task being carried out by others. Through a process of exposure, training, and reinforcement, machine learning algorithms may potentially infer how to accomplish tasks that have proved dauntingly challenging to codify with explicit procedures.

As a concrete example, consider the task of visually identifying a chair (discussed in Autor, forthcoming). An engineer applying a conventional rules-based programming paradigm might attempt to specify what features of an object qualify an object as a chair—it possesses legs, arms, a seat, and a back, for example. But one would soon discover that many chairs do not possess all of these features (for example, some chairs have no back, or no arms). If the engineer then relaxed the required feature set accordingly (chair back optional), the included set would grow to encompass many objects that are not chairs, such as small tables. The canonical approach to recognizing objects by pre-specifying requisite features—and more sophisticated variants of this approach—would likely have very high misclassification rates. Yet, any grade-school child could perform this task with high accuracy. What does the child know that the rules-based procedure does not? Unfortunately, we cannot enunciate precisely what the child knows—and this is precisely Polanyi’s paradox.

Machine learning potentially circumvents this problem. Relying on large databases of so-called “ground truth”—a vast set of curated examples of labeled objects—a machine learning algorithm attempts to infer what attributes of an object make it more or less likely to be designated a chair. This process is called “training.” When training is complete, the machine can apply this statistical model to attempt to identify chairs that are distinct from those in the original dataset. If the statistical model is sufficiently good, it may be able to recognize chairs that are somewhat distinct from those in the original training data, like chairs of different shapes, materials, or dimensions. Machine learning does not require an explicit physical model of “chairness.” At its core, machine learning is an atheoretical brute force technique—what psychologists call “dustbowl empiricism”—requiring only large training databases, substantial processing power, and, of course, sophisticated software.¹⁰

How well does machine learning work in practice? If you use a search engine or Google Translate, operate a smartphone with voice commands, or follow movie suggestions from Netflix, you can assess for yourself how successfully these technologies function. For example, if the majority of users who recently searched for the terms “degrees bacon” clicked on links for Kevin Bacon rather than links for best bacon cooking temperatures, the search engine would tend to place the Kevin Bacon links higher in the list of results. My general observation is that

¹⁰ Varian (2014) provides an introduction to machine learning techniques for economists.

the tools are inconsistent: uncannily accurate at times; typically only so-so; and occasionally unfathomable. Moreover, an irony of machine learning algorithms is that they also cannot “tell” programmers why they do what they do. IBM’s Watson computer famously triumphed in the trivia game of *Jeopardy* against champion human opponents. Yet Watson also produced a spectacularly incorrect answer during its winning match. Under the category of *US Cities*, the question was, “Its largest airport was named for a World War II hero; its second largest, for a World War II battle.” Watson’s proposed answer was Toronto, a city in Canada. Even leading-edge accomplishments in this domain can appear somewhat underwhelming. A 2012 *New York Times* article (Markoff 2012) described Google’s X Lab’s recent project (Le et al. 2012) to apply a neural network of 16,000 processors to identify images of cats on YouTube. The article’s headline ruefully poses the question, “How Many Computers to Identify a Cat? 16,000.”

Since the underlying technologies—the software, hardware, and training data—are all improving rapidly (Andreopoulos and Tsotsos 2013), one should view these examples as prototypes rather than as mature products. Some researchers expect that as computing power rises and training databases grow, the brute force machine learning approach will approach or exceed human capabilities. Others suspect that machine learning will only ever “get it right” on average, while missing many of the most important and informative exceptions. Ultimately, what makes an object a chair is that it is purpose-built for a human being to sit upon. Machine-learning algorithms may have fundamental problems with reasoning about “purposiveness” and intended uses, even given an arbitrarily large training database of images (Grabner, Gall, and Van Gool 2011). One is reminded of Carl Sagan’s (1980, p. 218) remark, “If you wish to make an apple pie from scratch, you must first invent the universe.”

Conclusions

Major newspaper stories offer fresh examples daily of technologies that substitute for human labor in an expanding—although still circumscribed—set of tasks. The offsetting effects of complementarities and rising demand in other areas are, however, far harder to identify as they occur. My own prediction is that employment polarization will *not* continue indefinitely (as argued in Autor 2013). While some of the *tasks* in many current middle-skill jobs are susceptible to automation, many middle-skill *jobs* will continue to demand a mixture of tasks from across the skill spectrum. For example, medical support occupations—radiology technicians, phlebotomists, nurse technicians, and others—are a significant and rapidly growing category of relatively well-remunerated, middle-skill employment. Most of these occupations require mastery of “middle-skill” mathematics, life sciences, and analytical reasoning. They typically require at least two years of postsecondary vocational training, and in some cases a four-year college degree or more. This broad description also fits numerous skilled trade and repair occupations, including plumbers, builders, electricians, heating/ventilating/air-conditioning installers, and

automotive technicians. It also fits a number of modern clerical occupations that provide coordination and decision-making functions, rather than simply typing and filing, like a number of jobs in marketing. There are also cases where technology is enabling workers with less esoteric technical mastery to perform additional tasks: for example, the nurse practitioner occupation that increasingly performs diagnosing and prescribing tasks in lieu of physicians.

I expect that a significant stratum of middle-skill jobs combining specific vocational skills with foundational middle-skills levels of literacy, numeracy, adaptability, problem solving, and common sense will persist in coming decades. My conjecture is that many of the tasks currently bundled into these jobs cannot readily be unbundled—with machines performing the middle-skill tasks and workers performing only a low-skill residual—without a substantial drop in quality. This argument suggests that many of the middle-skill jobs that persist in the future will combine routine technical tasks with the set of nonroutine tasks in which workers hold comparative advantage: interpersonal interaction, flexibility, adaptability, and problem solving. In general, these same demands for interaction frequently privilege face-to-face interactions over remote performance, meaning that these same middle-skill occupations may have relatively low susceptibility to offshoring. Lawrence Katz memorably titles workers who virtuously combine technical and interpersonal tasks as “the new artisans” (see Friedman 2010), and Holzer (2015) documents that “new middle skill jobs” are in fact growing rapidly, even as traditional production and clerical occupations contract.¹¹

This prediction has one obvious catch: the ability of the US education and job training system (both public and private) to produce the kinds of workers who will thrive in these middle-skill jobs of the future can be called into question. In this and other ways, the issue is not that middle-class workers are doomed by automation and technology, but instead that human capital investment must be at the heart of any long-term strategy for producing skills that are complemented by rather than substituted for by technological change. In 1900, the typical young, native-born American had only a common school education, about the equivalent of sixth to eighth grades. By the late 19th century, many Americans recognized that this level of schooling was inadequate: farm employment was declining, industry was rising, and their children would need additional skills to earn a living. The United States responded to this challenge over the first four decades of the 20th century by becoming the first nation in the world to deliver universal high school education to its citizens (Goldin and Katz 2008). Tellingly, the high school movement was led by the farm states. Societal adjustments to earlier waves of technological advancement were neither rapid, automatic, nor cheap. But they did pay off handsomely.

¹¹ A creative paper by Lin (2011) studies the growth of “new work” by documenting the differential growth of US employment in newly introduced Census occupation codes during the 1980s and 1990s in high-education and high-technology cities. New occupational titles are generally clustered across two categories: those associated with using new technologies such as web developer or database administrator; and novel personal services, such as personal chefs and stylists.

A final point, typically neglected in recent dismal prophecies of machine-human substitution, is that if human labor is indeed rendered superfluous by automation, then our chief economic problem will be one of distribution, not of scarcity. The primary system of income distribution in market economies is rooted in labor scarcity; citizens possess (or acquire) a bundle of valuable “human capital” that, due to its scarcity, generates a flow of income over the career path. If machines were in fact to make human labor superfluous, we would have vast aggregate wealth but a serious challenge in determining who owns it and how to share it. One might presume that with so much wealth at hand, distribution would be relatively straightforward to resolve. But history suggests that this prediction never holds true. There is always perceived scarcity and ongoing conflict over distribution, and I do not expect that this problem will become any less severe as automation advances. Are we actually on the verge of throwing off the yoke of scarcity so that our primary economic challenge soon becomes one of distribution? Here, I recall the observations of economist, computer scientist, and Nobel laureate Herbert Simon (1966), who wrote at the time of the automation anxiety of the 1960s: “Insofar as they are economic problems at all, the world’s problems in this generation and the next are problems of scarcity, not of intolerable abundance. The bogeyman of automation consumes worrying capacity that should be saved for real problems . . .” A half century on, I believe the evidence favors Simon’s view.

■ *This paper draws from an essay prepared for the Federal Reserve Bank of Kansas City’s economic policy symposium on “Re-Evaluating Labor Market Dynamics,” August 21–23, 2014, in Jackson Hole, Wyoming (Autor 2015) as well as the essay “The Paradox of Abundance: Automation Anxiety Returns” (Autor forthcoming). I thank Erik Brynjolfsson, Chris Foote, Frank Levy, Lisa Lynch, Andrew McAfee, Brendan Price, Seth Teller, Dave Wessel, participants in the MIT CSAIL/Economists Lunch Seminar, and the editors of this journal for insights that helped to shape my thinking on this subject. I thank Sookyo Jeong and Brendan Price for superb research assistance.*

References

- Acemoglu, Daron, David H. Autor, David Dorn, Gordon Hanson, and Brendan Price.** Forthcoming. “Import Competition and the Great U.S. Employment Sag of the 2000s.” *Journal of Labor Economics*.
- Akst, Daniel.** 2013. “What Can We Learn from Past Anxiety over Automation?” *The Wilson Quarterly*, Summer.
- Alpert, Andrew, and Jill Auyer.** 2003. “Evaluating the BLS 1988–2000 Employment Projections.” *Monthly Labor Review*, October, pp. 13–37.
- Andreopoulos, Alexander, and John K. Tsotsos.** 2013. “50 Years of Object Recognition: Directions Forward.” *Computer Vision and Image Understanding* 117(8): 827–91.

- Atkinson, Anthony B., Thomas Piketty, and Emmanuel Saez. 2011. "Top Incomes in the Long Run of History." *Journal of Economic Literature* 49(1): 3–71.
- Autor, David H. 2013. "The 'Task Approach' to Labor Markets: An Overview." *Journal for Labour Market Research* 46(3): 185–99.
- Autor, David H. 2014. "Skills, Education, and the Rise of Earnings Inequality among the 'Other 99 Percent.'" *Science* 344(6186): 843–51.
- Autor, David H. 2015. "Polanyi's Paradox and the Shape of Employment Growth." In *Re-Evaluating Labor Market Dynamics*, pp. 129–79. Federal Reserve Bank of Kansas City.
- Autor, David H. Forthcoming. "The Paradox of Abundance: Automation Anxiety Returns." In *Performance and Progress: Essays on Capitalism, Business and Society*, edited by Subramanian Rangan. London: Oxford University Press.
- Autor, David H., and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review* 103(5): 1553–97.
- Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103(6): 2121–68.
- Autor, David H., David Dorn, and Gordon H. Hanson. 2015. "Untangling Trade and Technology: Evidence from Local Labor Markets." *Economic Journal* 125(584): 621–46.
- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2006. "The Polarization of the U.S. Labor Market." *American Economic Review* 96(2): 189–94.
- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2008. "Trends in U.S. Wage Inequality: Revising the Revisionists." *Review of Economics and Statistics* 90(2): 300–323.
- Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics* 118(4): 1279–1333.
- Baumol, William J. 1967. "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis." *American Economic Review* 57(3): 415–26.
- Beaudry, Paul, David A. Green, and Benjamin M. Sand. 2014. "The Declining Fortunes of the Young since 2000." *American Economic Review* 104(5): 381–86.
- Beaudry, Paul, David A. Green, and Benjamin M. Sand. Forthcoming. "The Great Reversal in the Demand for Skill and Cognitive Tasks." *Journal of Labor Economics*.
- Bessen, James. 2015. "Toil and Technology." *Finance and Development* 52(1).
- Bowen, Harold R. (Chairman). 1966. "Report of the National Commission on Technology, Automation, and Economic Progress: Volume I." Washington: U.S. Government Printing Office.
- Brynjolfsson, Erik, and Andrew McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York and London: W.W. Norton & Company.
- Caroli, Eve, and John Van Reenen. 2001. "Skill-Biased Organizational Change? Evidence from a Panel of British and French Establishments." *Quarterly Journal of Economics* 116(4): 1449–92.
- Clark, Colin. 1951. *The Conditions of Economic Progress*, 2nd edition. London: Macmillan.
- Cortes, Guido Matias, Nir Jaimovich, Christopher J. Nekarda, and Henry E. Siu. 2014. "The Micro and Macro of Disappearing Routine Jobs: A Flows Approach." NBER Working Paper 20307, July.
- Douglas, Paul H. 1930. *Real Wages in the United States, 1890–1926*. New York: Houghton Mifflin Company.
- Foote, Christopher L., and Richard W. Ryan. 2014. "Labor-Market Polarization over the Business Cycle." Public Policy Discussion Paper 12-8, Federal Reserve Bank of Boston, April.
- Frey, Carl Benedikt, and Michael A. Osborne. 2013. "The Future of Employment: How Susceptible are Jobs to Computerization?" Oxford Martin School, September.
- Friedman, Thomas L. 2010. "The Election That Wasn't." *New York Times*, October 23.
- Goldin, Claudia, and Lawrence F. Katz. 2008. *The Race between Education and Technology*. Cambridge: Harvard University Press.
- Goos, Maarten, and Alan Manning. 2003. "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain." Center for Economic Performance Discussion Papers DP0604, December.
- Goos, Maarten, and Alan Manning. 2007. "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain." *Review of Economics and Statistics* 89(1): 118–33.
- Goos, Maarten, Alan Manning, and Anna Salomons. 2014. "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring." *American Economic Review* 104(8): 2509–26.
- Grabner, Helmut, Juergen Gall, and Luc Van Gool. 2011. "What Makes a Chair a Chair?" In *Computer Vision and Pattern Recognition 2011*, IEEE Conference, June 20–25, pp. 1529–36.
- Graetz, Georg, and Guy Michaels. 2015. "Robots at Work." IZA Discussion Papers 8938, Institute for the Study of Labor (IZA).
- Herald Press, The. 1966. "Skirting the Automation Question." February 7, p. 2. (*The Herald Press* of St. Joseph, Michigan.)

- Hoffman, William, and Leo Furcht.** 2014. *The Biologist's Imagination: Innovation in the Biosciences*. New York, NY: Oxford University Press.
- Holzer, Harry J.** 2015. "Job Market Polarization and U.S. Worker Skills: A Tale of Two Middles." Brookings Institution Economic Studies Working Paper, April 6.
- Hsieh, Chang-Tai, and Enrico Moretti.** 2003. "Can Free Entry Be Inefficient? Fixed Commissions and Social Waste in the Real Estate Industry." *Journal of Political Economy* 111(5): 1076–1122.
- Jackson, Kristin.** 1993. "The World's First Motel Rests upon Its Memories." *Seattle Times*, April 25.
- Jevons, William Stanley.** 1865. *The Coal Question; An Inquiry Concerning the Progress of the Nation, and the Probable Exhaustion of Our Coal Mines*. London: Macmillan and Co.
- Katz, Lawrence F., and Robert A. Margo.** 2014. "Technical Change and the Relative Demand for Skilled Labor: The United States in Historical Perspective." In *Human Capital in History*, edited by Leah Platt Bouston, Carola Frydman, and Robert A. Margo, 15–57. University of Chicago Press.
- Kremer, Michael.** 1993. "The O-Ring Theory of Economic Development." *Quarterly Journal of Economics* 108(3): 551–75.
- Le, Quoc V., Marc' Aurelio Ranzato, Rajat Monga, Matthieu Devin, Kai Chen, Greg S. Corrado, Jeff Dean, and Andrew Y. Ng.** 2012. "Building High-Level Features Using Large Scale Unsupervised Learning." In *Proceedings of the 29th International Conference on Machine Learning*, June 26–July 1, Edinburgh, Scotland, UK.
- Levy, Frank and Richard J. Murnane.** 2004. *The New Division of Labor: How Computers Are Creating the Next Job Market*. Princeton University Press.
- Lin, Jeffrey.** 2011. "Technological Adaptation, Cities, and New Work." *Review of Economics and Statistics* 93(2): 554–74.
- Markoff, John.** 2012. "How Many Computers to Identify a Cat? 16,000." *New York Times*, June 25.
- Mazzolari, Francesca, and Giuseppe Ragusa.** 2013. "Spillovers from High-Skill Consumption to Low-Skill Labor Markets." *Review of Economics and Statistics* 95(1): 74–86.
- Michaels, Guy, Ashwini Natraj, and John Van Reenen.** 2014. "Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years." *Review of Economics and Statistics* 96(1): 60–77.
- Mishel, Lawrence, Heidi Shierholz, and John Schmitt.** 2013. "Don't Blame the Robots: Assessing the Job Polarization Explanation of Growing Wage Inequality." EPI-CEPR Working Paper, November 19.
- Moravec, Hans.** 1988. *Mind Children: The Future of Robot and Human Intelligence*. Harvard University Press.
- Nordhaus, William D.** 2007. "Two Centuries of Productivity Growth in Computing." *Journal of Economic History* 67(1): 17–22.
- Pierce, Justin R., and Peter K. Schott.** 2012. "The Surprisingly Swift Decline of U.S. Manufacturing Employment." NBER Working Paper 18655.
- Polanyi, Michael.** 1966. *The Tacit Dimension*. New York: Doubleday.
- Ruggles, Steven, Trent J. Alexander, Katie Genadek, Ronald Goeken, Matthew B. Schroeder, and Matthew Sobek.** 2010. *Integrated Public Use Microdata Series: Version 5.0* [Machine-readable database]. University of Minnesota.
- Sachs, Jeffrey D., Seth G. Benzell, Guillermo LaGarda.** 2015. "Robots: Curse or Blessing? A Basic Framework." NBER Working Paper 21091, April.
- Sachs, Jeffrey D., and Laurence J. Kotlikoff.** 2012. "Smart Machines and Long-Term Misery." NBER Working Paper 18629, December.
- Sagan, Carl.** 1980. *Cosmos*. New York: Random House.
- Simon, Herbert A.** 1966. "Automation" (a letter in response to "Where Do We Go From Here?" March 17, 1966 issue) *New York Review of Books*, May 26.
- Smith, Christopher L.** 2013. "The Dynamics of Labor Market Polarization." Federal Reserve Board, Finance and Economics Discussion Series No. 2013-57, August.
- TIME.** 1961. "The Automation Jobless." February 24.
- US Bureau of the Census.** 1949. *Historical Statistics of the United States, 1789–1945*. U.S. Government Printing Office, Series D 134-144.
- Varian, Hal R.** 2014. "Big Data: New Tricks for Econometrics." *Journal of Economic Perspectives* 28(2): 3–28.