University Press Scholarship Online
Oxford Scholarship Online

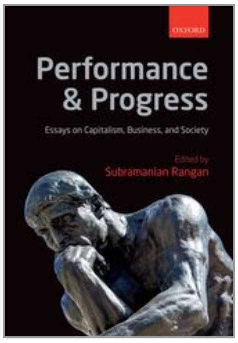

Performance and Progress: Essays on Capitalism, Business, and Society
Subramanian Rangan

Print publication date: 2015
Print ISBN-13: 9780198744283
Published to Oxford Scholarship Online: September 2015
DOI: 10.1093/acprof:oso/9780198744283.001.0001

Paradox of Abundance

Automation Anxiety Returns

David H. Autor

DOI:10.1093/acprof:oso/9780198744283.003.0017

Abstract and Keywords

Despite sustained increases in material standards of living, fear of the adverse employment consequences of technological advancement has recurred repeatedly. This represents a paradox of abundance: technological change threatens social welfare not because it intensifies scarcity but because it augments abundance. For most citizens of market economies, the primary income-generating asset they possess is their scarce labor. If rapid technological advances were to effectively substitute cheap and abundant capital for (previously) expensive and willful labor, society would be made wealthier, not poorer, in aggregate, but those who own labor but do not own capital might find it increasingly challenging to make a living. This chapter considers why automation anxiety has suddenly become salient in popular

and academic discourse. It offers informed conjectures on the potential implications of these developments for employment and earnings.

Keywords: abundance, automation, capital-biased technological change, labor share of income, task structure of production

The Paradox of Abundance

Anxiety about the adverse effects of technological change on employment has a venerable history. In the early nineteenth century, for example, a group of English textile artisans calling themselves the Luddites staged a machine-trashing rebellion. Their brashness earned them a place (rarely positive) in the lexicon. Economists have historically rejected what we call 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, this idea is demonstrably false. In 1900, for example, 41 percent of the United States workforce was in agriculture. By 2000, that share had fallen to 2 percent, after the Green Revolution revolutionized crop (p.238) yields. But the employment-to-population ratio rose over the twentieth century as women moved from home to market, and the unemployment rate fluctuated cyclically, with no long-term increase.

Despite sustained increases in material standards of living, fear of the adverse employment consequences of technological advancement recurred repeatedly in the twentieth 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’s 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 Press* 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” (p.239) (quoted in Akst 2014).[1] 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” (1965: 34–6).

In the five decades since the Ad Hoc Committee penned its open letter to the President, human labor has certainly not been rendered redundant, as these scholars had feared. But automation anxiety has clearly returned. Casual empiricism suggests that economists and public intellectuals have begun to question whether these earlier projections of technological unemployment were in fact flat-out wrong, as had been widely accepted, or whether instead they were simply ahead of their time in anticipating imminent employment challenges that in reality took several additional decades to materialize. For example, in a 2012 *New York Times* column titled “Rise of the Robots,” Paul Krugman cites the falling share of payments to labor in US national income as a harbinger of things to come: “If this is the wave of the future, it makes nonsense of just about all the conventional wisdom on reducing inequality. Better education won’t do much to reduce inequality if the big rewards simply go to those with the most assets.” Krugman is not alone among economists in invoking this concern. The *Economist* (Jan. 18, 2014) reports that

> Larry Summers, a former American treasury secretary [and Clark Medal winner in economics], looked at employment trends among American men between 25 and 54. In the 1960s only one in 20 of those men was not working. According to Mr. Summers’s extrapolations, in ten years the number could be one in seven. This is one indication, Mr. Summers says, that technical change is increasingly taking the form of “capital that effectively substitutes for labour.”

In a similar vein, MIT scholars Erik Brynjolfsson and Andrew McAfee argue in a 2011 book that humans are in danger of losing the “race against the machine.” And in a 2012 working paper, economists Jeffrey D. Sachs and Laurence J. Kotlikoff posit that “smart machines” may threaten us with “long-term misery.”[2]

Perhaps most telling is the finding of a recent poll of leading mainstream academic economists conducted by the Chicago Initiative on Global Markets regarding the impact of technology on employment and earnings.[3] Consistent with the canonical (p.240) 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” (see Figure 14.1). Yet, surprisingly, 43 percent of those polled endorsed (i.e. agreed with) the statement that “information technology and 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.[4] 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.

The concern that technological progress may harm a substantial fraction of workers presents a paradox of abundance. The paradox is that the threat to social welfare posed by technological change is the threat of *excess* rather than the threat of *scarcity*. Why is excess threatening? For most citizens of market economies, the primary income-generating asset they possess is their scarce labor. If rapid technological advances were to effectively substitute cheap and abundant capital for scarce and demanding labor, society would be made wealthier, not poorer. But this capital-biased technological progress would create a substantial income distribution problem: those who own labor but who do not own capital might have no means of making an adequate living.[5]

This would disrupt our central mechanism for economic organization and dramatically dis-equalize the income distribution, even at current high levels of inequality. Thus, the paradox: abundance threatens social welfare.

How did we reach a point where the robust faith of mainstream economists in the beneficence of technological advancement appears to have become decidedly tentative? And is there now a strong case for concern about the long-term consequences of advancing technologies—computers and robotics specifically—for employment and earnings?

This chapter presents evidence that labor scarcity has declined in rich countries. It then considers the current trajectory of technological advancement and considers why automation anxiety has suddenly become salient in popular and academic discourse. It ends by considering (more accurately, speculating on) implications for employment and earnings. (p.241)

(p.242)

Is Labor Scarcity Declining?

Three salient patterns in US and international data suggest that labor may indeed have become less scarce. A first is that labor's share of national income has declined in the large majority of countries since the early 1980s. Figure 14.2, reproduced from

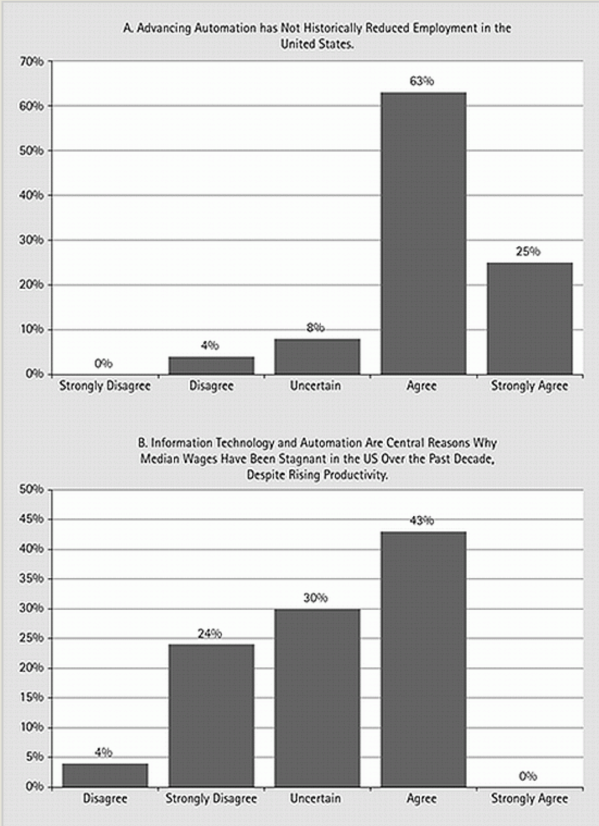

figure 14.1. Chicago Booth IGM Expert Poll: Impact of Automation on Employment and Wages

Notes: Survey date February 25, 2014. Details available at <http://www.igmchicago.org/igm-economic-experts-panel/poll-results?SurveyID=SV_eKbRnXZWx3jSRBb.>

Karabarbounis and Neiman (2014, KN hereafter), documents this trend specifically for the four largest world economies: the United States, Japan, China, and Germany. In all four countries, the labor share—specifically, the share of corporate gross value-added paid to labor—declined by roughly 2 to 4 percentage points per decade during the 1975–2010 period, with the precise time window differing by country according to data availability.

As evidenced by the remarks quoted from Krugman and Summers, many economists find these facts startling. Karabarbounis and Neiman attribute the decline in labor's share of national income to a fall in the price of investment relative to consumption goods—i.e. due to rising capital-labor substitution (a technological change). This form of technological change is of course only one possible explanation for this pattern, and there is as of yet no *direct* evidence linking the falling labor share of income to direct capital-labor substitution.[6] Nevertheless, if the patterns documented by KN prove robustly true and enduring, they suggest that something has profoundly changed in the macroeconomy that has reduced the "scarcity" value of labor.

A second pattern adding to the case for concern is the sharp falls in real wage levels of non-college workers in a number of advanced countries in recent decades, despite the *decline* in the relative supply of these workers. In the United States, this is seen particularly in the declining wages of non-college males evident in Figure 14.3. Between 1979 and 2012, real full-time weekly earnings of male high school graduates fell by approximately 15 percent while those of male high school dropouts fell by more than 25 percent. In a similar vein, Green and Sand (2013; Figure 14.1) document sharp falls in real wages in the bottom four deciles of the Canadian wage distribution between 1981 and 1996, while Card et al. (2013; Figure 14.1) document a substantial decline in the real daily wages of West German male workers between 1997 and 2009 from the median on downward.[7]

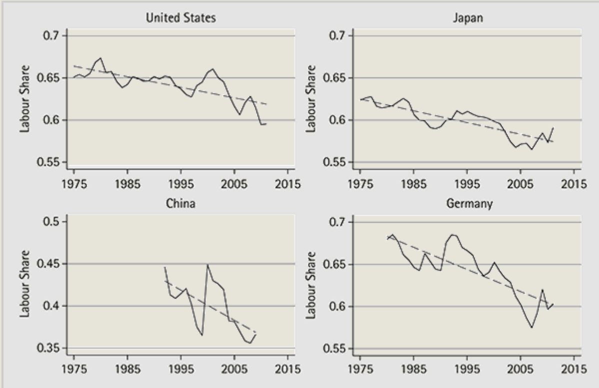

figure 14.2. Trends in the Labor Share in National Income in the Four Largest World Economies

Source: Karabarbounis and Neiman (2014), figure II.

How economically important are these wage declines? One gauge of their significance is their effect on labor force participation. Figure 14.4, which plots changes in employment-to-population rates between 1979 and 2008 among males ages 25 through (p.243)

39 by race and education group against changes in their real hourly wages, offers a third major cause for concern. Employment rates have fallen sharply among demographic groups that have seen the large fall in wages over the last three decades.[8] These declines are substantial, ranging from 7 to 10 percentage points among males with high school or lower education, and far greater among black males. Such employment declines would not necessarily be problematic if they were concentrated among groups with high and rising earnings. This would merely suggest that well-off groups were spending their growing resources on additional leisure—arguably a sign of the rising abundance of leisure that Keynes envisioned in 1930. The fact that employment rates have

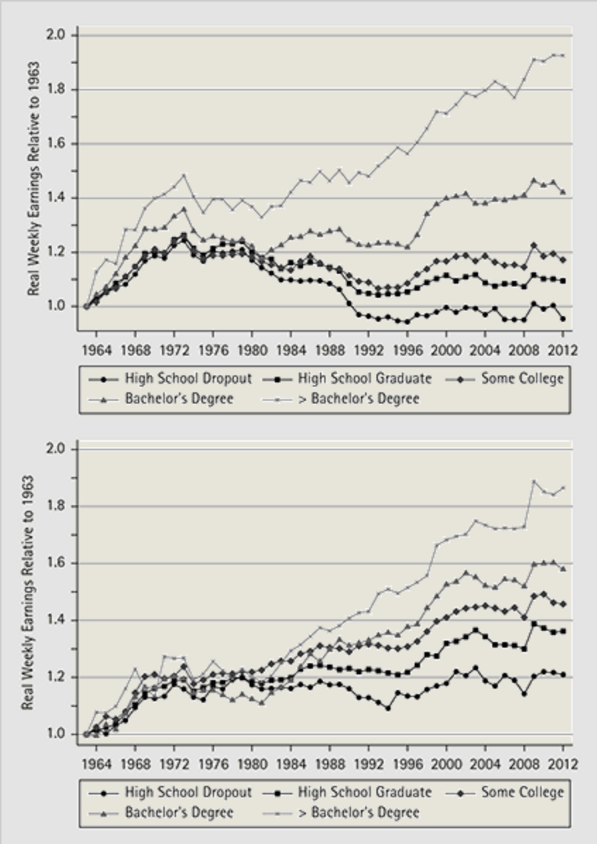

figure 14.3. Change in Real Wage Levels of Full-Time Male (top) and Female (bottom) Workers by Education, 1963–2012

Notes: Real earnings levels are plotted relative to their 1963 values. Wages are deflated to real 2012 values using the Personal Consumption Expenditure Deflator. Figure uses March CPS data for earnings years 1963–2012. Calculations hold constant labor market experience within each education group.

instead dropped steeply among demographic groups with low and falling earnings suggests that employer demand for less skilled workers has declined—so much so that many are either choosing not to work, or are unable to find gainful employment at prevailing wages.[9] Thus, the combined weight of the (p.244) evidence in Figures 14.2 through 14.4 lends credence to the concern that we have entered a realm where there is a growing surplus of labor—or at least a surplus of less-educated labor. These workers may not be “technologically unemployed” in the sense that scholars including Heilbroner or Myrdal had feared; it is plausible that many could (p.245) still command a positive market wage. But if a significant fraction of young, less-educated adults has substantially withdrawn from market employment due to falling demand for their skills, this may be operationally equivalent to technological unemployment.

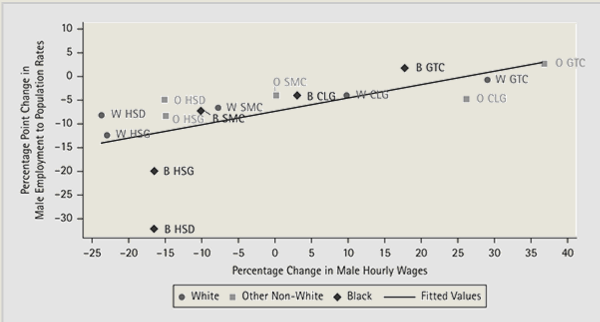

figure 14.4. Changes in Male Employment to Population Rates and Changes in Real Male Hourly Earnings among Males Aged 25–39, 1979–2008

Source: Census IPUMS 5 percent samples for years 1980, 1990, and 2000 and American Community Survey (ACS) 2009. (Reproduced from Autor and Wasserman, 2013: figure 6).

The Puzzle of Falling Wages

There is no economic law that says that wages must always rise. Under normal competitive conditions, an increase in the supply of a given skill group will reduce its market wage. The falling wages of low-skill workers could, therefore, reflect nothing more interesting than a rise in their relative supply. Yet, in essentially all advanced economies—and certainly in the US, UK, and Germany—this is not what has occurred; workers with tertiary education have become increasingly abundant in recent decades while workers with secondary or lower education have becoming increasingly rare. Provided that high- and low-skill workers are gross complements (formally, provided that the elasticity of substitution in production between these skill groups exceeds one), an increase in the relative supply of high-skill workers should *reduce* the real wages of high-skill workers and *raise* the wages of low-skill workers (see Acemoglu 2002, and Acemoglu and Autor 2011 for discussion). Neither has occurred.

A second candidate interpretation of these demographic and wage patterns is that there have been “skill-biased” demand shifts that have raised demand for high-relative (p.246) to low-skill labor. Indeed, a considerable body of evidence suggests that such shifts have occurred both in recent decades and throughout most of the twentieth century (see Katz and Autor 1999; Autor et al. 2008; Goldin and Katz 2008). However, while a skill-biased demand shift will raise the *relative* wages of high- relative to low-skill workers (again assuming the elasticity of substitution exceeds one), such a shift would *not* be expected to reduce real wages of low-skill workers. In fact, the opposite should occur: both high- and low-skill workers should experience an increase in earnings, though high-skill workers should gain by more.[10] The fact that real wages of high-skill workers have risen while those of low-skill workers have fallen in the face of a *falling* relative supply of low-skill workers is therefore inconsistent with *either* a supply-induced rise in the skill premium or a canonical skill-complementary technological change.

What else might be going on? It is likely that the causes for the sharp falls in real earnings among non-college workers are multiple, and it would be incorrect to conclude that technological change is the exclusive or even, necessarily, the primary explanation. One central factor that may have contributed to declining wages of less-educated workers is the globalization of labor markets, seen particularly in the greatly increased US trade integration with developing countries. Globalization has become particularly important for US labor markets since the early 1990s when China began its extremely rapid integration into the world trading system. Between 1987 and 2007, the share of total US spending on Chinese goods rose from under 0.5 percent to close to 5 percent. While the influx of Chinese goods lowered consumer prices, it also fomented a substantial decline in US manufacturing employment, contributing directly to the decline in production worker employment (Autor et al. 2013).

A second factor impinging on the earnings of non-college males is the decline in the penetration and bargaining power of labor unions in the United States. Unions have historically obtained relatively generous wage and benefit packages for blue-collar workers. Over the last three decades, however, US private-sector union density—i.e. the fraction of private-sector workers who belong to labor unions—has fallen by approximately 70 percent, from 24 percent in 1973 to 7 percent in 2011 (Card et al. 2004; Hirsch 2008). While the precise contribution of declining unionization to the evolution of male wage levels and wage inequality is a subject of ongoing debate, a number of studies place this contribution at 20 to 30 percent. Notably, because union membership has been historically quite concentrated among blue-collar workers, the majority of whom are males, the decline in union membership may have differentially affected non-college male earnings.

A third possibility, one which is the focus of this chapter, is that the ongoing substitution of computer-intensive machinery for workers performing routine task-intensive jobs has depressed demand for workers in both blue-collar production and (p.247) white-collar office, clerical, and administrative support positions, and reduced the set of middle-skill career jobs available to non-college workers more generally (Autor et al. 2003; Autor and Dorn 2013). I discuss this possibility in detail next.

It bears emphasis that these three forces—technological change, deunionization, and globalization—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.

How Computerization Changes Work: A Concrete Characterization

Economists frequently speak in abstract terms about capital-skill complementarity and capital-labor substitution—and with some justification, since these terms have precise meanings in the abstract production functions that the economics profession uses to represent economic processes. I find it useful, however, to conceptualize these terms concretely as reflecting distinctive technological phenomena with specific characteristics. Very roughly, one may characterize the recent phases of workplace computerization as undergoing three successive epochs: simulation, communications, and engagement.[11] The first is well understood, the second much less so, and the third reflects the current frontier. Its economic implications are a *terra incognita*.

Simulation

The notion of using computers to simulate (or replicate) codified, repetitive information-processing tasks stretches back to the dawn of the computer era. An early example 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 (p.248) of the first nuclear warheads.[12] But the scope of computer simulation is not limited to simulating *physical* processes. It includes simulating any work process that can be fully specified with an explicit procedure, and hence codified in a computer program. 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.[13]

The implications of computer simulation for work organization are reasonably well captured by the ideas set forth in Autor et al. (2003, ALM hereafter). ALM describe the process whereby computers substitute for workers in performing “routine” codifiable tasks. Routine tasks are characteristic of many middle-skilled cognitive and manual activities, such as bookkeeping, clerical work, and repetitive production tasks. Because the core job tasks of these occupations follow precise, well-understood procedures, they have in recent decades become increasingly codified in computer software and performed by machines. This has led to a substantial decline in employment in clerical, administrative support, and, to a lesser degree, production and operative employment.

But simulation as a computing paradigm has clear boundaries: programmers cannot write a program 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 humans understand tacitly and accomplish effortlessly for which they do not know the explicit “rules” or procedures. In the words of philosopher Michael Polanyi (1966), “We know more than we can tell.” When we ride upright on a two-wheeled bicycle, recognize different species of birds based only on a blurry glimpse from afar, write a compelling paragraph, or develop a hypothesis to explain a poorly understood phenomenon, we are engaging in tasks that we only tacitly understand how to perform.

As ALM discuss, the applicability of “simulation” to accomplishing human work tasks is particularly constraining in two broad areas, which they term “abstract” and “manual” tasks. These lie at opposite ends of the occupational skill distribution. Abstract tasks require problem-solving capabilities, intuition, and persuasion. They typically employ workers with high levels of education and analytical capability. The secularly falling price of accomplishing routine tasks using computer capital complements the “abstract” creative, problem-solving, and coordination tasks performed by highly educated workers such as professionals and managers, for whom data analysis is an input into production.

In contrast, manual tasks require situational adaptability, visual and language recognition, and in-person interactions. These tasks are characteristic of the jobs performed by janitors and cleaners, home health aides, construction laborers, security (p.249) personnel, and motor vehicle operators. They demand workers who are physically adept and, in some cases, able to communicate fluently in spoken language. They appear to require little in the way of formal education, however, at least relative to a labor market where most workers have completed high school.

This latter observation (low education and training requirements) applies with particular force to service occupations. Tasks such as food preparation and serving, cleaning and janitorial work, grounds cleaning and maintenance, in-person health assistance by home health aides, and numerous jobs in security and protective services, are highly intensive in non-routine manual tasks. These are not highly skilled activities by human labor standards, but they 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—which is likely to mute the potential for rapid wage growth in these occupations even in the face of rising demand.[14]

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—it is natural to suspect that computer “simulation” of routine job tasks has contributed to a “polarization” of employment opportunities. A large body of US and international evidence confirms this intuition: by reducing job opportunities in middle-skilled clerical, administrative, production, and operative occupations, computerization is strongly associated with employment polarization at the level of industries, localities, and national labor markets (Goos and Manning 2007; Autor and Dorn 2013; Michaels et al. 2014; Goos et al. forthcoming).

The implications of this process for employment and earnings are multivalent. For highly educated workers, computerization has almost certainly complemented their skills—raising their productivity and the scale of operations they can command, with attendant increases in relative and real earnings (Autor et al. 1998). For less-educated workers, the implications are ambiguous at best. On the one hand, the displacement of workers from middle-skill clerical, administrative support, production, and operative positions likely leads to downward occupational mobility toward less highly trained service positions. This undoubtedly places downward pressure on wages in these occupations. At the same time, it is possible for the real "value" of the output of (p.250) services to rise as societal wealth increases and the scarcity value of machine-produced output falls (e.g. think of large-screen TVs). Thus, while it is *possible* but far from certain for workers at all levels to benefit, the weight of the evidence suggests this has not for the most part occurred, particularly in the last decade.[15] My unproven hunch is that the net effect of the wave of computer "simulation" of workplace tasks has been to depress the earnings, and ultimately the employment, of less-educated workers.

Communications

Starting in the early 1980s, the advancing capabilities of computers in simulation were complemented by advances in telecommunications. Although large organizations such as airlines, banks, and (of course) the military had been harnessing telecommunications to connect computers for decades, price declines and technological advances in the 1980s and 1990s made computer communications ubiquitous and powerful. The commercialization of inexpensive local area networking technologies (Ethernet, specifically) in the early 1980s enabled firms to network computers within a workplace to share data and resources (e.g. file servers and printers). Before long, local area networks were interconnected in “wide area networks” within organizations, allowing the personal, mini, and mainframe computers belonging to a single organization to communicate across disparate locations over dedicated backbones. The opening of the internet to civilian and commercial use in 1995 provided firms with a set of protocols and non-dedicated infrastructure that ultimately enabled any digital device to communicate with any other internet-connected digital device anywhere in the world. Even more recently, the deployment of high-speed mobile networks has enabled digital devices to remain continuously connected to the internet over a large portion of the world’s populated land areas (and at sea or in the air via satellites).

How do these enhanced capabilities—ubiquity, high-speed communications, and a limitless set of “online” resources—expand or reshape the simulation paradigm? I do not pretend to have the complete answer to this question, but I see at least two profound consequences.[16] One is that the marriage of computing and communications makes it far easier for computers to take on a coordination or oversight role than was conceivable in the “simulation era”—for example, dispatching trucks, routing packages, orchestrating the flow of parts on an auto assembly line, or dynamically managing (p.251) the layout, restocking, and order retrieval from a warehouse.[17] These examples are all, arguably, simulation tasks in that computers are “simply” calculating, optimizing, and controlling following a procedure set down by a programmer. However, unlike the examples of payroll processing or census enumeration, computers in these applications are interacting in real time with the environment. Sensing and communications technologies give computers the ability to monitor a disparate set of activities in continuous time and issue instructions to coordinate and adapt these activities as new data arrive or conditions change.

One prominent application that builds on these capabilities is online sales. Prior to the communications era of computing, the primary functions of computers in retail sales were to track inventory and assist cashiers in customer checkout. The advent of cheap, ubiquitous computer communications vastly expanded the range of sales-related activities subject to computerization. Contemporary business-to-consumer websites showcase products in virtual electronic malls, recommend alternative or complementary purchases based on the search behavior of the current and past users, verify the buyer's identity, conduct the financial transaction, move the order into the back-end fulfillment system, and notify the purchaser, seller, and shipper of the transaction's status as it unfolds.

One can object that these activities are simply glorified "simulation": online sales systems are, in effect, carrying out the codified steps of tracking inventory, displaying products, transacting purchases, and tracking fulfillment.[18] But this interpretation strikes me as reductive. Fifteen years ago, one might have persuasively argued that computers could not effectively substitute for salespersons because they are unable to showcase products, make product recommendations, offer expertise and advice on product suitability and features, and more generally cannot "get to know" the customer. Those predictions would have been technically correct but substantively wrong. While it remains the case that online storefronts cannot measure your shoe size or help you to lace up a pair of oxfords, the virtues of convenience, broad selection, abundant information, and informative product recommendations (based on the behavior of countless prior buyers) appear in many cases to trump the virtues of in-person sales. Notably, the genesis of these advances is *not* simply that "simulation" has improved. The key enabler is communications: online storefronts serve the customer from any location and at any time, and provide remarkably rich decision-relevant information (photographs, product specifications, user reviews, how-to videos), in many cases exceeding what an expert salesperson could offer.

(p.252) Computer communications—and the internet in particular—also enable a set of information-based services that arguably had no close counterpart in the pre-communications area: search engines. Search engines draw on systems of network computers to provide services at zero marginal cost that, until recently, were both time- and resource-intensive to obtain: rapid, accurate search and retrieval of specialized information from encyclopedias, medical reference manuals, travel books, publications databases, and how-to libraries. Search engines have become such an essential tool that a substantial fraction of today's high school students have probably never looked up a historical fact, investigated a medical or recreational drug option, or checked the prevailing spring weather in another country using any tool but a search engine.

Again, one can argue that computers are merely "simulating" what a skilled research librarian would do if she had access to many of the world's best libraries, and also had time to read and memorize their full contents for instant recollection. But the absurdity of the comparison highlights a critical set of differences. The methods that search engines use to "search" for information are so different from how humans search for information (absent computers) that it is inaccurate to characterize computers as "simulating" human search. Humans do not read, memorize, and sort limitless amounts of information for later retrieval. Instead, they catalog where information is likely to be found (using the Dewey decimal system, travel guides, encyclopedias, journal indexes) and make directed, purposive searches within those locations to identify specific pieces of information. Humans have limited information absorption and recall capability, but they can use context and logic to quickly narrow the scope of a search to the logical locations where the information is likely to reside (e.g. to look up the historical population of Manhattan, I would consult old census volumes).

One focal contrast between human and machine search helps to highlight these differences. Human search techniques require prior organization and cataloging of information; attempting to search for a specific fact in a library where all of the books had been randomly distributed across shelves would be fruitless. Such a library would, however, pose no problem for a search engine; in fact, the World Wide Web presents an electronic facsimile of this type of library—a vast repository of disorganized information. Unlike humans, search engines are indiscriminate information collectors—absorbing vast amounts of data without specific organizational principles or explicit understandings of how one piece of information may be connected to another. When a user performs a query, the search engine offers its best guess at what the user is seeking based largely on statistics: what patterns the user's query matches, what similar queries this and other users have posted in the past. Stated compactly, human search is directive—guided by prior knowledge and context. Computer search, by contrast, is statistical and non-directive. And the differences between these approaches are dictated by the distinct information-processing capabilities of people and machines.

The power of online search also highlights the complementarity between successive waves of information technologies—specifically, simulation and communications. Search engines depend fundamentally on computer communications not only for (p.253) information delivery but also for primary data gathering. Google does not, to a first approximation, *create* the information it serves; it simply aggregates the countless information sources that others have made available through their computer systems. Thus, it is the very existence of computer networks that generates the resources that search engines search over. Search engines, and their close relatives, are meta-technologies that have virtuously—and arguably unexpectedly—emerged from the collective interaction of a vast number of computer systems, many of which are engaged in standard "simulation" tasks.

The power of this "meta" technology is increasingly evident beyond search. Automated "discovery" software reads reams of legal documents disgorged by companies undergoing lawsuits, identifies themes, catalogs contents, and attempts to thread together conversations based on email and paper chains (Markoff 2011). Fraud detection software applies statistical tools to flag suspicious patterns of transactions in real time, and often calls credit card holders to alert them of possible frauds. Recommendation engines suggest music and movies to consumers based on their expressed and revealed tastes, which are aggregated and compared with the browsing and rating tastes of countless other users.

While it would be foolhardy to attempt to infer general labor market implications from these high level observations, it is inarguable that the era of computer communications has substantially expanded the set of tasks beyond that which could be accomplished by computer "simulation." On the one hand, the information presentation and interaction seen in online sales allow computers to accomplish many interactive "manual" tasks that are *not* directly amenable to simulation in the canonical sense, such as order-taking and sales (i.e. the computer does not closely replicate what a human agent would do). On the other hand, the growing sophistication of statistical pattern recognition algorithms enables computers to encroach upon "expert" domains—work that has historically been the province of research librarians, paralegals, travel agents, and teachers.

Engagement

Computerization has recently entered a third era, "machine engagement," in which computers are emerging from their largely passive role as ever-ready information appliances to become increasingly "alert" machines—aware of people and objects, sensitive to contexts, and able to adjust plans accordingly to accomplish useful tasks.[19] One does not have to look hard to find early examples of "engaged" machines:

- Smartphones interpret and respond to voice commands based upon both verbal and contextual clues—where the user is currently located (e.g. home or work), (p.254) what events are scheduled on the calendar, what names are present in the address book, and what commands and queries the user has supplied in the past.
- Robotic vacuums (e.g. Neato Botvac) use lasers to scan and map rooms while vacuuming, thus plotting a purposive course over the autonomously mapped cleaning area rather than using the traditional "bump and turn" principle used by earlier generations of self-propelled devices.[20]
- Self-driving cars (e.g. the Google Car) semi-autonomously pilot conventional vehicles (retrofitted with sensors and actuators) along moderately complex suburban and urban roads—applying the accelerator, operating the steering

> wheel, complying with road signs and speed limits, and braking for unexpected hazards. Because robotic vehicles are never distracted, drowsy, or temperamental, it is a certainty that they will ultimately prove safer drivers than human operators.

These recent advances sit atop prior waves of computer simulation and telecommunications (as well as dramatic hardware price declines). Laser sensing and object recognition tools harness “simulation” software for digitizing physical environments. Location and contextual awareness technologies exploit mobile data connections to access digitized maps and search engines.

Distinct from earlier waves of computerization, recent advances in machine “engagement” with humans do *not* rely upon conventional computer simulation. Because these engagement tasks remain, to a substantial extent, unsolved problems in science and engineering, contemporary artificial intelligence has devised an “end run” around the problem. Rather than explicitly codifying such tasks, statistical machine learning algorithms inductively learn these tasks through a process of exposure, training, and reinforcement. This process enables computers to (in some cases) accomplish non-codified problems that, while remarkably mundane for humans, remain daunting challenges for engineering.

As one concrete example of machine learning, consider the challenge of object recognition, specifically the task of visually identifying a chair. Chairs come in innumerable varieties: some of have four legs, some have three, other have none; chairs may or may not have backs, may or may not rotate, swivel, or telescope, may or may not be upholstered, may or may not have arms; chairs may be comprised of any number of materials; and chairs may be highly stylized or unconventional. For example, the 1932 Zig Zag chair, designed by Gerrit Rietveld, is shaped like an upside down letter Z with an extra perpendicular ascender protruding from the top. It lacks distinct legs, arms, or an ergonomic seat or back. Nevertheless, most grade school children would immediately recognize the Zig Zag chair for what it is: a chair.

But this is not the case for most object recognition programs. Why not?

(p.255) Applying the "simulation" paradigm, a programmer might attempt to specify ex ante what features of an object suggest that it is a chair—it possesses legs, arms, a seat, and a back, for example. One could then program machines 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 simple "simulation" approach to object recognition—and many more sophisticated variants—would likely have very high misclassification rates.

Why is this ex ante approach unlikely to work? Ultimately, what makes an object a chair is that it is a device purpose-built to facilitate human beings in the act of sitting. Because there are an endless number of ways to accomplish this objective, it is likely almost impossible to pre-specify what attributes an object must possess to be a chair. Accordingly, humans (likely) recognize chairs not (simply) by comparing candidate objects to pre-specified feature sets but, instead, by reasoning about both the attributes of the object and the attributes of the human body to assess whether the candidate object is likely intended to serve as a chair (Grabner et al. 2011). 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 for a good seat. This implies that the problem of object recognition—at least as practiced by the human brain—is far deeper than the problem of determining whether objects have specific attributes; it likely requires reasoning about what an object is "for" and whether it is likely to serve that purpose. 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."

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 largest set of objects would be extremely challenging. 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?

Many contemporary object recognition tools circumvent the reasoning problem by exploiting what some would call “brute force”: applying statistical machine learning tools to infer by example what objects are likely to be chairs. Relying on very large databases of so-called “ground truth”—essentially, a vast set of curated examples of labeled objects—computers can be “trained” to recognize chairs (and other objects) by induction; i.e. they statistically infer what attributes of an object make it more or less likely to be designated a chair. This approach does not require either an explicit model of “chairness” or a model of the human anatomy; instead, it relies only on large training databases, substantial processing power, and of course sophisticated software. Machine-learning algorithms do not, at present, perform as well as grade school children in correctly classifying objects. But the underlying technologies—the (p.256) software, hardware, and training data—are all improving rapidly (Andreopoulos and Tsotsos 2013).

Not surprisingly, the long-term potential of machine learning to circumvent the reasoning problem 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 ultimately approach or exceed human capabilities. Others suspect that the machine learning approach will only ever get it right “on average” while missing many of the most important and informative exceptions. In either case, there is little disagreement that, at present, the ability of machines to “engage” in the human world is substantially constrained by (at least) three attributes of the candidate task:

1. *Structure in the environment*. Machine adaptability to variation in environment is, at present, far less complete, less accurate, and less reliable than human adaptability. It is natural, therefore, that the first (and current) primary application of commercial robotics is on production lines, where the environment is radically simplified and there are few variations in task requirements with which machines must contend (often only a handful of distinct operations). In production settings, industrial robots are typically bolted to the floor and surrounded by large cages that serve to protect nearby humans from their potent combination of superhuman speed and near-complete blindness to their environments.

2. *Degrees of freedom in dexterous interactions.* Though robots probably will eventually be able to walk up and down stairs, load and unload dishwashers, and fold towels, robotic dexterity will be far short of human dexterity for many years to come. It is unlikely that robots will cook fresh meals, sand and paint houses, cut hair, or wrap birthday presents anytime soon.

3. *Richness of perceptual information required to support completion of tasks.* Many mundane daily tasks are deeply dependent upon rich perceptual information. To remove, dust, and replace the objects on a shelf, untie a pair of shoes, or pack a set of items in a suitcase, an agent must recognize non-uniform objects, understand and respect their physical properties (e.g. clothes can be folded in a suitcase but shoes cannot), and make fine visual discriminations (e.g. are the shoelaces single or double-knotted?). These perceptual demands are trivial for human actors but are far outside the realm of machine capability at present.

Of course, a fourth constraint on all of these tasks is cost. While it might be technically feasible to build a robotic dishwasher loader/unloader in the near future, it will not be commercially viable to do so until the cost of numerous digital and mechanical components falls considerably.

What do these observations imply about the trajectory of capital-labor substitution? Again, it would be foolhardy to confidently project general equilibrium economic implications from these stylized characterizations. Nevertheless, it seems very likely that the scope of computer substitution into what ALM described as “manual” tasks is poised to greatly expand in the next ten years. I anticipate that we will see fewer (p.257) housekeepers and janitors, fewer waiters and busboys, fewer vehicle operators, fewer assembly line workers, fewer store stockers and warehouse workers, and fewer salespersons—even in “brick and mortar” shops. At the same time, there will remain core manual task-intensive jobs that are not subject to machine substitution anytime soon: child care, elder care, and health care; food preparation; construction and skilled repair; and numerous dexterous jobs that require high levels of adaptability, precision, and contextual awareness.

While the implications for the *aggregate* labor demand are ambiguous—since these technological advancements both substitute for and complement labor—their implications for *skill* demands appear more readily discerniable. Advances in machine engagement appear poised to have a far greater labor-substituting impact in low education, manual tasks than in high education, abstract tasks. These advances will likely amplify the paradox of abundance: by making low education labor that much less scarce, they will augment inequality even as they generate riches.

Conclusions

Generations of scholars and pundits have worried about the adverse labor market consequences of technological change. Generations of neoclassical economists have assured these thinkers that their worries are misplaced. Though I consider myself a neoclassical economist, I believe that economists’ bland reassurances are becoming less and less convincing. Technological advances have not created the mass unemployment that many feared. But my reading of the evidence is that they have significantly depressed wages among a substantial subset of workers, catalyzing sharp falls in labor force participation. Though declining participation in response to falling wages may be “voluntary,” it is definitely not welfare-improving relative to a setting where non-college workers might be drawn back into the labor force by higher wages. While it is dangerous to extrapolate far into the future based on current trends, I foresee the challenge facing non-college workers becoming more severe as “engaged” machinery increasingly subsumes manual tasks.

There will of course be encroachments upward as well: core job tasks of salespersons, educators, attorneys, engineers, and computer programmers will be increasingly subject to automation. I worry less about these worker groups, both because I think the rate of encroachment will be slower, and because these groups have greater resources and skills to adapt accordingly. But the changes will nevertheless be significant.

Some writers would at this point draw an analogy between the economic eclipse of horses by motorized vehicles in the first decades of the twentieth century and the coming obsolescence of human labor. But there is an important difference between these examples: horses do not own capital and people do. Horses were not made wealthier by the availability of machine substitutes for their labor, but people will be **(p.258)** (collectively) enriched. Thus, the paradox of abundance is not one of impoverishment but one of maldistribution. If technological advances make human labor substantially less scarce—as many have feared, and as Keynes eagerly anticipated—the challenge will not be finding jobs for people to do, but rather finding a means to distribute our abundant societal riches absent labor scarcity as a primary means of income distribution.

Acknowledgments

I thank Frank Levy, Brendan Price, Subi Rangan, Amartya Sen, Seth Teller, participants in the MIT CSAIL/Economists Lunch Seminar, and contributors to the Inaugural Assembly in London in Apr. 2014 for ideas and feedback that helped to shape the chapter.

References

Bibliography references:

Acemoglu, Daron (2002). Technical Change, Inequality, and the Labor Market. *Journal of Economic Literature,* 40(1): 7–72.

Acemoglu, Daron, and David Autor (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. In Orley Ashenfelter and David Card (eds), *Handbook of Labor Economics*, iv/B. Amsterdam: Elsevier, 1043–1171.

Akst, Daniel (2014). What Can we Learn from Past Anxiety over Automation? *Wilson Quarterly* (Summer [online only]).

Andreopoulos, Alexander, and John K. Tsotsos (2013). 50 Years of Object Recognition: Directions Forward. *Computer Vision and Image Understanding*, 117(8), 827–91.

Autor, David H., and David Dorn (2013). The Growth of Low Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review,* 103(5): 1553–97.

Autor, David H., and Melanie Wasserman (2013). *Wayward Sons: The Emerging Gender Gap in Education and Labor Markets*. Washington, DC: Third Way.

Autor, David H., Lawrence F. Katz, and Alan B. Krueger (1998). Computing Inequality: Have Computers Changed the Labor Market? *Quarterly Journal of Economics*, 113(4): 1169–213.

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): 1279–333.

Autor, David H., Lawrence F. Katz, and Melissa S. Kearney (2008). Trends in US Wage Inequality: Revising the Revisionists. *Review of Economics and Statistics,* 90(2): 300–23.

Autor, David H., David Dorn, and Gordon H. Hanson (2013). The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American Economic Review*, 103(6): 2121–68.

Bryniolfsson, Erik, and Andrew McAfee (2011). *Race Against the Machine*. New York: Digital Frontier Press.

Brynjolfsson, Erik, and Andrew McAfee (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: W. W. Norton & Co.

Card, David, Thomas Lemieux, and W. Craig Riddell (2004). Unions and Wage Inequality. *Journal of Labor Research*, 25(4): 519–59.

Card, David, Jörg Heining, and Patrick Kline (2013). Workplace Heterogeneity and the Rise of West German Wage Inequality. *Quarterly Journal of Economics*, 128(3): 967–1015.

Drum, Kevin (2013). Welcome, Robot Overlords. Please Don't Fire Us? *Mother Jones* (May/June [online only]).

Elsby, Michael W. L., Bart Hobijn, and Aysegül Sahin (2013). The Decline of the US Labor Share. *Brookings Papers on Economic Activity* (Fall): 1–52.

Garicano, Luis, and Esteban Rossi-Hansberg (2004). Inequality and the Organization of Knowledge. *American Economic Review*, 94(2): 197–202.

Garicano, Luis, and Esteban Rossi-Hansberg (2006). Organization and Inequality in a Knowledge Economy. *Quarterly Journal of Economics*, 121(4): 1383–435.

Goldin, Claudia, and Lawrence F. Katz (2008). *The Race between Education and Technology*. Cambridge, MA: Harvard University Press.

Goos, Maarten, and Alan Manning (2007). Lousy and Lovely Jobs: The Rising Polarization of Work in Britain. *Review of Economics and Statistics*, 89(1): 118–33.

Goos, Maarten, Alan Manning, and Anna Salomons (forthcoming). Explaining Job Polarization: Routine-Biased Technological Change and Offshoring. *American Economic Review*.

Grabner, Helmut, Juergen Gall, and Luc Van Gool (2011). What Makes a Chair a Chair? In *Computer Vision and Pattern Recognition 2011*. New York: IEEE Press Books, 1529–36.

Green, David A., and Benjamin Sand (2013). *Has the Canadian Labor Market Polarized?* University of British Columbia Working Paper, Nov.

Gregg, Paul, Stephen Machin, and Mariña Fernández-Salgado (2013). *Real Wages and Unemployment in the Big Squeeze.* Centre for Economic Performance Working Paper, Nov.

Grier, David Alan (2005). *When Computers were Human*. Princeton: Princeton University Press.

Heilbroner, Robert (1965). Men and Machines in Perspective. *Public Interest*, 1: 27–36.

Hirsch, Barry T. (2008). Sluggish Institutions in a Dynamic World: Can Unions and Industrial Competition Coexist? *Journal of Economic Perspectives*, 22(1): 153–76. Updated data to 2011 available at <>.

Karabarbounis, Loukas, and Brent Neiman (2014). The Global Decline of the Labor Share. *Quarterly Journal of Economics*, 129(1): 61–103.

Katz, Lawrence F., and David H. Autor (1999). Changes in the Wage Structure and Earnings Inequality. In Orley Ashenfelter and David E. Card (eds), *Handbook of Labor Economics*, iii/A. New York: Elsevier, 1463–555.

Keynes, John Maynard (1930). Economic Possibilities for our Grandchildren. In *Essays in Persuasion.* New York: Classic House Books, 358–73.

Krugman, Paul (2012). Rise of the Robots. *New York Times*, Dec. 8.

Markoff, John (2011). Armies of Expensive Lawyers, Replaced by Cheaper Software. *New York Times*, Mar. 4.

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.

Mokyr, Joel (1990). *The Lever of Riches: Technological Creativity and Economic Progress*. New York: Oxford University Press.

Polanyi, Michael (1966). *The Tacit Dimension*. New York: Doubleday.

Sachs, Jeffrey D., and Lawrence J. Kotlikoff (2012). *Smart Machines and Long-Term Misery.* NBER Working Paper, 18629, Dec.

Smith, Noah (2013). The End of Labor: How to Protect Workers from the Rise of Robots. *The Atlantic*, Jan. 14.

Solow, Robert (1965). Technology and Unemployment. *Public Interest*, 1: 17–26.

The Economist (2014). The Future of Jobs: The Onrushing Wave. Jan. 18.

The Herald Press (1966). Skirting the Automation Question. Feb. 7: p. 2.

Notes:

([1]) The three threats perceived by the ad hoc committee were: the cybernation revolution; the weaponry revolution; and the human rights revolution.

([2]) Of course, popular writing on the topic is far less circumspect. Journalist Kevin Drum opined in *Mother Jones* that we are becoming enslaved to our “robot overlords,” while Noah Smith laments in *The Atlantic* that we have reached “the end of labor.”

([3]) 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 Economic Advisors, and past and current editors of the leading journals in the profession.” Caveat emptor: the author is also a member of the panel.

([4]) Survey results are found at <www.igmchicago.org/igm-economic-experts-panel/poll-results?SurveyID=SV_eKbRnXZWx3jSRBb> (accessed Mar. 2014).

([5]) To clarify terminology, capital-biased technological change is a change in production technology that raises capital’s share of output.

([6]) KN’s model assumes an elasticity of substitution between labor and capital that exceeds one—a necessary condition for a fall in the price of capital goods to raise the capital share. Elsby et al. (2013) closely study the evolution of the US labor share over the period 1948–2013 and corroborate KN’s finding of a substantial decline in the labor share from the early 1980s forward. Their correlational evidence suggests, however, that outsourcing of labor-intensive tasks rather than capital-labor substitution is the largest proximate contributor to declining labor shares at the level of industries.

([7]) Gregg et al. (2013) report that a similar decline in low earnings has not occurred in the UK.

([8]) As reported in Autor and Wasserman (2013), over the entire 1979–2008 period, a 10% fall in wages for a demographic group is robustly associated with a 5.7 percentage point decline in its employment-to-population rate. The positive correlation between rising (or falling) wages and rising (or falling) employment rates holds in each of the last three decades (1979–89, 1989–99, and 2000–10), as well as before and during the Great Recession (2000–7 and 2007–10). The robust positive relationship between wage and employment changes is detected for all demographic subgroups: both sexes, all race groups, both younger and older workers, and both college and non-college workers.

([9]) This latter possibility would suggest that employment is “rationed” among low-skill groups, which cannot occur in a competitive neoclassical labor market setting. Contemporary labor markets are far from the textbook neoclassical model, however. The presence of wage rigidities (such as minimum wages or downward nominal wage rigidities), fixed hiring or firing costs, or significant search frictions, all make involuntary unemployment a plausible possibility.

([10]) The reason is q-complementarity. A high-skill-labor-augmenting technological change increases the effective supply of high-skill workers. Analogously to an increase in the capital/labor ratio, which increases the marginal product of capital in a standard two-factor production function, a high-skill-labor-augmenting technological change should raise the marginal product of low-skill labor.

([11]) This trichotomy is used by my MIT colleague, roboticist Seth Teller. While it is not common parlance in the computer science community, I find it extremely helpful for organizing ideas.

([12]) 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.

([13]) In many cases, the workers who performed these tasks were given the job title of “computer” (Grier 2005).

([14]) Interestingly, employment projections from the US Bureau of Labor Statistics also support the view that low-education service jobs are likely to be a major contributor to US employment growth going forward. The BLS forecasts that employment in service occupations will increase by 4.1 million, or 14%, between 2008 and 2018. The only major occupational category with greater projected growth is professional occupations, which are predicted to add 5.2 million jobs, or 17% (US Bureau of Labor Statistics, Current Employment Statistics, available at <http://www.bls.gov/ces>).

([15]) Autor and Dorn (2013) present evidence that the complementarity effect 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.

([16]) Papers by Garicano and Rossi-Hansberg (2004, 2006) explore how these two distinct eras of computerization—simulation and communications—may have distinct effects on the organization of knowledge hierarchies within firms.

([17]) Kiva Robotics, now owned by Amazon, offers a compelling example of how a warehouse can be reconceived as a dynamic filing system, with robotic “turtles” performing the filing—specifically, transporting shelves from location to location on their backs.

([18]) Similarly, ubiquitous smartphone-based navigation systems are “nothing more” than digitized maps married to route calculation software, off-the-shelf satellite global positioning circuitry, and real-time traffic and road hazard information. Nevertheless, this virtuous combination of data, calculation, and communication has turned vehicle operators from navigators to helmsmen whose primary function is to steer their vehicles safely through road hazards as the computer calls out routing commands.

([19]) Note that “alertness” does not mean machine consciousness—only that machines are aware of and responsive to the human environment, and to people in particular.

([20]) The Botvac employs a technology called Simultaneous Localization and Mapping (SLAM), where an autonomous machine builds up a map within an unknown environment. As Brynjolfsson and McAfee (2014) discuss, SLAM has been a holy grail of artificial intelligence researchers for decades.