

NBER WORKING PAPER SERIES

POLANYI'S PARADOX AND THE SHAPE OF EMPLOYMENT GROWTH

David Autor

Working Paper 20485

<http://www.nber.org/papers/w20485>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

September 2014

This paper is a draft 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. I thank Erik Brynjolfsson, Chris Foote, Frank Levy, Lisa Lynch, Andrew McAfee, Brendan Price, Seth Teller, Dave Wessel, and participants in the MIT CSAIL/Economists Lunch Seminar for insights that helped to shape the paper. I thank Sookyo Jeong and Brendan Price for superb research assistance. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2014 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.

Polanyi's Paradox and the Shape of Employment Growth

David Autor

NBER Working Paper No. 20485

September 2014

JEL No. J23,J24,J31,O3

ABSTRACT

In 1966, the philosopher Michael Polanyi observed, “We can know more than we can tell... The skill of a driver cannot be replaced by a thorough schooling in the theory of the motorcar; the knowledge I have of my own body differs altogether from the knowledge of its physiology.” Polanyi’s observation largely predates the computer era, but the paradox he identified—that our tacit knowledge of how the world works often exceeds our explicit understanding—foretells much of the history of computerization over the past five decades. This paper offers a conceptual and empirical overview of this evolution. I begin by sketching the historical thinking about machine displacement of human labor, and then consider the contemporary incarnation of this displacement—labor market polarization, meaning the simultaneous growth of high-education, high-wage and low-education, low-wages jobs—a manifestation of Polanyi’s paradox. I discuss both the explanatory power of the polarization phenomenon and some key puzzles that confront it. I then reflect on how recent advances in artificial intelligence and robotics should shape our thinking about the likely trajectory of occupational change and employment growth. A key observation of the paper is that journalists and expert commentators overstate the extent of machine substitution for human labor and ignore the strong complementarities. The challenges to substituting machines for workers in tasks requiring adaptability, common sense, and creativity remain immense. Contemporary computer science seeks to overcome Polanyi’s paradox by building machines that learn from human examples, thus inferring the rules that we tacitly apply but do not explicitly understand.

David Autor

Department of Economics, E17-216

MIT

77 Massachusetts Avenue

Cambridge, MA 02139

and NBER

dautor@mit.edu

I. Introduction

That computers are ubiquitous in contemporary life is self-evident. The share of information processing equipment and software in private, non-residential investment rose from approximately 8 percent to more than 30 percent between 1950 and 2012, with the largest leap occurring between 1990 and 2000.¹ It is hard to think of a prior historical episode where a single category of capital investment came so rapidly to dominate all others, now accounting for close to one in three business investment dollars.²

Given their ubiquity, it is tempting to infer that there is no task to which computers are not suited. But that leap of logic is unfounded. Human tasks that have proved most amenable to computerization are those that follow explicit, codifiable procedures—such as multiplication—where computers now vastly exceed human labor in speed, quality, accuracy, and cost efficiency.³ Tasks that have proved most vexing to automate are those that demand flexibility, judgment, and common sense—skills that we understand only tacitly—for example, developing a hypothesis or organizing a closet. In these tasks, computers are often less sophisticated than preschool age children. 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. Understanding this interplay is central to interpreting and forecasting the changing structure of employment in the U.S. and other industrialized countries. This understanding is also at the heart of the increasingly prominent debate about whether the rapid pace of automation threatens to render the demand for human labor obsolete over the next several decades.

¹ Author's calculations based on Bureau of Economic Analysis National Income and Product Account data.

² Fences used to demarcate property lines and secure livestock may be another example. In 1872, the value of fencing capital stock in the United States was roughly equal to the value of all livestock, the national debt, or the railroads, and annual fencing repair costs were greater than combined annual tax receipts at all levels of government (Hornbeck, 2010).

³ At the extreme, there are tasks that computers regularly perform that could not be accomplished by human labor at any price, for example, guiding a missile to intercept another missile in mid-flight.

This paper offers a conceptual and empirical overview of the evolving relationship between computer capability and human skill demands. I begin by sketching the historical thinking about machine displacement of human labor, and then consider the contemporary incarnation of this displacement—labor market polarization, meaning the simultaneous growth of high-education, high-wage and low-education, low-wages jobs—a manifestation of Polanyi’s paradox. I discuss both the explanatory power of the polarization phenomenon and some key puzzles that confront it. I finally reflect on how recent advances in artificial intelligence and robotics should shape our thinking about the likely trajectory of occupational change and employment growth.

A key observation of the paper is that journalists and expert commentators overstate the extent of machine substitution for human labor and ignore the strong complementarities that increase productivity, raise earnings, and augment demand for skilled labor. The challenges to substituting machines for workers in tasks requiring flexibility, judgment, and common sense remain immense. Contemporary computer science seeks to overcome Polanyi’s paradox by building machines that learn from human examples, thus inferring the rules that we tacitly apply but do not explicitly understand.

II. A Brief History of Automation Anxiety

Anxiety about the adverse effects of technological change on employment has a venerable history.⁴ In the early 19th century, a group of English textile artisans calling themselves the Luddites staged a machine-trashing rebellion in protest of the rapid automation of textile production, which they feared jeopardized their livelihoods. Their actions earned the term Luddite an (unflattering) entry in the popular lexicon. Economists have historically rejected the concerns of the Luddites as an example of the “lump of labor” fallacy, the supposition that an increase in labor productivity inevitably reduces employment because there is only a finite amount of work to do. While intuitively appealing, the notion that productivity gains reduce employment has received little historical support. The employment-to-population ratio, for

⁴ This section of the paper draws heavily on Autor (2014a), with some paragraphs quoted directly.

example, rose over the course of the 20th century as women moved from home to market, and the unemployment rate fluctuated cyclically, with no long-term increase.

Yet, despite sustained increases in material standards of living, fear of the adverse employment consequences of technological advancement has recurred repeatedly in the 20th century. In his widely discussed Depression-era essay “Economic Possibilities for our Grandchildren,” John Maynard Keynes (1930) foresaw that in a century’s time, “we may be able to perform all the operations of agriculture, mining, and manufacture with a quarter of the human effort to which we have been accustomed.” Keynes viewed these developments as posing short-term challenges, “For the moment the very rapidity of these changes is hurting us and bringing difficult problems to solve. . . . We are being afflicted with a new disease of which some readers may not yet have heard the name, but of which they will hear a great deal in the years to come—namely, *technological unemployment*.” But Keynes was sanguine about the long run, opining that “this is only a temporary phase of maladjustment,” and predicting that the fifteen-hour workweek (supporting a high standard of living) would be commonplace in a century’s time.


Keynes’ projection that the maladjustment was “temporary” was a bold one given that he was writing during the Great Depression. But the end of the Second World War seemed to affirm the rising prosperity that Keynes had foreseen. Perhaps more surprising is that “automation anxiety” recurred two decades after the Second World War, during what was arguably the height of American economic preeminence. In 1964, President Johnson empaneled a “Blue-Ribbon National Commission on Technology, Automation, and Economic Progress” whose charge was “to identify and assess the past effects and the current and prospective role and pace of technological change; to identify and describe the impact of technological and economic change on production and employment, including new job requirements and the major types of worker displacement, both technologically and economic, which are likely to occur during the next 10 years.”

While the commission ultimately concluded that automation did not threaten employment at that time, it recommended as insurance against this possibility, “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" (*The Herald Post*, 1966).

The blue-ribbon commission's sanguine conclusions did not entirely allay the concerns of contemporary social critics. In an open letter to President Johnson in 1966, the self-titled Ad Hoc Committee on the Triple Threat, which included Nobel laureates Linus Pauling (chemistry) and Gunnar Myrdal (economics), as well as economic historian Robert Heilbroner, opined that "The traditional link between jobs and incomes is being broken. . . . The economy of abundance can sustain all citizens in comfort and economic security whether or not they engage in what is commonly reckoned as work" (quoted in Akst, 2013).⁵ Writing separately in the *Public Interest* in 1965, Heilbroner argued that, "the new technology is threatening a whole new group of skills—the sorting, filing, checking, calculating, remembering, comparing, okaying skills—that are the special preserve of the office worker. . . . In the end, as machines continue to invade society, duplicating greater and greater numbers of social tasks, it is human labor itself—at least, as we now think of "labor"—that is gradually rendered redundant" (pp. 34–36).

Figure 1. Chicago Booth IGM Expert Poll: Impact of Automation on Employment and Wages



⁵ The three threats perceived by the ad hoc committee were: the cybernation revolution; the weaponry revolution; and the human rights revolution.

B. Information Technology and Automation Are a Central Reason Why Median Wages Have Been Stagnant in the US Over the Past Decade, Despite Rising Productivity.

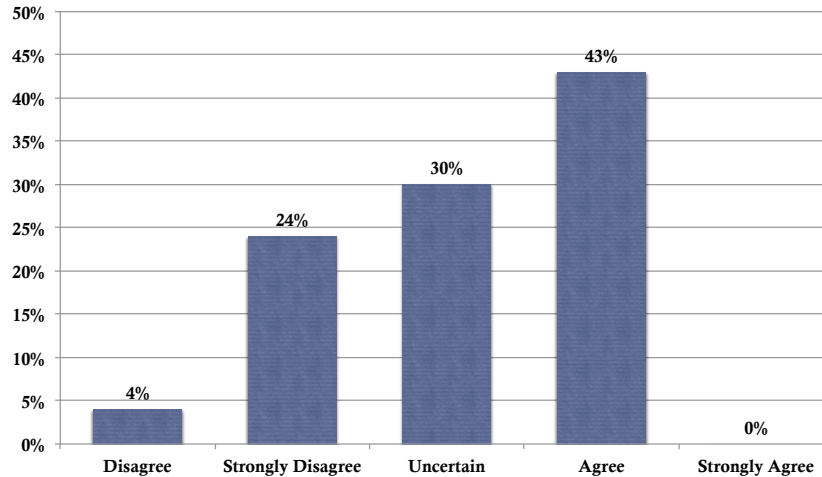


Figure 1 notes. Survey date February 25, 2014, available at http://www.igmchicago.org/igm-economic-experts-panel/poll-results?SurveyID=SV_eKbRnXZWx3jSRBb (accessed 3/26/2014)

In the five decades since the Ad Hoc Committee wrote its open letter, human labor has not been rendered redundant, as these scholars had feared. But automation anxiety has clearly returned. Perhaps most telling is the finding of a recent poll of leading academic economists conducted by the Chicago Initiative on Global Markets regarding the impact of technology on employment and earnings.⁶ Consistent with the canonical economic view that technology is, in the memorable phrase of Joel Mokyr, the “lever of riches,” a full 88 percent of economists in the poll either agreed or strongly agreed with the statement that “advancing automation has not historically reduced employment in the United States” (Figure 1). Yet, surprisingly, 43 percent of those polled endorsed (i.e., agreed with) the statement that “information technology and

⁶ The IGM webpage describes the panel members as follows: “Our panel was chosen to include distinguished experts with a keen interest in public policy from the major areas of economics, to be geographically diverse, and to include Democrats, Republicans, and Independents as well as older and younger scholars. The panel members are all senior faculty at the most elite research universities in the United States. The panel includes Nobel Laureates, John Bates Clark Medalists, fellows of the Econometric Society, past Presidents of both the American Economics Association and American Finance Association, past Democratic and Republican members of the President’s Council of Economics, and past and current editors of the leading journals in the profession.” Caveat emptor: The author is also a member of the panel.

automation are a central reason why median wages have been stagnant in the US over the past decade, despite rising productivity.” In contrast, only 28 percent disagreed or strongly disagreed. While I know of no comparable survey data from a decade earlier, I find these poll results stunning because they suggest that a plurality of mainstream economists has accepted—at least tentatively—the proposition that a decade of technological advancement has made the median worker no better off, and possibly worse off.

III. Employment Polarization: A Manifestation of Polanyi’s Paradox

To understand the wellspring of these concerns, it is useful to start from first principles. What do computers do? And how does their widespread adoption change what workers do?⁷ Anyone with children knows that computers appear “magical” to end users. But anyone who has written computer software knows that programming a computer to accomplish even the most rudimentary tasks is a tedious chore. Computers do not think for themselves, do not have common sense, do not compensate for programmer oversights and errors, and do not improvise solutions for unexpected cases. Fundamentally, computers follow procedures meticulously laid out by programmers. 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 precisely simulate these steps.

⁷ This essay’s singular focus on the impact of computerization on the labor market should not be taken to imply that computerization is the only important factor behind the employment and wage trends considered. Other important contributors include changes in the relative supply of college and non-college labor, rising trade penetration, offshoring and globalization of production chains, declines in labor union penetration, the falling ‘bite’ of the minimum wage, and shifts in tax policy. In addition, many of these forces work in tandem. Advances in information and communications technologies have directly changed job demands in US workplaces while simultaneously facilitating the globalization of production by making it increasingly feasible and cost-effective for firms to source, monitor, and coordinate complex production processes at disparate locations worldwide. The globalization of production has in turn increased competitive conditions for US manufacturers and US workers, eroding employment at unionized establishments and decreasing the capability of unions to negotiate favorable contracts, attract new members, and penetrate new establishments. This multi-dimensional complementarity among causal factors makes it both conceptually and empirically difficult to isolate the “pure” effect of any one factor. See Blinder (2009) and Blinder and Krueger (2013) for related theory and evidence on how the composition of job tasks—specifically, the demand for in-person interactions and physical proximity to customers—affect the potential for occupations to be offshored.

One early example of computer simulation was the use of punch card-driven computers at the Los Alamos National Laboratory to calculate the physical properties of explosions and implosions during the development of the first nuclear warheads.⁸ But the scope of computer simulation is not limited to simulating physical processes. 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 real 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. This remarkable cost decline creates strong economic incentives for firms to substitute ever-cheaper computing power for relatively expensive human labor, with attendant effects on employers' demand for employees. What are these effects?

The first-order effect is, of course, substitution. As the price of computing power has fallen, computers have increasingly displaced workers in accomplishing explicit, codifiable tasks—multiplication, for example. Autor, Levy and Murnane, 2003 (ALM hereafter) term these activities as "routine tasks," meaning tasks that follow an exhaustive set of rules and hence are readily amenable to computerization.⁹ Routine tasks are characteristic of many middle-skilled cognitive and manual activities, such as bookkeeping, clerical work, and repetitive production tasks. Because the core tasks of these occupations follow precise, well-understood procedures, they are increasingly codified in computer software and performed by machines. This force has

⁸ In many cases, the workers who performed these tasks were given the job title of "computer" (Grier, 2005). Prior to the Manhattan Project, an even earlier example of industrial-scale simulation was the use of mechanical "tabulators" to enumerate the 1890 Census of Population, which was stored on millions of punched cards.

⁹ Tasks such as performing a set of mathematical calculations, retrieving, sorting, and storing structured information, and precisely executing a repetitive physical operation in an unchanging environment, are routine in the sense of ALM (2003) not because they are mundane but because they can be fully codified and hence automated.

led to a substantial decline in employment in clerical, administrative support and, to a lesser degree, production and operative employment, as I document below.

But the scope for substitution is bounded: engineers cannot program a computer to simulate a process that they (or the scientific community at large) do not explicitly understand. This constraint is more binding than one might initially surmise because there are many tasks that we understand tacitly and accomplish effortlessly for which we do not know the explicit “rules” or procedures. I refer to this constraint as Polanyi’s paradox, following Michael Polanyi’s (1966) observation that, “We know more than we can tell.” When we break an egg over the edge of a mixing bowl, identify a distinct species of birds based only 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.¹⁰ 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.

At a practical level, Polanyi’s paradox means that many familiar tasks, ranging from the quotidian to the sublime, cannot currently be computerized because we don’t know “the rules.” At an economic level, Polanyi’s paradox means something more. The fact that a task cannot be computerized does *not* imply that computerization has no effect on that task. On the contrary: tasks that cannot be substituted by computerization are generally complemented by it. This

¹⁰ Computer scientists often refers to this same phenomenon as Moravec’s paradox, which is defined by Wikipedia as: “the discovery by artificial intelligence and robotics researchers that, contrary to traditional assumptions, high-level reasoning requires very little computation, but low-level sensorimotor skills require enormous computational resources. The principle was articulated by Hans Moravec, Rodney Brooks, Marvin Minsky and others in the 1980s. As Moravec writes, ‘it 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.’” (http://en.wikipedia.org/wiki/Moravec's_paradox, accessed 8/10/2014). I prefer the term Polanyi’s paradox to Moravec’s paradox because Polanyi’s observation also explains *why* high-level reasoning is straightforward to computerize and sensorimotor skills are not. High-level reasoning uses a set of formal logical tools that were developed specifically to address formal problems (e.g., counting, mathematics, logical deduction, encoding quantitative relationships). Sensorimotor skills, physical flexibility, common sense, judgment, intuition, creativity, spoken language, etc., are ‘built-in’ capabilities that the species evolved rather than developed. Formalizing these skills requires reverse engineering a set of activities that we normally accomplish using only tacit understanding. See Hoffman and Furcht (2014) for a discussion of the challenge that Polanyi’s paradox poses for scientific innovation.

point is as fundamental as it is overlooked. 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.¹¹

Concretely, consider the role played by mechanization in construction. By historical standards, contemporary construction workers are akin to cyborgs. Augmented by cranes, excavators, arc welders, and pneumatic nail guns, the quantity of physical work that a skilled construction worker can accomplish in an eight-hour workday is staggering. Naturally, automation has heavily substituted for human labor in performing construction tasks and, consequently, many fewer construction workers are required today to accomplish a given construction task than fifty years ago.

But construction workers have not been devalued by this substitution. Despite the array of capital equipment available, a construction site without construction workers produces nothing. Construction workers supply tasks such as control, guidance and judgment that have no current machine substitutes and which therefore become more valuable as machinery augments their reach. A worker wielding a single shovel can do a fairly limited amount of good or harm in an eight-hour day. A worker operating a front-end loader can accomplish far more. To a first approximation, automation has therefore complemented construction workers—and it has done so in part by *substituting* for a subset of their job tasks.

This example should not be taken to imply, however, that technological change is necessarily Pareto improving, even for construction workers. There are three factors that mitigate or augment its impacts:

¹¹ The extreme manifestation of this idea is the O-ring production function, discussed by Kremer (1993). In the O-ring production function, failure of any one step in the chain of production leads the entire production process to fail. Thus, 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, however, then the value of making link n more reliable as well rises.

1. Workers benefit 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 knows how to operate a shovel but not an excavator will generally experience falling wages as automation advances.
2. The elasticity of final demand can either dampen or amplify the gains from automation. Conceivably, productivity growth in construction could outstrip demand so that the value of further construction would fall even faster than output rose.¹² But this hypothetical response cannot capture the general case. Because household consumption has at least kept pace with household incomes over the very long run, we know that most technological improvements have ultimately translated into increased consumption rather than greater savings.
3. Labor supply changes can also mitigate wage gains. If the complementary tasks that construction workers supply are abundantly available elsewhere in the economy, it is plausible that a flood of new construction workers will temper wage gains emanating from complementarities between automation and human labor input.¹³

The construction example, writ large, explains a critical consequence of computerization that is typically overlooked in discussions of machine-worker substitution. Because machines both substitute for and complement human labor, focusing only on what is lost misses the central economic mechanism through which productivity growth raises the value of the tasks that workers uniquely supply.

I now return to Polanyi's paradox as it applies to computerization, focusing separately on two different margins of adjustment: changes in the occupational distribution (AKA job polarization) and changes in the wage distribution. I will argue that these occupational and

¹² Arguably, this has occurred with agricultural products over the long run: spectacular productivity improvements have met with declines in the share of household income spent on food.

¹³ While it is unlikely that supply effects would fully offset productivity-driven wage gains, there are perverse examples. Hsieh and Moretti (2003) find that new entry into the real estate broker occupation in response to rising house prices fully offsets average wage gains that would otherwise occur.

wage effects are likely to be distinct from one another for reasons hinted at in the discussion of construction labor above.

If computers largely substitute for routine tasks, how do we characterize the non-routine tasks for which they do not substitute? Autor, Levy and Murnane (2003) draw a distinction between two broad sets of tasks that have proven stubbornly challenging to computerize. One set includes tasks that require problem-solving capabilities, intuition, creativity and persuasion. These tasks, which ALM 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 of non-routine tasks that ALM identify are those requiring situational adaptability, visual and language recognition, and in-person interactions— which ALM term 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 are not highly skilled activities by human labor standards, they currently present daunting challenges for automation. Equally noteworthy is that many of the outputs of these jobs (haircuts, fresh meals, housecleaning) must be produced and performed on-site or in person (at least for now), and hence these tasks are not currently subject to outsourcing. Yet, because these jobs generally do not require formal education or extensive training beyond a high school degree, the potential supply of workers who can perform these jobs is very large.

Since 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—a straightforward implication of this reasoning is that computerization of routine job tasks may lead to the simultaneous growth of high-education, high-wage and low-education, low-wages jobs at the

expense of middle-wage, middle education jobs—a phenomenon that Maarten Goos and Alan Manning termed as ‘job polarization’ in a 2003 working paper. A large body of US and international evidence confirms the pervasive presence of employment polarization: computerization is strongly associated with employment polarization at the level of industries, localities, and national labor markets (Autor, Katz, and Kearney, 2006 and 2008; Goos and Manning, 2007; Autor and Dorn, 2013a; Michaels, Natraj, and Van Reenen, 2014; Goos, Manning, and Salomons, 2014).¹⁴


Figure 2 illustrates this pattern for the United States. The Figure plots percentage point changes in employment by decade for the years 1979 - 2012 for 10 major occupational groups encompassing all of U.S. non-agricultural employment.¹⁵ These ten occupations divide roughly into three groups. On the right-hand side of the figure are managerial, professional and technical occupations, which are highly educated, and highly paid. Between one-quarter and two-thirds of workers in these occupations had at least a four-year college degree in 1979, with

¹⁴ There are also two papers that do not find occupational polarization in the U.S. One is Katz and Margo (2013), who employ the extremely coarse occupation 1950 scheme provided by the IPUMS (Ruggles et al., 2010). The 1950 scheme enables long-term historical comparisons (Katz and Margo’s focus) at the expense of precision and is otherwise not normally applied to data from 1980 forward. The second paper is Mishel, Shierholz and Schmitt (2013), who offer an extended, and for the most part extremely careful, critique of the literature on technological change, employment and wage inequality. Their paper argues at length that the growth of low-wage service employment U.S. does not commence 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). While I concur with Mishel et al. that decadal revisions to the Census and OES occupational coding schemes (affecting the U.S. Census, Current Population Survey and Occupational Employment Statistics) make it essentially impossible to obtain fully consistent occupational employment counts over multiple decades, the adjustments that Mishel et al. apply to the data generate occupational patterns that appear anomalous. As later below in panel Figure 8 below (panel B), the data admit no ambiguity about the steep reallocation of non-college (high school or lower) workers from routine-intensive occupations (primarily production, operative, and clerical work) to manual task-intensive service occupations; indeed, the left-hand tail of the *employment* polarization phenomenon is driven by the changing occupational allocation of non-college workers. Mishel et al. acknowledge this reallocation (see Tables 6 and 8 of their paper) but view it as unimportant and argue that it is primarily explained by the secular decline in the fraction of non-college workers in the labor force (though it is unclear why a *decline* in the non-college share of the labor force would *increase* the share of non-college workers employed in service occupations).

¹⁵ More precisely, the figure plots 100 times log changes in employment, which are close to equivalent to percentage points for small changes.

the lowest college share in technical occupations and the highest in professional occupations (Acemoglu and Autor, 2011). Employment growth in these occupations was robust throughout the three decades plotted. Even in the deep recession and incomplete recovery between 2007 and 2012, these occupations experienced almost no absolute decline in employment.

Figure 2. Percentage Changes in Employment by Major Occupational Category, 1979 – 2012



Notes. 1980, 1990 and 2000 Census IPUMS files, American Community Survey combined file 2006 – 2008, and American Community Survey 2012. Sample includes the working-age (16-64) civilian non-institutionalized population. Employment is measured as full-time equivalent workers.

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 first two of this group of four are middle-skilled, white-collar occupations that are disproportionately held by women with a high school degree or some college. The latter two categories are a mixture of middle and low-skilled blue-collar occupations that are disproportionately held by males with a high school degree or lower education. While the headcount in these occupations rose in almost every decadal interval between 1979 and 2007, their growth rate lagged the economy-wide average and, moreover,

generally slowed across decades. These occupations were hit particularly hard after 2007, with absolute declines in employment between 2007 and 2012 ranging from 5 to 15 percent.


The leftmost three columns of Figure 2 depict employment trends in service occupations, which are 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 occupations categories. Despite their low educational requirements and low pay, employment has grown relatively rapidly in service occupations over the past three decades. All three broad categories of service occupations, protective service, food preparation and cleaning services, and personal care, expanded by double digits in the both the 1990s and the pre-recession years of the past decade (1999 - 2007). Notably, even during the recessionary years of 2007 through 2012, employment growth in service occupations was modestly positive—more so, in fact, than the three high-skilled occupations that have also fared comparatively well (professional, managerial and technical occupations). As noted by Autor and Dorn (2013a), the employment share of service occupations was essentially flat between 1959 and 1979. Thus, their rapid growth since 1980 marks a sharp trend reversal.

Cumulatively, these two trends of rapid employment growth in both high and low-education jobs have 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. One can quantify the consistency of this trend by correlating the changes in occupational employment shares across these ten occupational categories across multiple decades. The correlation between changes in occupational shares between 1979-1989 and 1989-1999 was 0.64, and for the decades of 1989-1999 and 1999-2007, was 0.67. Remarkably, the correlation between occupational share changes during 1999-2007 and 2007-2012—that is, prior to and during the Great Recession—was 0.80.

The polarization of employment across occupations is not unique to the United States. Evidence of this fact is presented in Figure 3, which plots changes in the share of employment between

1993 and 2010 within three broad sets of occupations—low-, middle-, and high-wage—covering all non-agricultural employment in 16 European Union economies. In all countries, middle-wage occupations declined as a share of employment, high wage occupations increased as a share of employment, and low-wage occupations gained in size relative to middle-wage occupations over this 17-year period.¹⁶ The comparability of these occupational shifts across a large set of developed countries—the United States among them—makes it likely that a common set of forces contributes to these shared labor-market developments.¹⁷ Simultaneously, the substantial differences among countries apparent in the data underscores that no single factor or common cause explains the diversity of experiences across the United States and the European Union.

Figure 3. Change in Occupational Employment Shares in Low, Middle and High Wage Occupations in 16 EU Countries, 1993 - 2010



¹⁶ In 14 of 16 countries, low-wage occupations increased as a share of employment.

¹⁷ Not only is the U.S. not unique in this regard, it is not even an outlier—falling roughly in the middle of the pack of this set of countries in the extent of employment polarization.

Figure 3 notes. Source: Goos, Manning and Salomons (2014, Table 2). 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.

IV. 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. This reasoning is appealing but incomplete because it ignores the role played by the three mitigating forces discussed above: complementarity; demand elasticity, and labor supply.¹⁸

Let's first consider the impact of computerization on wages in abstract task-intensive occupations such as managerial, professional and technical occupations. A key attribute of these occupations is that all draw upon large bodies of constantly evolving expertise, e.g., medical reports, legal cases, sales data, financial analysis, and economic statistics—so much so that many abstract task-intensive occupations employ skilled assistants and paraprofessionals to support their information processing tasks (e.g., medical secretaries, paralegals, and research assistants). Given this production structure, a clear prediction is that 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

¹⁸ Firpo, Fortin and Lemieux (2013) provide an analysis of the link between job tasks and wage polarization in the United States.

comparative advantage, with less time spent on acquiring and crunching information, and more time spent on interpreting and applying it.¹⁹


If demand for the output of abstract task-intensive activities were inelastic, however, these productivity gains might work to lower expenditure on these outputs, which could in turn mitigate wage gains. While it is hard to develop a strong theoretical prior on this possibility, all outward evidence suggests the opposite. As the output of the professions has risen, demand for their services has seemingly more than kept pace. A leading example is medicine, where expenditures for medical services have risen substantially as a share of GDP as the efficacy of medicine to address a larger set of ailments has expanded. But one can readily make similar arguments about finance, law, engineering, research, and design.²⁰

What about the labor supply? If workers could quickly move into the highly educated professions to capitalize on rising productivity, this would mute earnings gains. But of course, many professions require both college and graduate degrees (MBAs, JDs, MDs, Ph.D.s), meaning that the production pipeline for new entrants is five to ten years in length and, hence, supply almost necessarily responds slowly. Indeed, as discussed in Autor (2014b), young U.S. adults have responded remarkably sluggishly to the rising educational premium over the last thirty years — and this is particularly true for males, as shown in Figure 4. Thus, while the stock of workers with college and graduate degrees has certainly grown in response to rising productivity in these occupations, the supply response has not been nearly large enough to swamp the contemporaneous movements in demand.

¹⁹ By the same token, computerization substitutes for many of the support occupations that these professions employ.

²⁰ There are counterexamples as well. For example, computerization 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.

Figure 4. College Share of Hours Worked in the U.S. 1963-2012:
Workers with Less than 10 Years of Potential Experience



Notes. Autor (2014b, Figure S2) based upon March Current Population Survey data for earnings years 1963 – 2012.

Workers in abstract task-intensive occupations have therefore benefit from computerization via a virtuous combination of three forces: 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 computerization should raise earnings in occupations that make intensive use of abstract tasks and among workers who intensively supply them.

Do these same synergies apply to jobs that are intensive in manual tasks, such as janitors and cleaners, vehicle drivers, flight attendants, food service workers, and personal care assistants? In large part, the answer appears to be no. In contrast to workers in abstract task-intensive occupations, computerization has not greatly increased the reach or productivity of housekeepers, security guards, waiters, cooks, or home health aids. Because most manual task-intensive occupations are minimally reliant on information or data processing for their core tasks, there are very limited opportunities for either direct complementarity or substitution.

There are of course exceptions to this generalization: GPS and scheduling software allows truckers to minimize wasted mileage; calendar and contact software assists home health workers to more effectively manage time and bill hours; computerized ordering systems permit food service workers to rapidly tally customer tabs. But these information-intensive tasks are largely peripheral to these occupations' core job tasks.²¹ Ironically, manual task-intensive occupations enjoy relatively minimal direct benefits from computerization because they are *too well* insulated, offering limited opportunities for substituting or complementing human labor with information technology.

Say for the sake of argument, however, that demand for manual task-intensive occupations was rising due to rising societal income or changes in preferences. Would these demand increases likely translate into higher occupational earnings? The answer turns on both the elasticity of final demand and the elasticity of labor supply, as noted above. Much aggregate evidence suggests that final demand for manual task-intensive work—services in particular—is relatively *price* inelastic (Baumol, 1967; Autor and Dorn, 2013a). If so, productivity gains in manual task-intensive occupations 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), meaning that rising aggregate incomes will tend to increase demand for these activities. Computerization may therefore *indirectly* raise demand for manual task-intensive occupations by increasing societal income.²²

²¹ While surely there are important exceptions—for example, checkout cashiers using scanner driven point of sale registers that interface directly to centralized inventory systems—I see relatively few examples where information technology has fundamentally transformed the work activities or productivity of workers in manual task-intensive occupations. (I thank Dave Wessel for the cashier example.)

²² This can result from either of two economic forces. Clark (1951) argues that demand for services is non-homothetic: the expenditure share of services rising with income. Baumol (1967) argues that growing expenditure on services reflects unbalanced growth: because the relative prices of technologically lagging activities (e.g., haircuts and symphony orchestra performances) necessarily rise over time, an increasing share of societal income must be expended on these activities to maintain balanced consumption. Of course, Baumol's argument presupposes that demand for these activities is relatively inelastic—otherwise expenditure would fall as relative prices rose. Mazzorali and Ragusa (2013) present evidence consistent with Clark's view while Autor and Dorn (2013a) present evidence consistent with Baumol's thesis.

This is where the elasticity of labor supply becomes most critical, however. Due to their generally low education and training requirements, labor supply to manual task-intensive occupations is intrinsically elastic.²³ Consequently, wage increases in manual task-intensive occupations generally spur a robust supply response. Moreover, workers displaced from other sectors of the economy may readily obtain employment in manual task-intensive occupations due to their low entry requirements.


In short, while abstract task-intensive activities benefit from strong complementarities with computerization, relatively elastic final demand, and a low elasticity of labor supply, 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 computerization has strongly contributed to *employment* polarization, we would not generally expect these employment changes to culminate in *wage* polarization except in tight labor markets (Autor and Dorn, 2013a).²⁴

Figure 5 presents evidence consistent with this logic. Following the format of Figure 2, this figure depicts percentage point changes in mean weekly wages by occupation among full-time, full-year workers for 1979 through 2012, subdividing the time interval into the 1980s, 1990s, and the pre- and post-recession 2000s. To provide additional detail, Figure 6 plots wage changes across all occupational categories, weighted by initial size and smoothed for clarity. Specifically, the figure ranks all 318 detailed occupations from lowest to highest by their initial skill level (as measured by its 1979 mean hourly occupational wage), groups these detailed occupations into 100 bins of equal sizes, and plots smoothed changes in log earnings at each occupational percentile over each sub-period.

²³ Baumol (1967) observes that 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).

²⁴ Autor and Dorn (2013a) present evidence that the consumption complementarity effect (due to rising incomes) dominated the displacement effect on net between 1980 and 2005. But this effect was primarily driven by wage developments in the 1990s when labor markets were extremely tight. After 2000, the expansion of manual task-intensive service occupations accelerated but wages in these occupations fell.

Figure 5. Changes in Mean Wages by Major Occupational Category among Full-Time, Full-Year Workers, 1979 – 2012




Notes. Calculated using 1980, 1990 and 2000 Census IPUMS files; American Community Survey combined file 2006 – 2008, American Community Survey 2012. Sample includes the working-age (16-64) 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.

The right-hand two thirds of these wage figures look much like the plots of employment polarization. From 1979 through 2007, wages rose consistently across all three abstract task-intensive categories of professional, technical and managerial occupations.²⁵ By contrast, wage growth in the four middle-skill, routine task-intensive occupations was less rapid than in abstract task-intensive occupations and generally decelerated over time—with particularly anemic (and in two of four categories, negative) growth after 2000.

²⁵ Noting that these occupations are also the highest paid initially, their greater proportional earnings growth translates into even larger dollar growth.

Figure 6. Changes in Mean Wages by Occupational Skill Percentile among Full-Time, Full-Year Workers, 1979 – 2012



Notes. Calculated using 1980, 1990 and 2000 Census IPUMS files; American Community Survey combined file 2006 – 2008, American Community Survey 2012. The figure plots changes in mean log wages 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. Sample includes the working-age (16-64) 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.

The low-education, manual task-intensive occupations on the left-hand side of Figures 5 and 6 present a particularly intriguing pattern. Wage growth in these occupations was somewhat *more* rapid than in the routine task-intensive occupations in the 1980s—and even more so in the 1990s—which was roughly concordant with the pattern of employment polarization taking shape simultaneously. However, in the 2000s, employment and wage trends in manual task-intensive occupations diverged. While employment growth in these occupations exceeded that in all other categories between 2000 and 2007 (Figure 2), wage growth was generally negative—more so than almost all other categories (Mishel, Shierholz and Schmitt, 2013).

Why did wage growth in manual task-intensive occupations go from positive to negative after 1999? My strong hunch is that the explanation is shifting labor supply. Recent papers by Christopher Smith (2013), Cortes et al. (2014), and Foote and Ryan (2014) find that declining employment in routine task-intensive jobs has led middle-skill workers—both new entrants and those displaced from routine task-intensive jobs—to enter manual task-intensive occupations instead. This likely occurred particularly rapidly in the 2000s as flagging employment in middle-skill occupations combined with slack macroeconomic conditions spurred middle-skill workers to compete with less-educated workers for manual task-intensive jobs, thus checking the tendency for wages to rise in these occupations.

A final set of facts starkly illustrated by Figure 6 is that overall wage growth was extraordinarily 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.²⁶ Of course, wage growth was essentially zero at all percentiles during the recessionary years of 2007-2012.²⁷

V. Polarization: What Have We Learned from Another Decade of Data?

Although the polarization hypothesis can explain some key features of the U.S. and cross-national data, reality invariably proves more complicated than the theory anticipates. The clearest evidence for this general dictum is the unexplained deceleration of employment growth in abstract task-intensive occupations after 2000, which is discussed by Beaudry, Green and

²⁶ Since the 2000-2007 interval is two years shorter than the 1979-1989, one should multiple the 2000-2007 changes by 1.25 to put them on the same temporal footing. Net of this adjustment, wage growth is still considerably weaker at all percentiles than in the earlier two decades.

²⁷ Why are the rapidly rising earnings of the top 1 percent (e.g., Atkinson et. al. 2011) not strongly evident in Figure 6? There are two reasons, one reflecting substance, and the other data. Substantively, the plot depicts changes in earnings by *occupational* percentile rather *wage* percentile. Wage growth by occupational percentile is less concentrated than wage growth across wage percentiles since 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.

Sand (2013, 2014) and Mishel, Shierholz and Schmitt (2013).²⁸ This can be seen especially clearly in Figure 7, which, following the format of Figure 6, plots smoothed changes in the share of U.S. employment (rather than wages, as in Figure 6) at each occupational percentile. Since the sum of shares must equal one in each decade, the change in these shares across decades must total zero and, thus, the height at each skill percentile measures the growth in each occupation's employment relative to the whole.²⁹


Figure 7 contributes three nuances to the occupational polarization story above. A first, visible on the left-hand side of the figure, is that the pace of employment gains in low-wage, manual task-intensive jobs has risen successively across periods. Gains in these occupations were barely discernible in the 1980s, intensified in the 1990s, and accelerated again in the 2000s. A second nuance is that the occupations that are losing 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 1990s, the crossover point lay at approximately the 55th percentile. In the final two sub-periods, this rank rose still further to above the 75th percentile—suggesting that the locus of displacement of middle-skill employment is moving into higher skilled territories. The final empirical regularity highlighted by Figure 7 is that growth of high-skill, high-wage

²⁸ As discussed in footnote 14 above, Mishel, Shierholz and Schmitt (2013) dispute the factual basis and economic relevance of almost all empirical and theoretical conclusions of the polarization literature. Substantively, however, I believe their main contention is not that *employment polarization* has not occurred but rather that it has not contributed to *wage polarization*—or, more broadly, that occupational employment patterns are uninformative or irrelevant to the evolution of wage inequality. As discussed in section IV and noted in Autor and Dorn (2013), occupational polarization does *not* necessarily generate wage polarization since there are two countervailing forces operative: labor demand shifts stemming from consumption complementarities between goods and services (which tend to raise wages in service occupations over the long run); and labor supply shifts, stemming in part from movement of low-education workers out of middle-skill, routine-intensive occupations and into traditionally low-skill, manual-intensive occupations (which place downward pressure on wages in service occupations). Autor and Dorn (2013) provide evidence from local labor markets that occupational polarization contributed to wage polarization during the period 1980 through 2005. But this relationship is clearly not immutable since wage polarization reversed course after 2000 whereas employment growth in low-wage service occupation employment accelerated. While the evidence on *employment polarization* appears to me unambiguous, I leave it to the reader to assess whether these occupational employment shifts are helpful for understanding *wage polarization* or wage inequality more broadly.

²⁹ Due to the smoothing of the plotted series, this adding up property holds only as an approximation.

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 U-shaped growth of occupational employment came increasingly to resemble a downward ramp in the 2000s. Figure 8 takes a closer look at this phenomenon by plotting the distribution of occupational employment changes among college-educated (panel A) and non-college (panel B) workers across the three broad occupational categories above: manual-intensive, routine-intensive, and abstract-intensive.³⁰

Figure 7. Smoothed Employment Changes by Occupational Skill Percentile, 1979 – 2012




³⁰ Specifically, I collapse the ten categories in Figure 2 into three broader groupings. Manual occupations are personal care, food/cleaning service, and protective services. Routine occupations are operators/laborers, production, office/administrative, and sales. Abstract occupations are technicians, professionals, and managers.

Notes for Figure 7. Calculated using 1980, 1990 and 2000 Census IPUMS files; American Community Survey combined file 2006 – 2008, American Community Survey 2012. 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).

One would anticipate that a long-term rise in the college-educated workforce would eventually lead to a growth in college-educated workers in non-traditional occupations. This pattern is seen in the 1980s: the fraction of college-educated workers in both manual and routine task-intensive occupations rose modestly in this decade while the share in abstract-intensive occupations declines. In the subsequent decade of the 1990s, employment of college-educated workers polarized, with a sharp reduction in routine task-intensive employment, a steep rise in abstract task-intensive employment, and a modest rise in manual task-intensive employment. After 2000, however, occupational employment patterns of college-educated workers turned sharply downward, as discussed by Beaudry, Green and Sand (2013, 2014). Between 1999 and 2012, the fraction of college-educated workers employed in abstract occupations *fell* by more than a percentage point, the share employed in routine occupations fell by 0.4 percentage points, and the share employed in manual occupations rose by 1.5 percentage points. Among non-college workers, however, we see a much more consistent pattern of sharp reductions in routine employment and equally large gains in manual employment with essentially no gains in abstract employment except during the first decade of the sample (panel B). In net, these patterns suggest that the set of abstract task-intensive jobs is not growing as rapidly as the potential supply of highly educated workers. As Beaudry, Green and Sand (2013, 2014) highlight, the coalescence of these forces has likely led highly educated workers to seek less educated jobs, which in turn creates still greater challenges for the lower educated workers competing for routine and manual task-intensive work.

Figure 8. Changes in Employment Shares in Broad Occupational Categories, 1979 – 2012:
Workers with and Without a Four-Year College Degree



Notes. Calculated using 1980, 1990 and 2000 Census IPUMS files; American Community Survey combined file 2006 – 2008, American Community Survey 2012. Manual occupations are personal care, food/cleaning service, and protective services. Routine occupations are operators/laborers, production, office/administrative, and sales. Abstract occupations are technicians, professionals, and managers.

What explains the slowing growth of abstract task-intensive employment? One possible interpretation is that technological progress has encroached strongly upward in the task domain, such that it now strongly substitutes for the work done by professional, technical and managerial occupations. While one should not dismiss this possibility out of hand, contemporaneous data on computer and software investment militates against this interpretation. One would expect that a surge of new automation opportunities in highly paid work would catalyze a surge of corporate investment in computer hardware and software. Instead, the opposite occurred, as shown in Figure 9. After rising near-monotonically from approximately one-half of one-percent to almost five percent of gross domestic product between 1950 and 2000, private fixed investment in information processing equipment and software dropped by a full percentage point (more than 20 percent) between 2000 and 2002, and remained depressed thereafter. As of the first quarter of 2014, information processing equipment and software investment as a share of GDP was (only) at 3.5 percent, a level last seen in 1995 at the outset of the “dot-com” era.

Given that both employment growth in abstract task-intensive occupations and computer investment tailed off simultaneously, is it possible that one caused the other? That is roughly the view espoused by Beaudry, Green and Sand (2013). Building on Welch (1970) and Schultz (1975), they posit a model in which the introduction of a new technology generates demand for managerial and problem-solving skills (abstract tasks) during an *adoption* period, while the technology is installed, adapted, mastered and routinized, after which skill demands slacken since the challenge of adaptation gives way to the more quotidian tasks of operation and maintenance. The Beaudry, Green and Sand conceptual model can therefore rationalize both decelerating employment in abstract task-intensive work and slowing investment in information technology, though unfortunately the model generates few testable predictions beyond these aggregate facts.

An alternative interpretation of these facts, offered by Gordon (2012), is not that the computer revolution has been fully realized but rather that it has petered out. Gordon argues that the gains from information technology have been relatively superficial and short-lived, with few of


the monumental consequences for productivity or human welfare afforded by prior eighteenth and nineteenth century technological revolutions—transportation, power generation, communications, and sanitation. Thus, the slowdown in computer investment reflects the onset of rapidly diminishing marginal returns to information technology and an accompanying deceleration of productivity growth.³¹

While I do not have a strongly evidenced counter-explanation for these same facts, I am skeptical of both interpretations. Each would seem to imply that IT investment would plateau as the technology either attained maturity or reached a point of diminishing returns. But this does not accord with the evidence in Figure 9. Information technology investment abruptly reversed course after 2000, suggesting a rapid pullback in demand. Moreover, in the five years *prior to* this pullback, IT investment surged at an unprecedented rate, rising by approximately two-tenths of a percent of GDP in each year. What this pattern suggests to me is a temporary dislocation of demand for IT capital during the latter half of the 1990s followed by a sharp correction after 2000—in other words, the bursting of a bubble. The end of the “tech bubble” in the year 2000 is of course widely recognized, as the NASDAQ stock index erased three-quarters of its value between 2000 and 2003. Less appreciated, I believe, are the economic consequences beyond the technology sector: a huge falloff in IT investment, which may plausibly have dampened innovative activity and demand for high skilled workers more broadly.

It is also possible to read a more optimistic message from the trends in Figures 7 and 8. Employment in highly skilled occupations appears to have shown renewed growth after 2007, and this may augur renewed investment and innovative activity. Simultaneously, the ongoing contraction of middle-wage occupations and rapid expansion of lower-wage occupations provides less cause for comfort.

³¹ Gordon’s thesis does not address the labor market implications of the anticlimactic conclusion of the information technology revolution. Thus, there is no implied relationship between the slowdown of IT investment and the deceleration of employment growth in abstract task-intensive jobs.

Figure 9. 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)

VI. *The Future of Polanyi's Paradox*

Even as investment in information technology has slowed, popular and academic discussion of the potentially dire consequences of automation for employment has accelerated.³² MIT scholars Erik Brynjolfsson and Andrew McAfee argue in a 2011 book that workers are in danger of losing the “race against the machine.” In a 2012 working paper, Sachs and Kotlikoff present a model in which “smart machines” yield an economy of “long-term misery” because workers whose labor is devalued by automation are unable to make the human capital investments that would enable their children to profit from advancing technology. In a popular vein, journalist

³² This paragraph draws on Autor and Dorn (2013b).

Kevin Drum (2013) warns in *Mother Jones* that our “robot overlords” will soon take our jobs.³³ But where are these robot overlords? And if they are not here already—and all outward appearances suggest that they are not—should we expect their imminent arrival? In this final section, I discuss the progress of computing towards overcoming Polanyi’s paradox.

In the past decade, computerization has 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. Yet, Polanyi’s paradox remains relevant. Indeed, it helps to explain what has *not* yet been accomplished and to illuminate the current technological approaches used to enlarge the set of machine-feasible tasks. In my reading of the technology landscape, there are two overarching approaches that engineers employ to computerize tasks for which we “do not know the rules.” One approach, which I call environmental control, bows to Polanyi’s paradox. The second, machine learning, attempts to make an end-run around it.

A. 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. Why not robots? Because removing a broken windshield, preparing the windshield frame to accept a replacement, and fitting a replacement into that frame demand far more real-time adaptability than any contemporary robot can approach.

The distinction between assembly line production and the in-situ repair highlights the role of environmental control in enabling automation. While machines cannot generally operate autonomously in unpredictable environments, engineers can in some cases radically simplify the environment in which machines work to enable autonomous operation. The factory assembly line provides one such example of environmental control. But there are numerous examples that are so ingrained in daily technology that they escape notice. Present day

³³ In a more reflective vein, Noah Smith (2013) considers the challenges for income distribution were *if* advancing in robotics were to substantially devalue the stock of human capital.

automobiles, for example, are highly evolved machines—efficient, powerful, safe and reliable. But in another sense, they are remarkably helpless: they can only operate on smooth paved surfaces, something that is almost never found in nature. To enable their operation, humanity has adapted the naturally occurring environment by leveling, re-grading and covering with asphalt a non-trivial percentage of the earth’s land surface.³⁴ Environmental control for automobile travel has meant remaking the natural landscape to machine age specifications.

Executing non-routine tasks is a central obstacle in computer-based automation. Thus, environmental control in computer applications often means eliminating non-routine work tasks. One can see this process clearly in industrial robotics. Large online retailers, such as Amazon.com, Zappos.com and Staples.com, operate systems of warehouses that stock, pack and ship thousands of varieties of non-homogenous goods directly to consumers and businesses. These warehouses employ 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 to purchasers. There is at present no technologically viable or cost-effective robotic 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 from stocking, to storage, to picking and shipping. The dispatch software directs both a fleet of Kiva robots—essentially motorized, remotely controlled go-karts—and a set of *human* stockers and pickers who work in tandem with the robots. The Kiva robots circulate through a warehouse that is filled with uniform racks of freestanding stocked shelves. The robots’ sole task is to transport shelves from one location to another, which they accomplish by maneuvering under a rack of shelves, raising slightly to elevate the rack from the floor, motoring to a new location, and then lowering the rack.

³⁴ 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 Ohio (http://en.wikipedia.org/wiki/Impervious_surface, accessed 8/4/2014).

As objects arrive at the warehouse for stocking, the dispatch software directs robots to transport empty shelves to the loading area where they line up for loading. The software simultaneously directs *human* stockers to place merchandise on awaiting shelves at precise locations. Once stocked, shelves are sent back into the warehouse on their robotic carriers, where the dispatch software directs their dynamic placement to optimize product availability for expected product demand. As new orders arrive, the dispatch software sends robots to retrieve shelves containing needed items. The shelves line up in the packing area where they await a human picker who, directed by a laser pointer controlled by the dispatch software, picks 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 shelves scurry back to the warehouse floor (perhaps dynamically relocated) until needed again for packing or restocking.

Human flexibility is still required in the Kiva-operated warehouse: only workers handle merchandise; robots only move shelves. But the demand for human dexterity is dramatically reduced by automation: all non-routine motor tasks are performed during stocking and packing; all other goods movement, organization, storage and retrieval is delegated to robots, whose sole task is to shuttle shelves across a level surface (a routine task). Thus, Kiva applies environmental control to minimize the need for human flexibility.

While Kiva Systems provides a particularly clear example, the same principle of environmental control is often operative in unexpected places. Perhaps the least recognized—and most mythologized—is the Google Car. It is sometimes said by computer scientists that the Google car does not drive on roads but rather on maps. This observation conveys the fact that the Google car, unlike a human vehicle operator, cannot pilot on an “unfamiliar” road; it lacks the capability to process, interpret and respond to an environment that has not been pre-processed by its human engineers. Instead, the Google car navigates through the road network primarily by comparing its real-time audio-visual sensor data (collected using LIDAR) against painstakingly hand-curated maps that specify the exact locations of all roads, signals, signage, obstacles, etc. The Google car adapts in real time to obstacles (cars, pedestrians, 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 key static features of its pre-specified map (e.g., an unexpected detour, a police officer directing traffic where a traffic signal is supposed to be), then the car signals for its human operator to take command. Thus, while the Google car appears outwardly to be as adaptive and flexible as a human driver, it is in reality more akin to a train running on invisible tracks.

These examples highlight some of the limitations of current technology to accomplish non-routine tasks. They also illustrate the genius of human ingenuity in surmounting these obstacles. Humans naturally tackle tasks in a manner that draws on their inherent flexibility, problem solving capability, and judgment. Machines currently lack many of these capabilities, but they possess other facilities in abundance: strength, speed, accuracy, low cost, and unwavering fealty to directions. Engineering machines to accomplish human tasks does not necessarily entail equipping machines with human capabilities; instead, work tasks can, in some cases, be reengineered so that the need for specifically human capabilities is minimized or eliminated.

B. Machine Learning

There is an alternative route, however. Polanyi's paradox—"we know more than we can tell"—presents a challenge for computerization because conventional programming amounts to "telling" a computer precisely how to accomplish a task. If people tacitly understand how to perform a task but cannot "tell" a computer how to perform the task, then seemingly programmers cannot automate the task—or so the thinking has gone historically. But this understanding is shifting due to advances in machine learning. The simple idea of machine learning is to applying statistics and inductive reasoning to supply best guess answers in cases where formal procedural rules are unknown. Where engineers are unable to program a machine to "simulate" a non-routine 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. Thus, 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 one concrete example of machine learning, consider the challenge of task of visually identifying a chair.³⁵ Applying the conventional rules-based programming paradigm, an engineer might attempt to specify *ex ante* what features of an object qualify it as a chair—it possesses legs, arms, a seat, and a back, for example. One could then program a machine to identify objects possessing these features as chairs. But having specified such a feature set, one would immediately discover that many chairs that do not possess all features (e.g., no back, no legs). If one then relaxed the required feature set accordingly (e.g., chair back optional), the included set would clearly encompass many objects that are not chairs (e.g., tables). Thus, the canonical routine task approach to object recognition—and many more sophisticated variants—would likely have very high misclassification rates. Yet, any grade-school child could perform this task with very high accuracy. What does the child know that the rules-based procedure does not? Unfortunately, we do not know—this is precisely Polanyi’s paradox.

Machine learning potentially circumvents this problem. Relying on large databases of so-called ground truth—concretely, a vast set of curated examples of labeled objects—a machine learning algorithm can attempt to statistically infer what attributes of an object make it more or less likely to be designated a chair. This process is called training. Once training is complete, the machine can then apply this statistical model out of sample to potentially 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 (e.g., different shapes, materials, or dimensions). What makes the idea of machine learning powerful is that it 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.³⁶

How well does machine learning work in practice? If you use Google Translate, operate a smartphone with voice commands, or follow Netflix’ movie suggestions, you can assess for

³⁵ This example draws on the discussion in Autor (2014a).

³⁶ Levy and Murnane (2004) provide numerous illustrative examples of the automation of job tasks. For introductory material on machine learning written by and for economists, see Varian (2014).

yourself how successfully these technologies function.³⁷ My general observation is that the tools are inconsistent: uncannily accurate at times; typically, only so-so; and occasionally, unfathomable.³⁸ 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 the course of its winning match. Under the category of *U.S. 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 exemplary 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 (see Figure 10 for examples). The article's headline ruefully poses the question, "How Many Computers to Identify a Cat? 16,000."³⁹

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. Still, the long-term potential of machine learning for circumventing Polanyi's paradox is a subject of active debate among computer scientists. 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.

³⁷ By logging and analyzing the clicks of users in response to earlier queries, search engines also use machine learning to dynamically refine search results offered for subsequent queries. 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.

³⁸ A lovely irony of machine learning algorithms is that they also cannot "tell" programmers why they do what they do. The "decisions" that a machine learning program makes following training are something of a black box.

³⁹ As further evidence of the inchoate stage of machine learning (at least as of 2012), notice in Figure 10 that the image third down from the top and third over from the left is definitely not a cat and appears more likely to be a pair of coffee cups.

Figure 10. Images Identified as Cats by Google X Labs Team
Using a Neural Network of 16,000 Processors



Source: British Broadcasting Company (June 26, 2012, <http://www.bbc.com/news/technology-18595351>, accessed 8/4/2014).

To give this skepticism heft, return to the challenge of training a machine to recognize a chair. Ultimately, what makes an object a chair is that it is a device purpose-built for a human being to sit upon. This “purposiveness” may be difficult for a machine learning algorithm to infer, even given an arbitrarily large training database of images. As Grabner et al. (2011) argue, it is likely that humans recognize chairs not simply by comparing candidate objects to statistically probable feature sets but also by reasoning about the attributes of the object to assess whether it is likely intended to serve as a chair. For example, both a toilet and a traffic cone look somewhat like a chair, but a bit of reasoning about their shapes vis-à-vis the human anatomy suggests that a traffic cone is unlikely to make a comfortable seat. Drawing this inference, however, requires reasoning about what an object is “for” not simply what it looks like. Contemporary object recognition programs do *not*, for the most part, take this reasoning-based approach to identifying objects, likely because the task of developing and generalizing the approach to a

large set of objects would be extremely challenging.⁴⁰ One is reminded of Carl Sagan’s remark that, “If you wish to make an apple pie from scratch, you must first invent the universe.”

VII. Conclusions

A principle conclusion from the discussion above is that the challenges to computerizing numerous everyday tasks—from the sublime to the mundane—remain substantial. Let us assume, however, that a set of near-term breakthroughs enables rapid technological progress in non-routine manual and abstract domains. What does this augur for labor demand?

As chronicled in section II, there is a long history of leading thinkers overestimating the potential of new technologies to substitute for human labor and underestimating their potential to complement it. The green revolution displaced labor from farming. The industrial revolution replaced skilled artisanal labor with unskilled factory labor. The mass-produced automobile drastically reduced demand for blacksmiths, stable hands, and many other equestrian occupations. Successive waves of earth moving equipment and powered tools displaced manual labor from construction. In each case, groups of workers lost employment and earnings as specific jobs and accompanying skill sets were rendered obsolete.

Yet, short-term employment losses sparked by rising productivity were eventually more than offset by subsequent employment gains—in some cases in the innovating sectors, in many cases elsewhere. In 1900, for example, 41 percent of the United States work force was employed in agriculture. By 2000, that share had fallen to 2 percent, in large part due to productivity gains emanating from the Green Revolution (Autor, 2014b). It is unlikely, however, that farmers at the turn of the twentieth century could foresee that one hundred years later, health care, finance, information technology, consumer electronics, hospitality, leisure and entertainment would employ far more workers than agriculture.

⁴⁰ Could, for example, a machine that recognizes chairs by reasoning about their potential compatibility with human anatomy also be readily reprogrammed to recognize bicycles—or would it require another set of reasoning capabilities to determine whether the object could support a human being in the act of balancing while in motion?

Arguably, we stand at a similar moment today. One can find fresh examples daily in which technology substitutes for human labor in an expanding—though still circumscribed—set of tasks. The complementarities are always harder to identify. Despite these uncertainties, there are three inferences in which we can be fairly confident:

A first is that the technological advances that have secularly pushed outward the demand for skilled labor over many decades will continue to do so. As physical labor has given way to cognitive labor, the labor market's demand for formal analytical skills, written communications, and specific technical knowledge has risen spectacularly. If the nineteenth U.S. labor force were suddenly restored in the twentieth century, a large fraction of workers would be surely unemployable due to their exceedingly low levels of education—averaging approximately nine years of completed schooling (Katz and Goldin, 2008). While some have speculated that the advent of labor market polarization—particularly the growth of low-education, manual task-intensive jobs—indicates that the complementarity between higher education and technological change has come to an end, this reasoning is incorrect. Though computerization may increase the fraction of jobs found in manual task-intensive work, it is generally unlikely to rapidly boost earnings in these occupations for the reasons discussed above: an absence of strong complementarities and an abundance of potential labor supply. Thus, human capital investment must be at the heart of any long-term strategy for producing skills that are complemented rather than substituted by technology.

A second observation is that employment polarization will *not* continue indefinitely.⁴¹ While many middle skill *tasks* are susceptible to automation, many middle skill *jobs* demand a mixture of tasks from across the skill spectrum. To take one prominent example, medical support occupations—radiology technicians, phlebotomists, nurse technicians, etc.—are a numerically significant and rapidly growing category of relatively well-remunerated, middle skill employment. While not all of these occupations require a college degree, they do at least demand two years of post-secondary vocational training. Significantly, mastery of “middle

⁴¹ This discussion draws on Autor (2013), with some passages quoted directly.

skill” mathematics, life sciences, and analytical reasoning is indispensable for success in this training.

Why are these middle skill jobs likely to persist and, potentially, to grow? 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 the residual---without a substantial drop in quality. Consider, for example, the commonplace frustration of calling a software firm for technical support only to discover that the support technician knows nothing more than what is on his or her computer screen—that is, the technician is a mouthpiece, not a problem solver. This example captures one feasible division of labor: machines performing routine technical tasks, such as looking up known issues in a support database, and workers performing the manual task of making polite conversation while reading aloud from a script. But this is not generally a productive form of work organization because it fails to harness the complementarities between technical and interpersonal skills. Stated in positive terms, routine and non-routine tasks will generally coexist within an occupation to the degree that they are complements---that is, the quality of the service improves when the worker combines technical expertise and human flexibility.⁴²

This reasoning suggests that many of the middle skill jobs that persist in the future will combine routine technical tasks with the set of non-routine tasks in which workers hold comparative advantage---interpersonal interaction, flexibility, adaptability and problem-solving.⁴³ Medical support occupations are one leading example of this virtuous combination, but this example is not a singularity. This broad description also fits numerous skilled trade and repair occupations—plumbers, builders, electricians, HVAC installers, automotive technicians—marketing occupations, and even modern clerical occupations that provide coordination and decision-making functions rather than simply typing and filing. Indeed, even as some formerly middle skill occupations are stripped of their routine technical tasks and arguably deskilled—

⁴² Lawrence Katz memorably titles workers who virtuously combine technical and interpersonal tasks as “the new artisans” (see Friedman, 2010).

⁴³ 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.

for example the stockbroking occupation—other formerly high-end technical occupations are made accessible to workers with less esoteric technical mastery, 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, non-college jobs combining specific vocational skills with foundational middle skills—literacy, numeracy, adaptability, problem-solving and common sense—will persist in coming decades.

A final observation is that while much contemporary economic pessimism attributes the labor market woes of the past decade to the adverse impacts of computerization, I remain skeptical of this inference. Clearly, computerization has shaped the structure of occupational change and the evolution of skill demands. But it is harder to see the channel through which computerization could have dramatically reduced labor demand after 1999. As documented in Figure 9, the onset of the weak U.S. labor market of the 2000s coincided with a sharp *deceleration* in computer investment—a fact that appears first-order inconsistent with the onset of a new era of capital-labor substitution. Moreover, the U.S. labor market woes of the last decade occurred alongside extremely rapid economic growth in much of the developing world. Indeed, frequently overlooked in U.S.-centric discussions of world economic trends is that the 2000s was a decade of rising world prosperity and falling world inequality. It seems implausible to me that technological change could be enriching most of the world while simultaneously immiserating the world’s technologically leading nation.

My suspicion is that the deceleration of the U.S. labor market after 2000, and further after 2007, is more closely associated with two other macroeconomic events. A first is the bursting of the “dot-com” bubble, followed by the collapse of the housing market and the ensuing financial crisis, both of which curtailed investment and innovative activity. A second is the employment dislocations in the U.S. 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. As documented by Autor, Dorn and Hanson (2013), Pierce and Schott (2013) and Acemoglu et al. (2014), China’s rapid rise to a premier manufacturing exporter had far-reaching impacts on U.S. workers, reducing employment in directly import-competing U.S.

manufacturing industries and depressing labor demand in both manufacturing and non-manufacturing sectors that served as upstream suppliers to these industries.⁴⁴

Globalization, like technological change, is not typically Pareto improving, particularly in the short run. While the long-run effects of these developments should in theory be positive, the adjustment process, as with technological adaptation, is frequently slow, costly, and disruptive.

⁴⁴ Borjas and Ramey (1995) present evidence from the 1950s through 1980s that rising foreign competition in durable good industries increased U.S. wage inequality by eroding rents accruing to non-college workers.

References

- Acemoglu, Daron, and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics*, Vol. 4, Part B, edited by Orley Ashenfelter and David Card: 1043–1171. Amsterdam: Elsevier.
- Acemoglu, Daron, David H. Autor, David Dorn, Gordon Hanson, and Brendan Price. 2014. "Import Competition and the Great U.S. Employment Sag of the 2000s." MIT Working Paper, August.
- Akst, Daniel. 2013. "Automation Anxiety." *The Wilson Quarterly*, Summer.
- Alpert, Andrew, and Jill Auyer. 2003. "Evaluating the BLS 1988–2000 Employment Projections." *Monthly Labor Review*, October.
- Andreopoulos, Alexander, and John K. Tsotsos. 2013. "50 Years of object recognition: Directions forward." *Computer Vision and Image Understanding*, 117(8), 827–891.
- 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), 185–199.
- Autor, David H. 2014a. "The Paradox of Abundance: Automation Anxiety Returns." MIT working paper, May.
- Autor, David H. 2014b. "Skills, Education, and the Rise of Earnings Inequality Among the 'Other 99 Percent.'" *Science*, 344(6186), 843–851.
- Autor, David H., and David Dorn. 2013a. "The Growth of Low Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review*, 103(5): 1553–1597.
- Autor, David H., and David Dorn. 2013b. "How Technology Wrecks the Middle Class." *New York Times*, 24–27.
- 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): 2121–2168.
- Autor, David H, Lawrence F Katz, and Melissa S Kearney. 2006. "The Polarization of the U.S. Labor Market." *American Economic Review Papers and Proceedings*, 96 (2), 189–194.
- 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), 415–426.
- British Broadcasting Company. 2012. "Google Computer Works out How to Spot Cats." June 26. Available at: <http://www.bbc.com/news/technology-18595351>.
- Beaudry, Paul, David A. Green, and Benjamin M. Sand. 2013. "The Great Reversal in the Demand for Skill and Cognitive Tasks." NBER Working Paper #18901, December.
- Beaudry, Paul, David A. Green, and Benjamin M. Sand. 2014. "The Declining Fortunes of the Young Since 2000." *American Economic Review*, 104(5), 381–386.
- Blinder, Alan S. 2009. "How Many US Jobs Might Be Offshorable?" *World Economics* 10 (2): 41–78.
- Blinder, Alan S., and Alan B. Krueger. 2013. "Alternative Measures of Offshorability: A Survey Approach." *Journal of Labor Economics* 31 (2) (April): S97–S128.
- Borjas, George J., and Valerie A. Ramey. 1995. "Foreign Competition, Market Power, and Wage Inequality." *Quarterly Journal of Economics*, 110 (4), 1075–1110.
- Brynjolfsson, Erik, and Andrew McAfee. 2011. *Race Against the Machine*. Digital Frontier Press.
- Caroli, Eve and John Van Reenen. 2001. "Skill-Biased Organizational Change? Evidence from A Panel of British and French Establishments." *The Quarterly Journal of Economics*, 116(4), 1449–1492.
- Clark, Colin. 1951. *The Conditions of Economic Progress*. London: Macmillan.
- Cortes, Guido Matias, Nir Jaimovich, Christopher J. Nekarda, and Henry E. Siu. 2014. "The Micro and Macro of Disappearing Routine Jobs." NBER Working Paper #20307, July.
- Drum, Kevin. 2013. "Welcome, Robot Overlords. Please Don't Fire Us?" *Mother Jones*, May/June.
- Firpo, Sergio, Nicole M. Fortin, and Thomas Lemieux. 2013. "Occupational Tasks and Changes in the Wage Structure Occupational Tasks and Changes in the Wage Structure." University of British Columbia Working Paper, August.
- Foote, Christopher L. and Richard W Ryan. 2014. "Labor-Market Polarization Over the Business Cycle." Federal Reserve Bank of Boston Working Paper, April.
- Friedman, Thomas L. 2010. "The Election That Wasn't." *The 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): 118–133.
- Goos, Maarten, Alan Manning, and Anna Salomons. 2014. "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring." *American Economic Review*, 104 (8), 2509–2526.
- Grabner, Helmut, Juergen Gall, and Luc Van Gool. 2011. "What Makes a Chair a Chair?" in *Computer Vision and Pattern Recognition 2011*, IEEE. 1529 – 1536.
- Grier, David Alan. 2005. *When Computers Were Human*. Princeton, NJ: Princeton University Press.
- Heilbroner, Robert. 1965. Men and Machines in Perspective. *Public Interest*, 1: 27-36.
- The Herald Post*. 1966. "Skirting the Automation Question." February 7: p. 2.
- Hoffman, William, and Leo Furcht. 2014. *The Biologist's Imagination: Innovation in the Biosciences*. New York, NY: Oxford University Press.
- 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.
- Katz, Lawrence F., and Robert A. Margo. 2013. "Technical Change and the Relative Demand for Skilled Labor: the United States in Historical Perspective." NBER Working Paper No. 18752, February.
- Keynes, John Maynard. 1933. "Economic Possibilities for our Grandchildren (1930)." *Essays in Persuasion*: 358–373.
- Le, Quoc V., Marc Aurelio Ranzato, Matthieu Devin, Greg S. Corrado, and Andrew Y. Ng. 2012. "Building High-Level Features Using Large Scale Unsupervised Learning." In *Proceedings of the 29th International Conference on Machine Learning*. Edinburgh, Scotland, UK.
- Levy, Frank and Richard J. Murnane. 2004. *The New Division of Labor: How Computers Are Creating the Next Job Market*. Princeton, NJ: Princeton University Press.
- Markoff, John. 2012. "How Many Computers to Identify a Cat? 16,000." *New York Times*, June 25.
- Mazzolari, Francesco 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 25 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.
- Mokyr, Joel. 1990. *The Lever of Riches: Technological Creativity and Economic Progress*. New York: Oxford 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. 2013. "The Surprisingly Swift Decline of U.S. Manufacturing Employment." Yale Department of Economics Working Paper, November.
- 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]. Minneapolis: University of Minnesota.
- Sachs, Jeffrey D., and Lawrence J. Kotlikoff. 2012. "Smart Machines and Long-Term Misery." NBER Working Paper No. 18629, December.
- Schultz, Theodore W. 1975. "The Value of the Ability to Deal with Disequilibria." *Journal of Economic Literature*, 13(3), 827–846.
- Smith, Christopher L. 2013. "The Dynamics of Labor Market Polarization." Federal Reserve Board, Finance and Economics Discussion Series #2013-57, August.
- Smith, Noah. 2013. "The End of Labor: How to Protect Workers From the Rise of Robots." *The Atlantic*, January 14.
- Varian, Hal R. 2014. "Big Data: New Tricks for Econometrics." *Journal of Economic Perspectives*, 28(2), 3–28.
- Welch, Finis. 1970. "Education in Production." *The Journal of Political Economy*, 78(1), 35–59.