

February 7, 2020

Assessing the Impact of New Technologies on the Labor Market: Key Constructs, Gaps, and Data Collection Strategies for the Bureau of Labor Statistics

Contract No: GS-00F-0078M

Submitted to:

**U.S. Department of Labor
Bureau of Labor Statistics**

**Kristin Dell
Nicole Nestoriak**

2 Massachusetts Avenue, NE
Room G225
Washington, D.C. 20212

Submitted by:

Gallup

Jenny Marlar
Project Director

1001 Gallup Drive
Omaha, NE 68102
T 402-938-6249
Jenny_Marlar@gallup.com

Table of Contents

1. Overview	3
2. The Impact of New Technologies on the Labor Market: a Literature Review	4
2.1 Framing the Economic Effects of Technology	4
2.2 Defining New Technologies.....	6
2.2.1 Automation	6
2.2.2 Digital Technologies	7
2.2.3 Artificial Intelligence	7
2.3 Major Economic Theories on New Technologies	8
2.3.1 Lessons From the Industrial Revolutions	8
2.3.2 Skill-Biased Technological Change	10
2.3.3 Task-Based Models	12
2.3.4 Task-Based Models With Automation	13
2.4. Empirical Trends in the Labor Market	15
2.4.1 Broad Historical Patterns in the Structure of Work	15
2.4.2 Demand for Labor in Technology-Producing Industries	17
2.4.3 Demand for Labor in Technology-Using Firms	18
2.4.4 How New Technologies Are Affecting the Labor Share of Income	19
2.5 What Jobs Are at Risk of Technological Displacement?	20
2.6 Conclusion	22
3. Measuring the Impact of New Technologies on the Labor Market	25
3.1 Introduction	25
3.2 Key Constructs.....	25
3.3 Methodology.....	28
3.4 Measurement of Key Constructs	29
3.4.1 Productivity Growth	29
3.4.2 Diffusion of Technology.....	32
3.4.3 Advanced Industries.....	36
3.4.4 Skill and Labor Demand.....	39
3.4.5 Allocation of Tasks	41
3.5 Summary of Major Gaps in Existing BLS Data.....	42
4. Identifying and Filling the Gaps in BLS Data Products for Measuring the Effects of Technology on the Labor Market.....	44
4.1 Introduction	44
4.2 Topic Areas With Identified Gaps	44
4.2.1 The Labor Market for Skills, Occupations, and Tasks	44
4.2.2 Cataloging Tasks and Who or What Performs Them	52
4.3 Conclusion	58
5. References.....	60
Appendix.....	66
Appendix A: Acronyms and Abbreviations.....	66
Appendix B: BLS Data Products.....	68
Appendix C: Datasets From U.S. Principal Statistical Agencies	69
Appendix D: International Statistical Sources	71

1. Overview

The Bureau of Labor Statistics (BLS) of the U.S. Department of Labor is the principal Federal statistical agency responsible for measuring labor market activity, working conditions, and price changes in the economy. Its mission is to collect, analyze, and disseminate essential economic information to support public and private decision-making.¹

Private and public decisions related to labor markets and working conditions are increasingly being influenced by technological considerations. Spurred by a wave of technological developments related to digitization, artificial intelligence (AI), and automation, governments around the world have declared that the creation and deployment of these technologies present both important opportunities and challenges to their citizens.

Having better data related to the labor market and automation technologies could go a long way in helping address the concerns raised by technology. With these issues in the background, the BLS commissioned this report to identify constructs that would complement existing BLS products with a goal of ensuring that the necessary data exist that would allow stakeholders to assess the impact of automation on labor outcomes.

In the first section of this report, we review the social science literature on the effects of automation and related technologies on the labor force and identify some of the key theoretical constructs used in the scholarly literature.² In that section, we find some sources of agreement in the literature and highlight advances in our understanding of the topic. However, the section also underscores that there remains considerable uncertainty about how technology has affected the labor market in recent history and a great deal of uncertainty about the short- and long-term effects in the future. Many of the most influential papers in the field draw upon data sources and methods that have considerable limitations.

In the second section of this report, we review what international and domestic statistical agencies are doing to measure the key theoretical constructs linked to the labor market effects of technology.³ We describe several areas where data gaps are large relative to the theoretical importance of the constructs.

This report concludes with an examination of the key constructs with a view to filling the data gaps identified. The aim is to present options that the BLS should consider in meeting the agency's mission to further the understanding of how technology is affecting and is likely to affect the labor market. Many of the constructs identified are closely related to data that the BLS already collects, and so the final section looks to leverage existing sources as much as possible. In an attempt to create fully developed data collection strategies, we examined publicly available data on the costs of operating various BLS surveys to provide a qualitative assessment of how each option may broadly impact BLS costs.⁴

The primary lesson learned from this report is that researchers and, by extension, policymakers lack the data necessary to fully understand how new technologies impact the labor market. No individual agency or statistical system, in the U.S. or abroad, has developed a comprehensive approach to collecting data on all key constructs needed to assess the impact of AI, automation, and digitization on labor outcomes. These agencies face the challenge of measuring rapidly evolving technologies, as well as the difficulty of parsing a fragmented literature that has not hitherto provided clear guidance on what data are needed. However, by fully reviewing existing research and data collection efforts, this report takes an initial step at bridging this crucial divide and helping the BLS become a leader in this domain.

¹ U.S. Department of Labor, Bureau of Labor Statistics, <https://www.bls.gov/bls/infohome.htm> (accessed October 16, 2019).

² "The Impact of Automation on the Labor Market: A Literature Review," Submitted to U.S. Department of Labor, Bureau of Labor Statistics on June 12, 2019 by Gallup, Inc.

³ "Measuring the Impact of New Technologies on the Labor Market: Technical Paper" Submitted to U.S. Department of Labor, Bureau of Labor Statistics on October 16, 2019 by Gallup, Inc.

⁴ "FY 2020 Congressional Budget Justification Bureau of Labor Statistics," <https://www.dol.gov/sites/dolgov/files/general/budget/2020/CBJ-2020-V3-01.pdf> (accessed October 17, 2019).

2. The Impact of New Technologies on the Labor Market: a Literature Review

2.1 Framing the Economic Effects of Technology

The concept of technology is at the heart of macroeconomic analysis. In standard macroeconomic growth models, labor and capital are the key factors of production that generate economic value (Jones 2016). Basic macroeconomic accounting subtracts the value of these measurable factors (the cost of labor and capital) from Gross Domestic Product (GDP) and describes the residual as productivity growth. In these neoclassical models, this residual productivity growth is the only long-term driver of higher living standards, and it is commonly referred to as “technology.” In the simplest versions of this framework, technology makes labor more productive and results in higher average wages and purchasing power. As this review will discuss, scholars have deepened and complicated this framework in recent years, but a unifying theme is that technology is closely linked to productivity growth.

Aggregate productivity growth has historically led to wage growth, but there are theoretical reasons why this may not hold in the future. One possibility is that an increasingly large share of GDP (or productivity gains) could go to capital instead of labor, rewarding investors but not workers. Secondly, even if some share of productivity gains goes to workers, the benefits could be unevenly distributed by level of skill or type of tasks performed. This review will discuss how economists have tried to assess the plausibility of these and related scenarios.

Since technology is so closely related to productivity, the review starts with how economists have interpreted productivity growth trends and how they relate to technological change. In the 18th and 19th centuries, technologies associated with the Industrial Revolution dramatically reduced the costs of producing food, clothing, and other goods—and through recording devices, radio, film, television, planes, and automobiles, the costs of communication and transportation. Gordon (2017) found that the most economically important innovations occurred from 1870 to 1970, a period associated with very rapid growth. Since then, he posited, productivity growth has slowed because digital technologies are fundamentally less economically important than those that preceded them, and indeed productivity growth has slowed across advanced industrial economies since the 1980s. For example, in the United States, productivity grew at a rate of 2.8% on an annual average basis between 1947 and 1973, but since then, it has been much slower, with the exception of the 2000 to 2007 period. From 2007 to 2017, average annual productivity growth was 1.3% (Bureau of Labor Statistics 2019a). Based on these considerations and related analysis, Gordon (2017) concluded that new technologies are having little impact on the economy and hence the labor market.

Cowen (2011) has advanced a similar argument that previous technological advances were far more impactful than recent ones. Atkinson and Wu (2017) provided empirical evidence on this point by showing that recent decades have resulted in lower rates of creation and destruction of new occupations relative to previous eras in economic history.

From the point of view of these scholars, the latest wave of advanced technologies (i.e., digital technology, artificial intelligence (AI), and automation) is unlikely to affect labor markets nearly as much as the technological changes of prior generations.

However, other economists and scholars have reached what could be described as the opposite conclusion—arguing that new technologies have already started to profoundly transform the labor market and will likely accelerate in their effects. Klaus Schwab (2016), founder and executive chairman of the World Economic Forum, has gone as far as to label the current period of technological advancement the Fourth Industrial Revolution, emphasizing the rapid pace of change. Consistent with Schwab's (2016) conceptualization, Gill Pratt (2015), who formerly managed a robotics program for Defense Advanced Research Projects Agency, compared the latest wave of technologies to the Industrial Revolution, and wrote: "[T]his time may be different. When robot capabilities evolve very rapidly, robots may displace a much greater proportion of the workforce in a much shorter time than previous waves of technology. Increased robot capabilities will lower the value of human labor in many sectors." Pratt listed several key advances he believes are driving technological changes: growth in computing performance, innovations in computer-aided manufacturing tools, energy storage and efficiency, wireless communications, internet access, and data storage. Brynjolfsson and McAfee (2014) have advanced similar arguments and claimed that information technology inhibited job creation after the Great Recession and is leading to income inequality and reduced labor demand for workers without technical expertise. Responding to arguments from those who see a slowing pace of innovation as the explanation for reducing productivity growth, they state: "We think it's because the pace has sped up so much that it's left a lot of people behind. Many workers, in short, are losing the race against the machine."

Defenders of this high-impact view of technology are left with how to reconcile it with the observed slowdown in productivity growth. Brynjolfsson, Rock, and Syverson acknowledged the lack of strong growth in macroeconomic productivity data but argued that this is a measurement challenge associated with general purpose technologies like AI, which they indicated, require significant complementary investments that are slow to reveal themselves as productivity advances. The argument draws on David (1990), who listed a number of reasons why computer and information-processing technology may require a long lead time before showing up in measured productivity statistics—including diffusion lags, measurement error, the slow depreciation of previous technologies, and the complexity of organizational restructuring.

A related point is found in Bessen (2015), who, in writing about machines of the Industrial Revolution, indicated that new inventions are slow to deploy because they require significant practical refinements and investments in human capital to be competitive with existing practices. However, new analysis of historical data suggests that new technologies, such as software and industrial robots, did have a significant impact on labor markets. Webb (2019) finds, using the text of both job task descriptions and patents to construct a measure of the exposure of tasks to automation, that occupations that were highly exposed to previous automation technologies saw large declines in employment and wages over the relevant periods.

Yet, with respect to computers and other information technology, the most significant advances were arguably made decades ago as Gordon (2017) has argued, and the technologies have already been widely adopted without dramatically affecting productivity trends. Detailed efforts by Byrne, Fernald, and Reinsdorf (2016) to examine the measurement challenges related to new technology found that even aggressive assumptions about missing data or poorly measured gains yield only small changes in productivity growth statistics.

Acemoglu and Restrepo (2019) provide an alternative response to this conundrum that bridges the two points of view. They argue that automation creates two effects that raise the demand for labor: a productivity effect that expands the demand for labor by making labor more efficient in the tasks left to it after automation (e.g., lawyers can compose briefs more efficiently using the internet and writing software) and a reinstatement effect that creates new demand for labor as automation creates new tasks (e.g., developing software, managing information networks). Automation also creates a displacement effect that substitutes for labor (e.g., online airline and hotel booking platforms reduce demand for travel agents). The net effect on labor depends on the combination of these three effects, which vary by technology. For weak technologies that are only modestly productivity-enhancing and don't generate new tasks, the displacement effect can dominate the productivity effect. One example Acemoglu and Restrepo (2019) provide is the automation of call processing centers. Humans tend to be better at solving customer problems than automated voice recognition software but less efficient on net because of higher costs. If this scenario is indicative, automation could erode the labor share of GDP without generating significant economic growth.

Other scholars who characterize the potential of new technologies as high impact sidestep the issue of productivity decline by pointing out that the effects are coming in the near or not-too-distant future. In influential research, Frey and Osborne (2017) predicted that 47% of U.S. workers are at risk of having their jobs computerized. Subsequent studies by Organisation for Economic Co-operation and Development (OECD) economists have put the figure considerably lower using more detailed information on how tasks vary by occupation (Arntz, Gregory, and Zierahn 2016).

Overall, even at the highest conceptual level, economists and scholars remain divided as to the effects of new technology on the macroeconomy and the labor market. Moreover, this literature has not yet been well integrated with the potential effects of other macroeconomic changes such as immigration, trade, and demographic patterns, all of which put additional constraints and pressures on the labor market (Schwab 2016). The remainder of this review seeks to clarify the key issues of the theoretical debate, including how technologies will affect the distribution of income between workers of different occupations or of different education levels, and how technologies will affect the returns to capital compared to labor. The review will also examine efforts to resolve these issues in empirical work and point to the key constructs that will need to be addressed.

2.2 Defining New Technologies

2.2.1 Automation

Automation is the broadest definition of new technologies, and thus will serve as the main focus for this discussion.

Automation is defined as the substitution of non-human value for human production value. Following Acemoglu and Restrepo (2019), we define automation as “development and adoption of new technologies that enable capital to be substituted for labor in a range of tasks” (p.3).

Said otherwise, it refers to any instance where capital replaces labor as the sources of value in the chain of production. The production of any good or service can be defined as the performance of specific tasks, each of which, in theory—though not reality—can be performed by a human or something else—such as a calculator, computer, algorithm, or machine. The automated component is simply the source of value-added that is not performed directly by humans.

To take an intuitive example, a vending machine performs the specific tasks of registering a customer's request for a drink or snack, processing payment, and dispensing the product. It does not, however, grow the ingredients, harvest and process the ingredients, design, manufacture, or market the final product, transport the components or final product, install or replace the product, repair or maintain itself, or protect itself from theft, though those are all valuable tasks in the value-chain for delivering a candy bar to a customer from a vending machine.

This review will discuss automation with respect to the Industrial Revolution's technologies but will focus on technologies introduced in recent decades and associated with computers and computer-controlled machines. AI, machine learning, and digitization are specific manifestations of automation.

2.2.2 Digital Technologies

Digitization refers to the translation of information into a form that can be understood by computer software and transmitted via the internet (Goldfarb and Tucker 2019). The closely related concept of “digitalization” encompasses this meaning of digitization but is used more broadly to refer to the diffusion of digital technologies (technologies that process or transmit digital information) into business operations and the economy (Muro et al. 2017; Charbonneau, Evans, Sarker, and Suchanek 2017). Many of the aforementioned technological changes described by Pratt (2015) are relevant to the growing importance of digitization, especially internet access and speed, wireless communication, processing speed, and data storage efficiency. Taken together, these changes have encouraged automation. Many services once performed by humans—like the trading of financial assets, banking, accounting, processing orders for food and retail goods, the coordination of transportation, creating and confirming reservations at restaurants or accommodations, searching publications and media content, and monitoring energy usage—are now routinely handled by software as a result of digitization.

In Goldfarb and Tucker’s (2019) review of the literature on how digital technologies shape economic activity, they conclude that digital technologies lower the costs of five important economic activities: 1) search for specialized labor or products, 2) replication, reproduction, and copying, 3) acquiring or sharing goods and information, 4) “tracking” or identifying individuals and their preferences, and 5) verification or assessing the quality of products and services. These activities could be broadly classified as the transmission of information.

When applied to the labor market, their review suggests a number of consequences, which will be addressed in more detail subsequently. The effect on search costs should facilitate the specialization of labor and allow for specialized niche producers; the lowering of reproduction costs creates ambiguous effects for individuals and companies that own digitized intellectual property (e.g., video, literary, music). As long as they can enforce their ownership, it may increase profit margins because they can scale up distribution with relatively little cost; however, using that distribution channel also opens them up to more intense competition. The falling cost of obtaining and sharing information could increase the value of workers who deal in specialized knowledge or subject matter expertise, as it would allow them to acquire data and interpret data more readily, via consulting services, for example. The falling cost of search, information, tracking, and verification has already increased the demand for application developers, website developers, and computer programmers whose skills are needed to allow individuals and companies to participate in the growing digital economy and take advantage of the trends mentioned.

2.2.3 Artificial Intelligence

There is no consensus in the economics or computer science literature as to what is exactly meant by “artificial intelligence.” From our perspective, AI is a subset of automation technologies that is distinct from digitalization but may utilize digitization technologies. We define AI as the automation of cognitive tasks that are part of the production chain; we distinguish cognitive tasks from those that involve only the manipulation of physical objects. In this sense, trucking services do not presently use AI in the truck itself, but a driverless truck would, since driving is a cognitive task whereas transportation itself is not. This is consistent with a more general definition offered by an expert committee on AI, which defined it as “that activity devoted to making machines intelligent” (Stone et al. 2016, p.12).

The expansion of automation into cognitive tasks through AI raises questions about the feasibility of automating tasks that were previously thought to be non-automatable. Famous recent examples of AI include IBM’s creation of Deep Blue and Watson, algorithms that defeated champions in chess and the game show Jeopardy, respectively. Other examples include the development of AlphaGo by DeepMind, a Google-acquired company, which created an algorithm that could defeat the leading Go champion, and Siri, acquired by Apple, which developed sophisticated voice-recognition software.

Agrawal, Gans, and Goldfarb (2019) examine the economics literature on AI and argue that the important advances in AI could be characterized as machine learning and essentially involve prediction. The argument is that the recent advances in AI have been mostly limited to machine learnings, and machine learning entails prediction (as in estimating an outcome based on measurable variables or features), using techniques such as maximum likelihood, neural networks, and reinforcement learning. Agrawal, Gans, and Goldfarb (2019) indicate that the economic effects of machine learning on labor are varied, complex, and not well understood at this point. Some predictions may replace human decision-making and labor, but many others complement human labor and make it more efficient. They summarize their conclusion by stating:

“Overall, we cannot assess the net effect of artificial intelligence on labor, even in the short run. Instead, most applications of artificial intelligence have multiple forces that impact jobs, both increasing and decreasing the demand for labor. The net effect is an empirical question and will vary across applications and industries” (Agrawal, Gans, and Goldfarb 2019, p.34).

Other scholars have a more expansive interpretation of what has been accomplished in AI (Brynjolfsson and McAfee 2014; Ford 2015; Stone et al. 2016). A committee of inter-disciplinary scholars, including economist Erik Brynjolfsson, have written more broadly about the tasks that can be performed by AI, including object and activity recognition, language translation, and robotics (Stone et al. 2016). They claim that AI has mostly affected workers performing routine tasks in the middle of the wage distribution, but that AI is likely to gradually automate cognitive tasks that are not currently considered routine, which may include, for example, legal case searches performed by first-year lawyers. In summary, they speculate that “AI is poised to replace people in certain kinds of jobs ... However, in many realms, AI will likely replace tasks rather than jobs in the near term, and will also create new jobs ... AI will also lower the costs of many goods and services, effectively making everyone better off” (Stone et al. 2016 p.8).

2.3 Major Economic Theories on New Technologies

Modern macroeconomic growth theory generally posits a positive relationship between productivity and wage growth. The underlying implication is that technological change—which is generally regarded as the fundamental cause of growth in neoclassical models—should improve the real incomes of workers. This section will review historic patterns and theoretical conditions when this basic prediction may hold or fail, depending on factors like the extent to which new technology displaces work and rewards capital over labor or is skill-biased. After briefly discussing the Industrial Revolution, we discuss the literature on skill-biased technological change, followed by the introduction of task-based frameworks, and the application of task-based frameworks to automation.

2.3.1 Lessons From the Industrial Revolutions

In the midst of the Industrial Revolution, Karl Marx (1867) famously stated that the accumulation of capital led to the impoverishment of laborers. His argument was that investment in machines was both labor-displacing and labor-complementing. He believed business owners invest in labor-saving machines when wages get too high, thus creating a “reserve army of labor” that would bid wages back down. Yet, the expansion of production also required workers. As he worded it: “Capital works on both sides at the same time. If its accumulation, on the one hand, increases the demand for labour, it increases on the other the supply of labourers by the ‘setting free’ of them” (Marx 1867, sect. 3, last para.).

Economic historians have since rejected Marx’s prediction that the real wages of workers would remain stagnant in modern economies. As Keynes (1978) predicted, living standards have increased considerably and unemployment resulting from technological process proved to be only temporary. There is no debate among economists that living standards are dramatically higher now than in the 19th century in rich countries. Not only is the purchasing power of income orders of magnitude higher, but ordinary people, workers, and business owners enjoy far greater health and longevity (Deaton 2016). Marx’s predictions were also strikingly inaccurate even during his own era. Data from Gregory Clark’s (2005) research on the Industrial Revolution shows that the wages of workers rose rapidly in England. In fact, from 1850 to 1900, real wages of building workers doubled in England as capital accumulation and education increased.

Assessing the Impact of New Technologies on the Labor Market

More detailed accounts of specific sectors on the cutting edge of new technologies reveal similar dynamics of rising wages and living standards for workers, as new technologies diffused. Economic historian James Bessen calculated the real hourly wages for weavers and spinners, roles that were using cutting-edge technologies in factory settings. From 1830 to 1860, these wages remained relatively stagnant, but grew rapidly from 1860 to 1890. Bessen's (2015) explanation was that labor markets were relatively uncompetitive during the earlier phase, and workers had fewer alternative sources of employment (consistent with Marx's perspective), but as technological change expanded economic growth and created new sources of employment, even workers with modest skills, like spinners, saw their wages increase, and those with more specialized technical skills—weavers—benefited disproportionately.

Aside from average wage patterns, economists are also interested in understanding the consequences of technological innovation on the income distribution. Income inequality fell dramatically for England following the Industrial Revolution, as documented by Clark (2008) and Lindert (1986). In the U.S., Lindert and Williamson (2016) found that income inequality rose for much of the 19th century (from 1800 to 1860) from a low start, plateaued until around 1910 and declined sharply thereafter until the 1970s. This is consistent with evidence from Goldin and Katz (2010) that the wages of high-school educated workers grew more rapidly than the wages of college-educated workers from 1915 to 1980. Piketty, Saez, and Zucman (2017) found the same broad reduction in income inequality as measured by the share of national income held by the top 1% of earners, which fell from 20% to 10% from 1930 to 1980 (World Inequality Database). This was a period of rapid innovation and productivity growth. A major focus of the economics literature in recent years has been devoted to explaining why income inequality started rising again around 1980.

Economic historians have also examined and debated to what extent the technologies of the first and second waves of the Industrial Revolution could be regarded as leading to an increase or decrease in the demand for skills. More formally, scholars have examined whether or not technology is complementary with skilled labor.

Within manufacturing and production work, Attewell (1992) reviewed the literature on the “deskilling” and skill upgrading effects of technology and concluded that the evidence of systematic deskilling is mixed, though there are many specific examples of it occurring. Attewell (1992) found thin evidence for one interpretation of the deskilling theory: investment in machines in order to reduce reliance on skilled workers. Instead, he refers to many instances of workers developing advanced skills in relation to the use and maintenance of the new machines. Consequently, with respect to the early to mid-20th century, Attewell (1992) noted many examples of new skills emerging in response to technological demands placed on production workers, even when the technologies performed tasks that required high levels of training and skills. Across the economy at large, Attewell (1992) noted a stronger consensus across studies that technology led to a higher demand for skills, as shown, for example, by the secular rise in educational attainment.


Some scholars, such as Goldin and Sokoloff (1982) and Humphries (2013) interpreted the heavy reliance of women and children during the mid-19th century as evidence that industrial machinery disproportionately increased demand for low-skilled entry-level work and was thus deskilling. Nuvolari (2002) identified the deskilling effects of technology as one of the motivations for political resistance to factory work in England during the 18th and 19th centuries. Likewise, De Plejit and Weisdorf (2017) identified a large shift in the share of occupations from semi-skilled to unskilled (using early 20th century categorizations), with only a small shift toward high-skilled workers. By the early 20th century, however, technology and skill became complementary, according to Goldin and Katz (1998), who found a positive relationship between the skill level of blue-collar workers and investment in capital per worker from 1909 to 1940. Their interpretation of the evidence posits the early 20th century as a turning point for skill-technology complementarity.

One complication to this debate is how skill is defined. Disputing the deskilling literature, Bessen (2015) used worker-level microdata to show that employers invested in training for female workers during the 19th century, and they became far more productive with experience. Moreover, as production became more mechanized, the remaining non-automated tasks demanded higher rather than lower levels of skill, as evidenced by the rising positive relationship between on-the-job learning and worker productivity. Bessen (2011) interpreted this and related evidence as indicating that early industrial technologies required and rewarded high levels of skill that could be acquired on-the-job but did not require formal education.

2.3.2 Skill-Biased Technological Change

Well-established literature in labor economics holds that technology has increased the productivity of workers with college education more so than workers with less education. This fact explains the rise of earnings for college-educated workers relative to the earnings of non-college-educated workers, despite the increase in the labor supply of college-educated workers.

The basic empirical facts for this theory are illustrated in Figure 2.1. Starting around 1980, the relative incomes of college-educated workers have increased relative to workers with a high school education, adjusting for other obvious observable factors. This is known as the college earnings premium, and it has increased from 34% in 1980 to 68% 2018. A surprising aspect of this rising premium is that the share of hours worked by college-educated workers has nearly doubled from 20% in 1979 to 39% in 2018. In a simple supply-demand framework, this suggests that demand for college-educated workers has outpaced the steady increase in supply.



Source and Note: Author analysis of IPUMS-CPS. Premium is calculated by year-specific regression of log of total personal income on a categorical variable for having at least four years of college relative to high school (which is the reference category), a binary variable for less than high school, a binary variable for some college but less than four years, gender, the number of hours worked last week, 10 age categories in roughly five-year bands, with workers aged 20-24 in the reference group. Experience premium is coefficient on age category of workers aged 50-54 (the highest across all years) relative to workers aged 20-24. Population is restricted to workers under the age of 65. The share of national hours worked by college-educated workers is the sum of hours worked per year (assuming hours worked per week is consistent across 52 weeks of the year) for college-educated workers divided by the total number of hours worked.

This literature has a long history, with many scholars contributing to theoretical and empirical studies. Among the seminal papers on this is work from Katz and Murphy (1992), who found that the rising college earnings premium could be linked to evidence of rising relative demand within industries for college-educated workers. In explaining the college premium over an even longer period, Goldin and Katz (2010) emphasized a slowdown in the growth rate of the supply of college-educated workers, in the midst of stronger increases in demand, as being primarily responsible for the rising premium since 1980.

Other research makes a more direct link between new technologies and changes in the wage structure. Krueger (1993) found that use of computers was associated with more highly educated workers and predicted higher earnings levels and growth. Closely related work by Autor, Katz, and Krueger (1998) found further evidence that industries that have invested more heavily in computers exhibited a sharper rise in demand for educated workers. Michaels, Natraj, and Van Reenen (2014) further demonstrated that investment in information and computer technology and technology-related research and development can account for up to a quarter of the growth in the college earnings premium and demand for educated labor.

While the theory has been influential, not all economists agree with this interpretation of the college premium or its larger implication in establishing skill-biased technological change. DiNardo and Pischke (1997) ran what might be considered a placebo test using a similar analysis as that performed by Krueger (1993). They found that the labor market returns to pencil use show similar properties as the return to computer use, despite the fact that nearly every worker could, in principle, use a pencil. They caution against interpreting the computer wage premium as suggesting that computer skills specifically have led to higher wages. Rather, they suggest that selection bias—more skilled workers may be more likely to use both a computer and a pencil—is a likely possibility. If so, skill-biased technological change may simply be observing a rising relative demand for education that is unrelated to technology—or at least any direct effect of technology on the productivity of college-educated workers.

Card and DiNardo (2002) also conveyed skepticism of the skill-biased technological change literature, pointing out that the framework relies entirely on supply, demand, and technology, ignoring well-established sources of variation in earnings, such as unions, efficiency wage premiums, minimum wage laws, and economic rents. Empirically, the theory also runs into various problems, they argued, in that wage inequality narrowed during the 1990s even as information technology deepened in its diffusion and adoption by every conventional measure. They also reviewed complications in applying the skill-biased technology framework to patterns within age, race, and gender groups, where the college wage premium plays out differently in ways that are inconsistent with computer or technology use. Lemieux (2006) also documented empirical challenges to the framework and related rising variance in wages to demographic changes in the workforce.

Bessen's (2015) work on the Industrial Revolution analyzed the complex relationship between technology, formal education, and skill (which comes from a mix of formal education, training, and experience). In Bessen's account, new technologies with broad applications—like information technology but also many of the machines of the Industrial Revolution—create demand for workers with high levels of cognitive ability, as measured by their ability to rapidly learn and master new skills. Thus, the earliest workers using weaving machines were far more literate (or educated) than the general population, even though literacy was not needed to use the machine. This relationship creates an education premium, but ultimately, specific technical skills—not the sort of education acquired through college—are needed to adopt and make use of new technologies. When the technical skills became standardized, formal education became less valuable, and factory workers gradually became less literate. As Bessen wrote:

“Thus, while demand for college graduates has grown in relative terms, it appears to be mainly because college-educated workers are better at learning new, unstandardized skills on the job, not because their college education conferred specific technical skills ... As technical knowledge in these occupations becomes increasingly standardized, more and more workers will be able to acquire the needed skills without a college education, in formal training provided by employers or vocational and technical skills” (Bessen 2015, p.145-146).

Despite the decline in formal education (literacy), Bessen (2015) found that skill was highly rewarded during the industrial revolution. Weavers were paid based on how many pieces of cloth they produced, and employers and employees invested heavily in on-the-job-training. He found that more experienced (and hence skilled) workers produced substantially more cloth per hour.

Consistent with Bessen's (2015) theory that technological change tends to reward informal skills, the premium for experience has also increased since 1980 and, while difficult to compare, is larger than the college premium as measured here. Figure 2.1 measures the experience premium as the difference in log income between workers aged 50-54 compared to workers aged 20-24. That premium increased from roughly 67% in 1980 to 91% in 2018.

The results of a firm-level study from Bresnahan, Brynjolfsson, and Hitt (2002) are relevant to Bessen's (2015) interpretation. They found evidence that firm investments in information technology increase worker autonomy and the use of smaller teams, which also coincided with higher demand for skilled workers. In this way, the adoption of technology has rewarded skilled workers, individuals whom managers trust to be more autonomous.

From another point of view, the evolution of the income distribution in the U.S. has raised theoretical concerns. In light of work by Thomas Piketty and collaborators (2015), the rise of income shares going to the top 1% or top 10% of income earners has been difficult to reconcile with a skill-biased technological change framework. In an era when roughly one-third of workers have a college degree, it is not obvious why most of the gains would have gone to the top 1% to 10%. One potential solution is to appeal to the increasing importance of capital income, which may accrue to the owners of technology-producing companies, but that would not account for the fact that labor income (rather than capital income) drove rising levels of inequality from 1980 to 2000, as measured by top 1% shares.

Consistent with the basic tenets of capital-skill complementarity and skill-biased technological change, a skill premium, whereby advancing technology favors relative demand for highly skilled workers, has been evidenced as a widespread global trend—from India (Berman, Somanathan, and Tan 2003) to Hungary (Kezdi 2002), in low- and middle-income countries (Conte and Vivarelli 2011), and in a range of both developed and developing countries (Burnstein and Vogel 2010). However, technology advancement is not the only relevant factor: Imported skill-enhancing technology (Conte and Vivarelli 2011), foreign investment (Kezdi 2002), purchase of foreign machinery for technological progress (Gorg and Strobl 2002), and trade and multinational production (Burnstein and Vogel 2010) all appear to contribute to the skill premium. For example, Burnstein and Vogel (2010) identified that U.S. trade and multinational production between 1966 and 2006 accounts for one-ninth of the 24% rise in the U.S. skill premium. In their model, multinational production appears to be more influential: The rise in the skill premium caused by trade is 1.8% in the U.S. (and other skill-abundant countries) and 2.9% in skill-scarce countries, whereas, considering multinational production, the rise in the skill premium is 4.8% in the U.S. (and other skill-abundant countries) and 6.5% in the skill-scarce countries.

2.3.3 Task-Based Models

In a review and critique of the skill-biased technological change literature, Acemoglu and Autor (2011) described two important features of these models. One is that skills are typically thought of (or at least measured) as entirely determined by college education. Second, technology is modeled as factor augmenting—meaning it makes skilled workers more productive but does not replace labor. Drawing further implications, the canonical model suggests technological change increases the wages of high- and low-skilled groups. Thus, the skill framework cannot explain why wages for low-skilled workers might experience real decline, unless relative supply increases. Since the relative supply of low-skilled workers has fallen (see Figure 2.1), this framework does not explain why real wages for low-skilled workers have declined or been stagnant, depending on the measures used.

Acemoglu and Autor (2011) proposed a task-based framework where tasks are defined as units of work activity that produce output and skills refer to the ability to perform tasks, which may or may not be related to college education. In this view, occupations are “bundles of tasks.” Looking at it from this richer framework confers several advantages, including the ability to understand occupational growth patterns that differ within education levels and the capacity to model machines as potential substitutes for work—rather than only labor-augmenting. Still, this framework does not imply that machines always replace human workers. Rather, a machine may perform some aspect of an occupation—a specific task—while the human worker performs the other tasks associated with the occupation.

Several papers have used some version of this task-based framework to investigate how technology has affected the labor market. Autor, Levy, and Murnane (2003) argued that computers are especially efficient relative to labor at routine tasks, and the falling price of computers has lowered wages and demand for routine human labor while increasing relative compensation of non-routine labor. Autor and Price (2015) extended the earlier analysis and found additional evidence that the task content of U.S. labor has become less routine (measured as more analytical and interpersonal) from 1980 to 2009, whereas the largest losses in employment shares have come from routine cognitive work and manual labor. Atalay et al.’s (2017) findings confirmed the broad patterns of shifts toward non-routine labor using a database of job postings that they constructed from 1960 to 2000. Spitz-Oener (2006) analyzed how the task content of occupations has changed in western Germany and found a large shift within occupations toward non-routine tasks and away from routine tasks. This change, she stated, explains a larger fraction of the shift toward more educated workers.

In thematically related work, Autor and Dorn (2013) ranked occupations by wages—which they used as a proxy for skill—and found that occupational growth was faster from 1980 to 2005 for high- and low-paying occupations relative to “middle-skilled” occupations, which would include many production occupations within the manufacturing sector and clerical jobs across a wide range of industries. This is referred to as job polarization. They posited that computers have displaced routine manual and cognitive tasks, but not low-skilled services, which are often not routine. Other studies have found evidence consistent with this pattern for the U.S. (e.g., Autor, Katz, and Kearney 2006; Autor, Katz, and Kearney 2008; Autor and Dorn 2013; Holzer 2015). Goos and Manning (2007) and Goos, Manning, and Salomons (2014) found a similar polarizing pattern of job growth by occupation for Britain (from 1979 to 1999). Katz and Margo (2014) found evidence that the “hollowing out” pattern also described the U.S. labor market from 1850 to 1910.

Growth projection estimates suggest that these patterns may continue. Bughin et al. (2018) took a task-based approach to project skill-demand shifts in hours worked per week from 2016 to 2030. According to their estimates, there will be an 11% decline for physical and manual skills and a 14% decline for basic cognitive skills (e.g., basic literacy). In contrast, they project a 9% increase for higher cognitive skills (e.g., creativity, problem-solving), a 26% increase for social and emotional skills, and a 60% increase in technological skills.

Similarly, an analysis of O*NET occupation skills and work activities for the 20 occupations projected to grow fastest between 2016 to 2030 (Bureau of Labor Statistics 2019b) revealed that 95% require advanced cognitive skills and 85% require socioemotional skills, whereas 65% require basic cognitive skills and only 15% require manual labor skills (Maese 2019). Among these fast-growing jobs, there are differences in skill demand based on median wage (using 2018 median wage): High-paying jobs are more likely to require advanced technical skills (e.g., data mining, network monitoring), to involve technological work tasks (e.g., analyzing data) and advanced cognitive work tasks (e.g., making decisions and solving problems, thinking creatively), and are less likely to require manual labor tasks (e.g., repairing and maintaining mechanical equipment) (Maese 2019).

2.3.4 Task-Based Models With Automation

The aforementioned papers on skill-biased technological change and task-based literature have either assumed technology increases labor demand or only threatens workers who perform routine tasks. More recent research broadens the role of technology to include the ability to perform any task; this meets the definition of automation described in Section 2.2.1.

Assessing the Impact of New Technologies on the Labor Market

The model introduced by Acemoglu and Restrepo (2019) describes three classes of technology: automation, new task generation, and factor-augmenting technologies (which increase the productivity of labor or capital in doing any task). A new technology may contain all or multiple aspects of these effects. An industrial machine might automate some assembly tasks performed by humans but create demand for new tasks in programming, installation, maintenance, and repair. It is less realistic to imagine technologies that make labor or capital better at any task, as the authors point out.

Acemoglu and Restrepo's (2019) framework classifies the effects of these groups of technologies. When a technology automates tasks, it creates demand for labor through a productivity effect, reduces demand for labor through a displacement effect, and has ambiguous effects, depending on how it changes the composition of work done by industry. A technology might also increase the number of tasks performed in the economy, which reinstates demand for labor and creates a productivity effect.

In the empirical section of the paper, Acemoglu and Restrepo (2019) examine trends in U.S. data and distinguish between 1947 to 1987 and 1987 to 2017. In the earlier period, they measure a displacement effect from new technologies that amounted to 0.48% per year, which was offset by a reinstatement effect and strong productivity growth (2.4% per year). The net result was rising real wages (2.5% per year) and strong labor demand. In the period since 1987, wage growth has been much weaker (1.3% per year) as a result of weaker productivity growth (1.5% per year), a slowdown of the reinstatement effect (from 0.47% to 0.35% per year), and an acceleration of the displacement effect (from 0.48% to 0.70%). Using industry-year variation within the U.S., they find that the proxy measures for the use of automation and reliance on routine tasks within an industry predict larger displacement effects and smaller reinstatement effects. However, they also find that industries that rely more heavily on new occupations or occupations with new tasks have larger reinstatement effects.

These results are consistent with earlier empirical work from Acemoglu and Restrepo (2017) on industrial robots, one specific form of automation technology. Using data on robots by industry for the U.S.—while identifying variation using European-industry trends to reduce reverse causality—they found that labor force participation fell in the commuting zones most exposed to robots, where exposure runs via initial employment at the regional and industry level multiplied by an index of industry-specific increases in robots per worker.

Using similar data but a different modeling and empirical strategy, Borjas and Freeman (2019) contrast the labor competition effects of robots to the effects of immigration and find that robots have a greater impact. Specifically, they find that industries with robots displaced two to three workers per robot or three to four workers in the most exposed groups, such as workers who either have low levels of education or perform highly automated tasks (as defined by workers through their responses to an O*NET item on work context: “How automated is your current job?”). They do not find negative wage or labor displacement effects for college-educated workers or workers in jobs that are not automated. They suggest these effects have historically been small from a macroeconomic perspective because the number of robots is small, but the effects could become macroeconomically significant if robots become much more widely adopted.

Going beyond the U.S., other cross-country empirical work suggests the productivity and reinstatement effects have substantially outweighed the displacement effect—at least for industrial robots. Graetz and Michaels (2018) acquired data on the purchase of industrial robots by country and industry and conducted an analysis across 17 countries from 1993 to 2007. They modeled robots as perfect substitutes for certain human tasks and assumed companies adopt robots when the profits from doing so exceed the cost of purchasing the robots. Their empirical analysis concluded that the adoption of robots increased GDP per hour worked (or productivity) with no effect on labor demand in the affected industries. Presumably, labor demand would have increased in other industries. In other words, industries operating in countries that were especially prone to adopt robots did not experience job growth that was any different than job growth in industries and countries with low adoption rates. Graetz and Michaels (2018) found that robot adoption predicts wage growth and lower prices for consumers, but employment shifts from low-skilled workers to middle- and higher-skilled workers. They used several techniques to verify whether their analysis could be interpreted as a causal effect and found evidence that it is.

Caselli and Manning (2019) introduce an alternative theoretical model that also draws on a task-based framework and defines technology broadly to be any capital investment that reduces the direct or indirect costs of something purchased by consumers. They then lay out a series of parsimonious assumptions and work out the logical outcomes with respect to effects on average wages. They assume interest rates are not affected by technology, so that the supply of capital is not constrained. Next, they distinguish between investment goods and consumer goods. They reason that if the price of investment goods (e.g. machines) falls relative to consumer and intermediate goods, workers must benefit, though not necessarily all, and the returns to investment capital will fall (though not necessarily the capital-labor ratio). When they further assume that workers can seamlessly switch occupations and retrain, they reason that all workers stand to gain from technological change. In reality, workers typically face a modest wage penalty after experiencing a layoff even six years later, suggesting that transitions are not seamless (Couch and Placzek 2010).

Still, Caselli and Manning's (2019) analysis suggests that most plausible scenarios involving technological change will result in benefits to most workers. Yet, historical data analyzed by Webb (2019) indicates that occupations that were highly exposed to previous automation technologies experienced large declines in employment and wages. This suggests that AI, which the author finds is directed at high-skill tasks, may lead to the long-term substitution of high-skilled workers in the future.

2.4. Empirical Trends in the Labor Market

2.4.1 Broad Historical Patterns in the Structure of Work

The theoretical work described above identifies how economists believe technology is affecting labor markets, usually after attempting to isolate technological effects from other factors. However, regardless of the effect technology has had on the labor market, readers may want a broader sense of long-term labor market trends, irrespective of the underlying causal mechanisms.

The Industrial Revolution and subsequent era of high productivity growth coincided with a major transformation of work in the U.S. In 1850, roughly half of workers were classified into farming or related agricultural occupations. By 1970, when Robert Gordon (2017) located the end of an economic revolution, the share of workers in farming occupations had fallen to just 4%. These data are shown in Figures 2.2A-2.2F. Farming jobs were largely replaced with work in professional occupations, non-professional service occupations, and clerical services. Blue collar work peaked as a share of total employment around the middle of the 20th century and saw large losses—as a share of total employment—before the introduction of information technology. Since 1980, almost all of the net changes have been in professional services, with small gains from non-professional services. Consistent with the task-based framework of Acemoglu and Autor (2011), clerical occupations, which are typically classified as routine and automatable, peaked as a share of total employment in 1980 and have declined steadily with the spread of information technology. Professional service occupations, meanwhile, are classified as non-routine and cognitively demanding, and therefore most likely to be resistant to displacement by automation.

Economists at the Bureau of Labor Statistics have summarized related trends from 1910 to 2000 (Wyatt and Hecker 2006). In addition to the patterns described above, they noted the massive rise of jobs in healthcare occupations from 1910 (0.4 million jobs) to 2000 (9.1 million jobs), and the large-scale disappearance of private household jobs from 2.3 million in 1910 to 0.5 million in 1990. Similarly, Pilot (1999) documented the emergence of novel occupational categories, as well as the redescription, and disappearance of others from 1948 to 1998. Overall, 52 of the 209 occupational categories listed in 1948 were listed with the same detail in 1998 and an additional 78 with a change in detail or description of the category. Finally, 79 occupations were not listed in 1998 (including “adding machine servicemen” and “blacksmiths”). Analyzing data by industry, Fuchs (1980) described the broad pattern of transition away from agriculture toward manufacturing and services. He cited and largely agreed with earlier theoretical views that the rise of services follows from a higher income elasticity of demand for services relative to goods (meaning an extra dollar of income translates more readily into service demand rather than goods demand). Fuchs (1980) accepted this but indicated that the higher productivity of goods is also a factor. In Acemoglu and Restrepo's (2019) framework, economic growth, and the higher productivity of goods-producing sectors create a productivity effect and a reinstatement of labor effect outside of the goods-producing industries.

Figure 2.2A Employment by Occupation as a Share of Total U.S. Employment, 1850-2017
(Professionals and Managers)




Figure 2.2B Employment by Occupation as a Share of Total U.S. Employment, 1850-2017
(Sales and Service Providers)




Figure 2.2C Employment by Occupation as a Share of Total U.S. Employment, 1850-2017
(Clerical)




Figure 2.2D Employment by Occupation as a Share of Total U.S. Employment, 1850-2017
(Operatives and Laborers)




Figure 2.2E Employment by Occupation as a Share of Total U.S. Employment, 1850-2017
(Craft and Production)




Figure 2.2F Employment by Occupation as a Share of Total U.S. Employment, 1850-2017
(Agricultural)



Source: Author analysis of IPUMS USA, occ1950 variable, which is a consistent classification of occupations across all census years. Those occupations were reclassified into the higher-level categories shown above by the authors.

2.4.2 Demand for Labor in Technology-Producing Industries

Even as manufacturing employment has fallen in relative importance in the U.S., “advanced industries” have continued to provide a stable source of employment and a growing source of economic output. As defined by scholars at the Brookings Institution (Muro et al. 2015), advanced industries are defined by their high levels of investment in R&D and their propensity to employ workers in science, technology, engineering, or mathematical (STEM) occupations (Muro et al. 2015). These industries are responsible for the creation of most of the new technologies in the U.S., as measured by patents, as well as all commercial information technologies, since the industries that produce these technologies (e.g., computer manufacturing, software, information, telecommunications, and computer services) are explicitly included.

These 50 industries are found in the manufacturing, energy, and service sectors, but the service side has accounted for all the net employment growth while employment in advanced manufacturing has declined. Together, advanced industries accounted for 8.7% of total employment in 2015, with 60% of those in services. In fact, four service industries account for the largest share of employment and over one-quarter of the total: computer systems design; architecture and engineering; management, scientific, and technical consulting; and scientific research and development (Muro et al. 2015).

The educational requirements to work in these industries are relatively high, often STEM-focused, and increasing. In 1968, 76% of workers in advanced industries never attended college. By 2013, that share fell to 25%. At the same time, all of the job losses for those without college have been in the high school or less category. The share of advanced industry workers with some college or an associate degree has slightly increased in recent decades and stands at 25% (Muro et al. 2015). Along these lines, Rothwell (2013) finds that roughly half of workers who are in STEM-intensive occupations have less than a bachelor's degree, but often some form of postsecondary training or education. Across all levels of education, Muro et al. (2015) document a substantial wage premium for workers in advanced industries, possibly as a result of the value of intellectual property in these industries.

Projecting forward, these results suggest that demand for tertiary education will continue to dominate employment opportunities in technology-producing industries. However, a sizable and perhaps even growing share of jobs will be available for those with technical certifications or other sub-bachelor's level credentials, some of which could be provided through high school education (Krigman 2014).

More recent work from Muro et al. (2017) found that demand for workers with high levels of skill in digital technologies has increased considerably in recent years, and the share of all workers in these occupations increased from 4.8% in 2002 to 23% in 2016. Moreover, they found that even occupations with medium or low digitalization skill requirements became more digital intensive from 2002 to 2016, suggesting that the task content of occupations has shifted broadly toward the need for increased digital literacy. Along these lines, the U.S. Department of Education, the U.S. Bureau of Labor Statistics, and scholars all predict further growth in STEM jobs (U.S. Department of Education 2016; Fayer, Lacey, and Watson 2017; Wisskirchen et al. 2017).

Within technology-producing firms, there may also emerge a demand for entirely new professions. Specifically, new work will require professionals capable of training AI systems to perform intelligent tasks, such as teaching natural language processors and language translators, teaching customer service chatbots to mimic and detect the subtleties and complexities of human communication, and teaching AI systems (e.g., Siri and Alexa) to show compassion and to understand humor and sarcasm. There will also be a need for professionals who maintain and sustain AI systems to ensure that they are operating as intended and to address unintended consequences. In fact, this need exists today; currently, less than one-third of companies are confident in the fairness and transparency of their AI systems and less than half are confident in the safety of their systems (Smith and Anderson 2017 and Wilson, Daugherty, and Morini-Bianzino 2017). Also in demand will be workers who can bridge the gap between high-tech professionals, and the technologies they create, and businesspeople and consumers to help explain and provide clarity about AI systems. As implied by the European Union's General Data Protection Regulation, companies will need data experts, such as Data Protection Officers, to protect privacy rights and related issues; in other words, companies will require professionals who can communicate technical details to non-technical professionals and consumers (Wilson, Daugherty, and Morini-Bianzino 2017).

2.4.3 Demand for Labor in Technology-Using Firms

There is a dearth of comprehensive analyses relating the adoption of new technologies to the demand for labor overall and by skill level. At the macroeconomic scale, De Long and Summers (1991) have found that investment in equipment—which captures the technologies used in the production of goods and services—increases the growth rate of GDP. Those results could be thought of as general effects that go beyond the firms that create or use technology to the economy as a whole, but they do not directly relate to the labor market.

Whether at the macroeconomic or industry level, a principal methodological difficulty in measuring the causal effect of technology adoption stems from identifying whether technological adoption is a driving force of outcomes or a response to a complex and potentially unobservable set of temporal, country, industry, or firm characteristics. A few papers have made explicit attempts to overcome these challenges by using firm-level data.

In quasi-experimental research, Gaggl and Wright (2017) tested the causal effect of information technology investment on labor demand by taking advantage of a tax incentive offered by the United Kingdom government between 2000 and 2004, which allowed small businesses to deduct information technology investment expenses from their tax bills. Treatment effects were identified by using the eligibility cutoffs in a regression-discontinuity design. The causal effect of the investment was to significantly increase employment, wages, and productivity. A decomposition of the employment effect showed a modest decrease in routine cognitive workers (in administrative positions), a sharp increase in non-routine cognitive workers, and no change in manual workers.

Harrigan, Reshef, and Toubal (2016) also addressed concerns about causality in the adoption of technology. They used historical occupational data on the presence of “techies” within firms to predict future adoption of technology and identified the causal effect on job polarization and job growth. They defined “techies” as a set of workers in occupations that involve the installation, management, maintenance, and support of information and communications technology. These workers are mostly “in-house” and not brought in as consultants; it is difficult for firms to scale up IT use without them. The result of this research is that IT use in France predicts skill upgrading—that is, a higher percentage of managerial and professional workers relative to lower-paid workers.

Cortes and Salvatori (2015) likewise used unusually detailed data at the firm level in the United Kingdom and found that the adoption of new technology, as reported by firm managers, was correlated with employment growth from 1998 to 2011, but did not predict a loss of routine work. They did not attempt to address concerns about the endogenous adoption of technology.

The papers described above are limited to non-industrial technologies. The results from papers using industrial machines are more mixed (Section 2.3.4), with some finding negative effects on labor demand and others finding no effect (Acemoglu and Restrepo 2017; Borjas and Freeman 2019; Graetz and Michaels 2018). The consensus among these papers, however, is that industrial machines have displaced lower-educated workers. Though a methodological limitation is each of those papers relies on country- or industry-level data, thus, the analysis may be biased by firm-level characteristics and difficult to relate to macroeconomic patterns that may affect firms differently for a variety of reasons.

Business survey evidence also suggests a positive relationship between labor demand and the adoption of new technology. Bughin et al. (2018) surveyed executives from large organizations and found that only 6% expect their workforce in the U.S. and Europe to shrink as a result of automation and AI. In fact, 17% expect their workforce to grow. Companies that see themselves as more extensive adopters of technology were somewhat more likely than those who described themselves as early adopters to project employment growth over the next three years. Qualitative evidence from industrial technology-using firms in the research of Kianian, Tavassoli, and Larsson (2015) also pointed to a positive relationship between labor demand and novel technologies.

Across a wide range of industries, technology also expands access to work through internet-enabled remote work, thus expanding the labor market and employment to those who may need flexibility or opportunities to work from home, work flexible hours, or engage in other alternative arrangements, such as opportunities for women, youth, older workers, and disabled workers (Millington 2017; World Bank Group 2016). Cramer and Krueger (2016) found that Uber’s mobile platform has increased the efficiency of the taxi service industry, thereby increasing the demand for drivers’ labor.

2.4.4 How New Technologies Are Affecting the Labor Share of Income

As defined above, automation is capital that replaces tasks performed by labor. As such, the adoption of automating technologies across the economy could affect the distribution of income between labor and capital. A change in national income in favor of capital would disproportionately reward individuals who earn income from investment and business ownership. A number of factors will moderate whether labor or capital rises as a share of national income in response to technological adoption: the absolute and relative price of capital compared to labor, how technologies affect labor productivity, the types of tasks performed by technology, and the adaptability of labor to perform new tasks.

Autor and Salomons (2018) used industry-level data across developed countries to estimate the effects of total factor productivity growth—which they liken to the adoption of new technologies—on employment and the labor share of income. They found that the direct effect of total factor productivity growth was a decrease in employment in the industries experiencing productivity growth (a direct effect), but an increase in other industries through an indirect effect. Moreover, they found that productivity growth has coincided with a reduction in the labor share of income in industries experiencing productivity growth.

Autor and Salomons' (2018) finding that highly productive industries have experienced a declining labor share of income is consistent with Elsby, Hobijn, and Şahin's (2014) review of the labor share literature, in that they found changes within industries—particularly a declining share of income going to labor within manufacturing—explained most of the overall decline. Yet, Elsby, Hobijn, and Şahin (2014) rejected the hypothesis that the substitution of capital for labor explains the fall in the labor share, which they argued does not match the trends that theory would predict. Instead, they tentatively suggested that offshoring has played a significant role, rather than technology.

One limitation of Autor and Salomons' (2018) method, as it applies to understanding technology, is that total factor productivity growth can come through channels that are not tied to technology. The quality of human capital, access to international markets, industrial organization, the degree of misallocation, and the institutional context all affect measured total factor productivity growth but may be related to actual technology in complex ways (Jones 2016). On the other hand, an encouraging result of Autor and Salomons (2018) that goes beyond this limitation is that productivity growth—whether it comes from technology or other sources—has not coincided with net losses in employment.

Using a different accounting-based approach, Eden and Gaggl (2018) estimated the information and communications technologies replaced a large number of routine workers (which they defined as the following occupations: sales, office, clerical, administration, production, transportation, construction, and installation, maintenance, and repair) with capital equipment. They also found that information technology rebalanced the labor share of national income in the U.S. toward capital and workers in non-routine occupations.

2.5 What Jobs Are at Risk of Technological Displacement?

A strand of the literature acknowledges that the Industrial Revolution was generally beneficial to workers and living standards. Workers employed in agriculture were able to move into other sectors of the economy as agricultural production and food processing became more mechanized and efficient.

Yet, this history is not guaranteed to repeat itself, especially if newer technologies have fundamentally different characteristics. For example, the power of automation to perform complex cognitive tasks—via AI—distinguishes it from automated technologies of the Industrial Revolution. Likewise, digital technologies can almost instantly transmit data anywhere in the world from any place, which is not something that pre-digital technologies could do. Possible implications involve a reduction of demand for human monitoring and control activities. Along these lines, recently, economists classified driving an automobile to be impervious to automation (Autor, Levy, and Murnane 2003), but self-driving cars are already a reality only a decade and a half later.

Motivated by these and related concerns, economists have made various attempts to study how the adoption of current or future technologies could affect the labor market. In a widely cited paper, Frey and Osborne (2017) provided a review of recent advances in machine learning, AI, and related technologies that collectively are performing both cognitive and manual tasks that would have been difficult to anticipate even a decade ago. In their economic model, they relaxed the literature's constraint that technology is limited to routine tasks and considered three engineering bottlenecks (perception, creative intelligence, and social intelligence) as significantly and practically limiting automation.

In attempting an empirical analysis of these concepts, Frey and Osborne (2017) confronted obvious measurement challenges. Their solution was to combine a subjective assessment with an objective source of information on the task content of occupations (from O*NET) and the level of skill required by the occupations, with respect to the three bottlenecks. The subjective evaluation consisted of expert categorization of a subset of occupations (70 of 702) by participants in a machine learning conference at Oxford University. Each participant was asked to rate an occupation as automatable based on the answer to this question:

“Can the tasks of this job be sufficiently specified, conditional on the availability of big data, to be performed by state-of-the-art computer-controlled equipment?” (Ibid, 30.)

The binary answers to these questions were then modeled as a function of the O*NET-based scores on the bottlenecks. The best-fitting models were then used to calculate an automatable score for all 702 occupations, using the features of jobs that best predicted automation as assessed by the experts. They classified occupations as high-risk if the estimated probability of automation is 70% or higher and low-risk if it is under 30%. This exercise led to the conclusion that 47% of U.S. jobs are at high risk of automation within the next two decades. They found that many jobs in office and administrative support, transportation, and services are at risk, despite the latter not typically being considered routine. Additionally, Webb (2019) finds that AI, in contrast with previous new technologies like software and robots, is directed at high-skill tasks. This research suggests that highly skilled workers may be displaced at a higher rate given the current rate of adoption of AI.

Frey and Osborne (2017) acknowledged that this estimate is not a prediction about the percentage of jobs that will actually be automated, because they explicitly did not model the relative costs of capital versus labor, nor did they consider that technology might partially automate a job. A further limitation is that they did not consider the research and development costs of these potential applications. Thus, as others have pointed out, their result was not a measure of what is economically feasible, so much as an estimate of what is technologically feasible (Arntz, Gregory, and Zierahn 2016).

Two papers from OECD economists have attempted to refine Frey and Osborne’s (2017) estimates and apply them to a larger group of developed countries.

Arntz, Gregory, and Zierahn (2016) used Frey and Osborne’s (2017) occupational results as their main dependent variable and calculated the probability of automation based on the underlying characteristics of the worker and his or her job. Crucially, they allowed job tasks within the same occupational category to vary and have independent effects on the probability of automation, using data from the OECD Program for the International Assessment of Adult Competencies (PIAAC) exam. This approach acknowledged two important things: occupations contain multiple tasks, and even within the same occupation, workers do not perform exactly the same functions at the same level of complexity. Their results showed that jobs that involve more complex tasks are less automatable, especially those involving tasks such as influencing, reading, writing, and computer programming. Moreover, human capital—measured by education level, experience, and cognitive ability—lowers the risk of working in an occupation deemed automatable by Frey and Osborne (2017).

Their final estimate, which they cautioned likely overstates the actual likelihood of automation, predicts that only 9% of workers in the U.S., and in the average OECD country, face a high risk of losing their job to automation within an unspecified number of years—estimated by Frey and Osborne (2017) to be roughly 10 to 20. This is likely to be an overestimate because they did not consider, as the authors pointed out, the slow pace of technological adoption, nor the economic incentives for companies to produce or adopt the technology.

Nedelkoska and Quintini (2018) followed the above methods closely—but allowed tasks to vary based on classification of tasks into Frey and Osborne's (2017) bottlenecks—instead of the more general set of occupational characteristics used by Arntz, Gregory, and Zierahn (2016). Nedelkoska and Quintini's (2018) estimate was nonetheless close to the latter's in that they found 14% of jobs are at risk of automation across 32 countries. They emphasized, however, that roughly half of the jobs across OECD countries could be affected by automation, as some aspect of the job is likely to be changed.

Other relevant research in this area looks at how the skills of occupations have changed in relationship to technology. After analyzing changes in the important job tasks across occupations from 2006 to 2014, MacCrory, Westerman, Alhammedi, and Brynjolfsson (2014) concluded that skills that complement technology (e.g. familiarity with equipment) are increasing along with skills that do not currently compete with machines (e.g. interpersonal). On the other hand, skills that do compete with machines (e.g. manual and perception) are declining, in that workers were less likely to report these skills as important in 2014 relative to 2006 within the same occupation group. The upshot of this analysis is that job tasks are gradually drifting away from work that competes with machines, and these are more likely to be jobs that require less formal education.

After reviewing the latest trends in technology across various industries, Ford (2015) predicted that technology will create massive social disruption, but he also concluded that new technologies “will primarily threaten lower-wage jobs that require modest levels of education and training” (p.26). He was, however, skeptical about the possibility of absorbing millions of displaced workers into high-wage jobs, which he believed will also face competition from machines that are increasingly able to perform cognitive tasks like writing, data analysis, and problem solving.

2.6 Conclusion

Economists have wrestled with the implications of technological change from the beginning of the Industrial Revolution to the present, and yet no theoretical consensus has emerged as to what effects new technologies will have on the labor market.

One of the key constructs in the literature relates to the classification of work into theoretically tractable categories. The early literature limited this to skilled and unskilled labor, which proved too simplistic, as middle-skilled jobs in manufacturing and clerical work experienced lower demand than workers in low-skilled service occupations. The task framework helped fill this gap, by explaining the hollowing out of certain “middle-skilled” occupations, as a result of technological displacement. However, this framework is inadequate for explaining the demand for low-skilled jobs. As Manning (2004) argued, demand for low-paying service jobs can be thought of as stemming from local demand from highly paid workers, adding considerable nuance to the skill-biased technological change framework. If technology raises productivity disproportionately for professional workers, labor demand for lower-paying service jobs may be an inevitable result. Along these lines, Goos, Konings, and Rademakers (2016) also found evidence that STEM jobs create spillovers for non-STEM jobs. Acemoglu and Restrepo's (2019) reinstatement and productivity effects capture some of these potential channels, but additional theoretical and empirical work would be fruitful.

Further refinements of the theory led to the classification of work based on combinations of cognitive, manual, routine, and non-routine task performance. More recently still, the recognition of advances in AI and related technologies has led to the suggestion that perception, creativity, and social intelligence are important cross-cutting characteristics of tasks, and recent work by OECD economists point to specific job tasks that may or may not fit into those categories (Arntz, Gregory, and Zierahn 2016).

Another key construct in the literature involves the classification and measurement of technology. Some papers have attempted to use all-encompassing measures derived from theoretical models (Autor and Salomons 2018 and Acemoglu and Restrepo 2019), such as total factor productivity. While theoretically appealing from a modeling framework, measured productivity growth is caused by multidimensional factors that go beyond technology but are difficult to measure or incorporate comprehensively in a theoretically tractable model. Other scholars have made progress on the empirical study of technology by restricting the analysis to specific technologies, such as industrial robots, computers, software, data storage devices, and related IT equipment. The literature would benefit from both clear evidence of effects from specific technologies, as well as work that attempts to estimate more comprehensive effects across all new technologies, drawing from the more detailed evidence.

Empirical evidence on the direct effects of technology on labor demand has resulted in ambiguous results, with some papers—particularly those looking at industrial machines—suggesting substitution and competition effects dominate between workers and technologies, leading to few jobs in the industries using the technology. On the other hand, studies looking at the adoption of information technology tend to find productivity effects dominate and technology leads to greater labor demand. This difference suggests either that technology has different effects on labor in different sectors, or that unmeasured sector-level factors are the driving force behind either the adoption of technology or the demand for labor exhibited by those sectors.

Ultimately, the more pressing issue is to consider the full effect of technology—including how productivity advances related to technological creation or innovation, followed by adoption—generates gains that flow back into the economy and create demand for labor and upward—though possibly uneven—advances in living standards. Analysis of the indirect effects of technology on labor demand within or outside of technology-using industries has been largely left to theoretical work, as the empirical challenges are significant, with Caselli and Manning (2019), Acemoglu and Restrepo (2019), and Autor and Salomons (2018) providing noteworthy attempts at quantification. The consequences of the Industrial Revolution strongly suggest that technology raises both living standards and the demand for labor, but as pointed out above, many scholars are not convinced that new technologies will have the same effects.

As for which type of workers are most at risk from automation, the Industrial Revolution provides somewhat mixed evidence, with some scholars arguing that the early stages of the Industrial Revolution reduced demand for craft and skilled workers, while others argue that mechanization increased demand for intangible and informal skills. That ongoing debate illustrates the need for clear and compelling definitions of skill that go beyond highest level of education.

Since the 20th century, the consensus among scholars is that workers with lower levels of education or skill whose occupations require the performance of largely routine tasks have experienced the greatest economic threat so far. Likewise, based on the more limited evidence from technological forecasts of Frey and Osborne (2017) and related papers, routine work by less educated workers is also at highest risk over the next 10 to 20 years in terms of experiencing lower wage growth and displacement.

However, with respect to AI, it is unclear what tasks new machines could perform and what pressure that could place on occupations and workers with higher levels of education. Concerns about technology's deskilling effects could resurface if even a small number of occupations that are today considered "skilled" become displaced through automation on a large scale, which may soon be the case, as suggested by Webb (2019).

A challenge for future research entails further refining what characteristics of workers and jobs are most robust to technological change and better understanding and tracking the potential impact of AI. The most compelling empirical work finds that cognitive ability, education, training, and experience provide a relative but not absolute bulwark against job and income loss. This suggests that technology will continue to put pressure on income inequality, as the skill-biased technological change literature and its offshoots predicted. The strongest analyses have made progress by drawing upon unusually detailed data on worker tasks and cognitive ability at the individual level or detailed investment and employment data at the firm level. To push the empirical and theoretical literature forward, scholars will need granular data on skills, tasks, and capital investments. The need for new types or sources of data is especially apparent when attempting to forecast the effects of AI and other advanced automation technologies.

Before concluding, it is worth mentioning that regardless of where the theoretical literature leads in settling various issues such as technology-skill complementarity, technological developments and market dynamics have never operated independently of political forces. Among other policies, mass public education, training, trade, immigration, and public support for research and development have all affected the development of new technologies and the capacity of entrepreneurs and workers to use them.

3. Measuring the Impact of New Technologies on the Labor Market

3.1 Introduction

The proliferation of new technologies, such as digitization, automation, and artificial intelligence (AI), has motivated efforts by international organizations and individual countries to ensure these new technologies have a positive impact on individuals in the labor market. The Organization for Economic Co-operation and Development (OECD) has developed AI principles and recommendations to governments, such as “AI should benefit people and the planet by driving inclusive growth, sustainable development and well-being.”⁵ Similarly, the American AI Initiative states, “The United States must train current and future generations of American workers with the skills to develop and apply AI technologies to prepare them for today’s economy and jobs of the future.”⁶

The outcomes stated in these governing principles require appropriate data to assess how digitization, automation, and AI are affecting labor outcomes and the broader economy. Some national and international statistical agencies have responded to these developments by collecting data on new constructs, such as AI adoption, that complement existing labor market indicators. However, given the complex relationship between skill level, occupational tasks, and the adoption of AI, a variety of data sources are needed to more fully understand the dynamics between technology and the workplace.

This section documents how the BLS currently measures key constructs identified in the literature review.⁷ Additionally, this section examines data collection efforts by other U.S. and international statistical agencies to determine whether there are existing measures that can inform BLS data collection efforts. A review of data products produced by these statistical agencies suggests that there are notable gaps in BLS data products, specifically regarding the classification of skills, task performance, and the adoption of new technologies. While the data collection efforts of other agencies provide instructive approaches and indicators to help address these data gaps, no agency has developed a comprehensive approach to measuring the impact of new technologies on labor market outcomes.

The first part of this section defines each of the key constructs evaluated in the report, followed by a discussion of the methodological approach used to review relevant data products from U.S. and international statistical agencies. The following subsection analyzes how each of the key constructs identified in the literature review is currently captured by statistical agencies, including a description of key data collection efforts. Lastly, we summarize gaps between what BLS currently captures in its data products and the desired set of constructs.

3.2 Key Constructs

A review of the existing literature highlights how economists have attempted to understand the impact of new technologies on labor markets from the Industrial Revolution to the present. While no consensus has emerged in this literature regarding the effects of technological change, this work has highlighted the importance of accurately classifying and measuring technology and work.

Tables 3.1 and 3.2 document the key constructs identified in this literature, as well as the definitions for each construct and the relevant theoretical concepts to which the construct is linked. Additionally, the tables list the gaps in BLS data that have been identified after evaluating BLS data products and external datasets collected by other U.S. agencies and international statistical agencies (see Section 3.3 for a description of the methodology used to identify these gaps).

⁵ “OECD Principles on AI.” Retrieved from: <https://www.oecd.org/going-digital/ai/principles/>.

⁶ “Executive Order on Maintaining American Leadership in Artificial Intelligence.” Retrieved from: <https://www.whitehouse.gov/presidential-actions/executive-order-maintaining-american-leadership-artificial-intelligence/>.

⁷ Submitted to U.S. Department of Labor (DOL), Bureau of Labor Statistics on June 12, 2019 by Gallup, Inc.

Table 3.1 Construct Summary Table – Adoption of New Technologies

Construct	Definition	Necessary Components of Key Theoretical Concepts ⁸	Related Gaps in BLS Data
Productivity Growth	Growth in how efficiently production inputs, such as labor and capital, are being used in an economy to produce a given level of output.	<ul style="list-style-type: none"> Productivity Effect 	<ul style="list-style-type: none"> Metropolitan measures of productivity Adjustments for intangible investments
Diffusion of Technology	The use of new technology, such as digitization, automation, and AI, in production and other market activities.	<ul style="list-style-type: none"> Skilled-Biased Technological Change Job Polarization Capital-Skill Complementarity Reinstatement Effect Displacement Effect 	<ul style="list-style-type: none"> Data to differentiate between technologies Data that track purchase and adoption of new technologies over time
Advanced Industries	Industries that tend to create new technologies, defined by high levels of research and development (R&D) and high employment of workers in STEM (science, technology, engineering, and math) occupations.	<ul style="list-style-type: none"> Skilled-Biased Technological Change Job Polarization Capital-Skill Complementarity Reinstatement Effect Displacement Effect 	<ul style="list-style-type: none"> Linked micro-data between existing BLS and external datasets to better assess the impact of technology on labor outcomes at the establishment level

Since technological change is the independent variable of interest in this analysis, the first set of constructs focuses on the adoption of new technologies. The first key construct is productivity growth, which is a measure of how efficiently production inputs (both capital and labor) are used in an economy to produce a given level of output. Theoretically, productivity growth is interpreted as an outcome of new technology, and so productivity growth is sometimes used as a proxy for technological change. Historically, technology has produced wage growth and a higher standard of living. However, some have posited that productivity gains may increasingly go to capital (i.e., investors) instead of labor or only to highly skilled workers or those who perform certain tasks. Productivity growth is a key construct for a better understanding of this relationship as technology changes and becomes more widely adopted in the economy.

Productivity growth is key to measuring the *productivity effect* identified in the literature (Acemoglu and Restrepo 2019). This effect suggests that there will be an increase in productivity as a result of labor being more productive at remaining non-automated tasks.

⁸ Each concept is defined in the text. For a full discussion of each theoretical concept, see Section 2.

Table 3.2 Construct Summary Table – Labor Market Outcomes

Construct	Definition	Necessary Components of Key Theoretical Concepts	Related Gaps in BLS Data
Labor Demand	The need for workers in a particular job market.	<ul style="list-style-type: none"> • Reinstatement Effect • Displacement Effect 	<ul style="list-style-type: none"> • Real-time labor demand across occupations, skill-levels, and task profiles
Skill Demand	The need for labor with specific skills or abilities to perform specific tasks.	<ul style="list-style-type: none"> • Skilled-Biased Technological Change • Capital-Skill Complementarity • Job Polarization 	<ul style="list-style-type: none"> • Data on cognitive ability, non-cognitive ability, and other skills
Tasks	Roles, responsibilities, and activities performed by a worker in the production of a good or service.	<ul style="list-style-type: none"> • Job Polarization • Productivity Effect • Reinstatement Effect • Displacement Effect 	<ul style="list-style-type: none"> • Well-documented data on tasks performed by workers over time • Links to external datasets on tasks performed by machines

Additional constructs focus on the adoption of specific technologies within the economy. First, the diffusion of technology, or businesses' use of new technology in production or other market activities, is key to understanding how fast and to what extent technology impacts the labor market and the broader economy. The diffusion of technology is also a dimension of key theoretical concepts related to the technology/labor market relationship, such as the reinstatement and displacement effects (Acemoglu and Restrepo 2019). The *reinstatement effect* is a relationship where automation increases the number of tasks performed in the market by humans, increasing the demand for labor. The *displacement effect* is a relationship where automation replaces some portion of the tasks performed by humans, decreasing the demand for labor.

Some research, however, suggests that diffusion of technology has a heterogenous effect on labor market outcomes across industries—labor demand is likely to be different for firms in technology-producing industries relative to firms that simply use technology (or firms that neither create nor use technology). These technology-producing firms have been termed “advanced industries,” characterized by high levels of investment in R&D and their propensity to employ workers in STEM occupations (Muro et al. 2015).

In addition to constructs related to the adoption of new technologies, this analysis includes several key labor market outcomes: labor demand, skills demand, and allocation of tasks. First, labor demand is a core construct needed to understand the effects of new technology on employment. Acemoglu and Restrepo's (2019) classification of the primary effects of technology (discussed above) on labor market outcomes uses demand for labor, or the need for employees and workers in a particular job market, as the key outcome of interest.

The central models of technological change and its impact on labor markets, however, rely on different ways of classifying labor into theoretically tractable categories. Since new technologies have heterogeneous effects on labor, the literature's central focus has been on which occupational traits have the greatest impact on labor outcomes when new technologies are introduced. The first strand of the literature focuses on skills, positing that new technologies shift the skills needed by employers. Therefore, a key construct is skill demand, or the need for labor with specific skills or abilities. While early research focused on varying demand for skilled versus unskilled labor, more advanced conceptualizations have included middle-skill workers, as well as a focus on the advantages of specific skills, such as advanced cognitive skills, technological, and socio-emotional skills, over basic cognitive and manual skills.

These skills frameworks introduce theoretical constructs, such as *capital-skill complementarity*, which is the extent to which advancing technology (capital) complements or displaces highly skilled workers. Another theoretical construct is *skilled-biased technological change*, which is the proposed relationship between technology and skill that results in an increase in the relative demand for skilled workers related to unskilled workers as a result of their relationship to new technologies.

Additional work, such as Acemoglu and Autor (2011), proposed a task-based framework where tasks are defined as units of work activity that produce output (in this framework, skills refer to the ability to perform tasks). This task construct is necessary to understand occupational growth patterns that differ within skill levels and the capacity to model machines as potential substitutes for work—rather than only labor-augmenting. Theoretical constructs related to the task-based framework include the productivity, reinstatement, and displacement effects discussed above. Another construct is *job polarization*, in which there is a rising demand for workers in low- and high-skilled occupations relative to workers in middle-skilled occupations as a result of technology that has displaced routine manual and cognitive tasks.

3.3 Methodology

To assess how these key constructs are currently captured by BLS, all of the agency's relevant data products were identified. They were then analyzed for current data collection efforts on the key constructs and individually reviewed to assess any current gaps in information on productivity growth, adoption of technology, advanced industries, labor demand, skill demand, or tasks.

This gap analysis was informed by an evaluation of relevant datasets produced by U.S. and international statistical agencies. In the U.S., all Principal Statistical Agencies of the Federal Statistical System are included in this analysis, along with other Federal agencies that collected data related to education, workforce development, and science and technology.⁹

Internationally, this paper assesses national statistical agencies of non-U.S. countries as well as the data products of intergovernmental organizations (e.g., Eurostat, OECD, World Bank). National statistical agencies of all 36 OECD member states¹⁰ are included in this analysis, representing high-income economies that may produce statistical products relevant to the U.S. economy. Additionally, some non-OECD countries are included if they received the top score (100) on the methodology assessment of statistical capacity on the World Bank's Statistical Capacity Index in 2018.¹¹ All international organizations listed in the UN's Global Inventory of Statistical Standards were also included in this analysis.¹²

⁹ "Statistical Programs of the United States Government." 2018. Executive Office of the President, Office of Management and Budget. Retrieved from: www.whitehouse.gov/sites/whitehouse.gov/files/omb/assets/information_and_regulatory_affairs/statistical-programs-2018.pdf.

¹⁰ "Member Countries." Retrieved from: <https://www.oecd.org/about/members-and-partners/>.

¹¹ Armenia, China, India, Moldova, and Ukraine received the top rating (100) in 2018. Retrieved from: <http://datatopics.worldbank.org/statisticalcapacity/SCIdashboard.aspx>.

¹² "List of International Organizations." Retrieved from: <https://unstats.un.org/unsd/iiss/List-of-International-Organizations.ashx>.

Once these statistical agencies were identified, they were reviewed for data collection efforts relevant to the following individual-level measures: employment, occupation, wages and compensation, skill level (including educational attainment, cognitive ability, and skill related to literacy, numeracy, and computer skills), and occupational tasks. Additionally, these agencies were reviewed for any establishment-, firm-, or household-level data collection efforts that may capture the adoption of technology by sector and industry, the classification of workers within establishments by occupation, income, or education, as well as R&D activities and expenditures.

A key finding from the review of international agencies is that there is significant overlap and coordination between the national statistical agencies included in this analysis. Given that most of the countries included in this technical report are members of the EU and/or OECD, current data products are fairly consistent across countries, and there are very few unique country-level datasets relevant to the topics of this study. In some instances, national statistical agencies are required to conduct surveys and submit the data to Eurostat, following common guidelines to ensure comparability across member states. For instance, in the case of the European Union Labour Force Survey, member states must select the sample, develop questionnaires, conduct household interviews, and forward the data to Eurostat, following guidelines such as using common concepts, definitions, and classifications.¹³

Beyond the EU, OECD members often coordinate large-scale data collection efforts, such as the Survey of Adult Skills and the Survey on the Use of Information and Communication Technology (ICT) by Businesses. These relevant datasets are collected across countries by national statistical agencies and likely supplant other data collection efforts on key topics at the country level. Therefore, there is significant cross-national consistency in the data collection efforts relevant to the key constructs on technological change and labor market outcomes.

3.4 Measurement of Key Constructs

3.4.1 Productivity Growth

Description of Measures. The most commonly used measure of productivity growth is labor productivity, or the output per unit of labor. The total volume of output is measured using GDP data, and total hours worked is typically used as the unit of labor. However, in some countries with less reliable labor market data, agencies use the less granular measure of the number of employed persons as the unit of labor.

In addition to labor productivity, in which labor is the only input, multifactor productivity (MFP) is also widely measured by national and international statistical agencies. MFP is the output per unit of a set of combined outputs, including capital (e.g., equipment, land, structures, and existing inventory), labor, energy, materials, and purchased business services. Any MFP growth reflects an increase in outputs that cannot be accounted for by the change in combined inputs. Therefore, increases in MFP can be interpreted as productivity gains attributed to factors such as new technologies, R&D, or other organizational improvements.

Existing BLS Data Products. Two BLS programs produce labor productivity measures for sectors of the U.S. economy. The Major Sector Productivity program maintains quarterly and annual measures of output per hour (labor productivity) for private business and nonfarm business from 1947 to the present (BLS 2008). Quarterly and annual labor productivity measures are also available for manufacturing (total, durable, and nondurable sectors) from 1987 to the present and for nonfinancial corporations from 1958 to the present.

¹³ Council Regulation (EEC) No. 577/98 of March 9, 1998.

Additionally, the Industry Productivity program produces annual labor productivity measures for U.S. industries at each North American Industry Classification System (NAICS) level,¹⁴ including comprehensive coverage of mining, utilities, manufacturing, wholesale trade, retail trade, and accommodation and food services. To calculate the output portion of labor productivity, BLS uses Bureau of Economic Analysis (BEA) measures of U.S. output (real GDP) by major sector. Specifically, BLS uses real business sector output, which excludes outputs produced by the government, nonprofit institutions, employees of private households, and the rental value of owner-occupied dwellings.¹⁵

For the industry productivity measures, output is measured as real sectoral output, the total value of goods and services leaving the industry. Industry output measures are constructed using a series of U.S. Census Bureau industry surveys, including the Annual Survey of Manufactures, Census of Manufactures, Annual Retail Trade Survey, Census of Retail Trade, Monthly and Annual Wholesale Trade Surveys, Census of Wholesale Trade, Service Annual Survey, and Census of Service Industries.¹⁶

Labor input is primarily derived from the BLS Current Employment Statistics (CES) program, which provides monthly survey data on employment and average weekly hours of workers employed in the nonfarm business sector.¹⁷

In addition to labor productivity measures, BLS publishes two sets of MFP measures for the major sectors and subsectors of the U.S. economy, each using a distinct methodology (BLS 2007). The first is MFP for major sectors, while the second set measures MFP for total manufacturing (Major Sector Productivity program) and 18 three-digit NAICS manufacturing industries (Industry Productivity program).¹⁸ These annual measures are available from 1987 to the present.¹⁹

External Data Products. Productivity growth is a construct that has been consistently measured in major developed economies since the post-World War II period. The methodology for measuring labor productivity and MFP has been standardized cross-nationally with Eurostat and national agencies publishing productivity per person employed and productivity per hour worked. The OECD provides detailed productivity measures across its member countries, as well as data from accession countries, key partners, and some G20 countries.²⁰ These measures include GDP per hour worked (1971 to the present) and MFP (1985 to the present).²¹

¹⁴ NAICS levels includes sector (two-digit code), subsector (three-digit code), industry group (four-digit code), NAICS industry (five-digit code), and national industry (six-digit code). See the following for sector definitions: <https://www.census.gov/programs-surveys/economic-census/guidance/understanding-naics.html>.

¹⁵ The nonfarm business sector output also excludes the farm sector.

¹⁶ "Industry Productivity Measures." <https://www.bls.gov/opub/hom/inp/pdf/inp.pdf>.

¹⁷ BLS also uses data from the Current Population Survey (CPS) and National Compensation Survey (NCS) for its quarterly productivity measures and industry productivity measures (BLS 2008).

¹⁸ The data inputs used to calculate MFP vary significantly between different sector and industry levels. More details about data sources are available at: <https://www.bls.gov/mfp/mprover.htm#data>.

¹⁹ Not all countries produce annual industry MFP measures. For instance, Australia reports MFP in two-year increments: <https://www.abs.gov.au/AUSSTATS/abs@.nsf/Lookup/5260.0.55.002Main+Features12017-18?OpenDocument>.

²⁰ See OECD (2019) for a comprehensive overview of trends in productivity levels and growth in OECD and some non-OECD countries, including a breakdown of productivity by industry and firm size.

²¹ Data not available for every country on an annual basis. Retrieved from: https://stats.oecd.org/viewhtml.aspx?datasetcode=PDB_GR&lang=en.

The only differences internationally, when compared with BLS, is that some national governments publish public sector productivity data. For instance, the Office for National Statistics (ONS) publishes “public service” productivity estimates, which include healthcare, education, social services, defense, police and public safety in the U.K. These data are published to track productivity in government and non-profit establishments (BLS only produces estimates on the business sector as part of its Major Sector Productivity program). However, like BLS, ONS is unable to measure output for some service areas. Therefore, ONS utilizes the “outputs-equals-inputs” convention. This method assumes that productivity remains constant and growth will always be zero, which does not provide meaningful information on productivity in these areas. This approach represents 38% of total public service output in the U.K., highlighting the degree to which outputs-equals-inputs drives public sector productivity measures.²² While public sector productivity would be a valuable measure, given that new technologies also impact productivity outside of the business sector, it is not recommended that BLS adopt this approach without more accurate measures of outputs in these areas.

Outside the OECD, the International Labour Organization (ILO) measures labor productivity for every country in the world, including some autonomous territories, regions, and groups of countries (e.g., the G7, G20, and EU). For comparability, given the lack of data collected within some countries, ILO annually publishes output per worker (rather than output per hour worked), where output is measured using GDP in constant international and U.S. dollars.

Analysis of Data Gaps. Productivity growth is the most widely available time-series cross-sectional economic indicator relevant to the study of new technologies. The main gap in BLS data products, however, is that productivity measures are not available at more granular geographic and temporal units with the same methods. In the U.S., the BLS recently produced experimental data on state productivity, reported as output per hour worked. These data are being evaluated for possible improvement and better alignment with national productivity figures (Pabilonia et al. 2019).

The BLS does not currently produce productivity statistics at the metropolitan area level, but the BEA provides output per capita measures at the state, metropolitan, and county scales in the U.S., which can be used as productivity measures, though geographic differences in labor force participation complicate comparisons. These data are available going back to 1929 for states and 1969 for metropolitan areas and counties.²³ By using the number of residents rather than the number of workers or hours worked, these productivity statistics are not as accurate as the national productivity statistics or the experimental BLS state level productivity measures. However, recent research has underscored the importance of geo-industrial clusters on labor markets (Park et al. 2019), so a more fine-grained productivity measure would provide a better understanding of these local dynamics.

BLS faces additional measurement challenges with subnational productivity, in that hours worked data are often limited to place of residence rather than place of work. This is because commonly used household surveys, such as the CPS and American Community Survey, capture place of residence, which may not equal place of production. The Australian Bureau of Statistics²⁴ and researchers at BLS (Pabilonia et al. 2019) have developed sub-national measures of productivity, yet use simplifying assumptions related to the potential measurement error related to the use of household surveys.

²² “Public services productivity.” Retrieved from:

<https://www.ons.gov.uk/economy/economicoutputandproductivity/publicservicesproductivity>

²³ U.S. Bureau of Economic Analysis, “MAINC1 Personal Income Summary: Personal Income, Population, Per Capita Personal Income.” Retrieved from: <https://apps.bea.gov/iTable/iTable.cfm?reqid=99&step=1#reqid=99&step=1>.

²⁴ “Experimental Estimates of State Productivity.” Retrieved from:

<https://www.abs.gov.au/ausstats/abs@.nsf/Lookup/5260.0.55.002Feature+Article12016-17>.

An additional data gap focuses on the underestimation of output. Specifically, some research has suggested that new technologies, such as AI, require significant investments. However, there is a period, which can be lengthy, when measurable resources are committed, and measurable outputs are forgone to develop new, unmeasured inputs to complement new technologies. Thus, productivity appears to decline in the short term while these investments are made, and then true productivity growth is eventually overestimated once the hidden inputs (new technologies) generate measurable output (Brynjolfsson et al. 2018).

While researchers have shown how accounting for intangible investments correlated with measurable investments can substantially impact estimates of productivity growth, additional work is required to better understand how current productivity measures may introduce bias in the face of unmeasured inputs. BLS data products do not currently include any adjustments for these inputs, which may lead to mismeasurement related to R&D, software, and hardware investments.

Beyond the scope of BLS, there are additional challenges when comparing U.S. productivity to other countries. For instance, ILO uses national output measures obtained from country-level agencies. However, for countries outside the OECD, common principles are not always followed in national accounts estimates, leading to some inconsistency. Therefore, a major weakness of this data cross-nationally is a lack of comparability caused by differences in the treatment of output in services sectors, differences in how agencies correct output measures for price changes, and differences in the degree of coverage of informal or underground economic activities in each country.

3.4.2 Diffusion of Technology

Description of Measures. The primary measure of the diffusion of technology is the use or purchase of automation equipment by firms. This information is self-reported by individual businesses on surveys conducted by national and international statistical agencies. Some proprietary datasets from industry associations also collect sales and installation data from technology suppliers. For instance, the International Federation of Robotics (IFR) tracks data on the number of industrial robot installations performed by suppliers (Graetz and Michaels 2018). While these data have unique information on robot price and density, using the “gold standard” definition of industrial robots²⁵, the data are not available at lower levels of disaggregation. Additionally, IFR does not have consistent country-industry-level data on service robots.

Existing BLS Data Products. BLS does not currently collect data that directly measure the diffusion of technology. However, BLS data may be used to provide more detailed employment information about the industries and establishments that purchase automation equipment. Specifically, the Occupational Employment Statistics (OES) program produces employment and wage estimates for specific occupations and industries in the U.S.

The OES program surveys approximately 180,000 to 200,000 establishments per panel (every six months), collecting data on the full sample of 1.2 million establishments every three years. The OES program produces employment and wage estimates for about 800 occupations, as well as estimates for approximately 415 industry classifications at the national level. These data enable BLS to evaluate how the diffusion of technology, measured using external datasets, is influencing employment levels and wages within specific industries.

Additionally, improved ability to link OES microdata with other data products produced by U.S. statistical agencies (e.g., the Annual Business Survey (ABS), discussed below) would expand the ability to assess the impact of technology adoption on establishment-level employment and wage dynamics (see Section 4 for a full discussion of linking establishment-level and occupational data).

²⁵ The IFR definition refers to a “manipulating industrial robot as defined by ISO 8373: An automatically controlled, reprogrammable, multipurpose manipulator programmable in three or more axes, which may be either fixed in place or mobile for use in industrial automation” (Graetz and Michaels 2018).

External Data Products. Outside of BLS, there are two primary sources of data on adoption of new technologies in the U.S. First, the BEA uses Census-collected data on fixed assets, which are used continually in the process of production for an extended period of time. These data include statistics on net stocks, beginning with 1925, and are reported by industry, legal form of organization, and asset type. Of particular relevance to the diffusion of technology is the private fixed investment of equipment, including assets such as information processing equipment (e.g., computers and communication equipment), industrial equipment, and transportation equipment.

Second, the ABS was launched in 2017 as a joint project between U.S. Census Bureau and the National Science Foundation's (NSF) National Center for Science and Engineering Statistics.²⁶ The survey covers all nonfarm employer businesses filing the 941, 944, and 1120 tax forms, and participants were selected from a sample of 850,000 employer businesses in the first year of the survey in 2017.²⁷ Approximately 300,000 employer businesses will be sampled annually in years 2018-2021. Data are collected via web only.

On the 2019 survey, the ABS includes a module on technology and intellectual property, which asks businesses to report their use of technologies such as AI, cloud-based computing systems and applications, and robotics. Given the complexity of these technologies, the ABS defines each technology. For instance, AI is defined as "a branch of computer science and engineering devoted to making machines intelligent. Intelligence is that quality that enables an entity to perceive, analyze, determine response and act appropriately in its environment," while robotics are distinguished as "automatically controlled, reprogrammable, and multipurpose machines used in automated operations in industrial and service environments."²⁸

Each business is asked whether it used these new technologies in production processes for goods or services during the past three years. Use is categorized as high, moderate, or low, and respondents can respond that they tested but did not use these technologies in production or service. Those who reported using AI, robotics, or other technologies were asked about their motivation for adopting these technologies, including the following options:

- to automate tasks performed by labor
- to upgrade outdated processes or methods
- to improve the quality or reliability of processes or methods
- to expand the range of goods or services
- to adopt standards and accreditation
- some other reason

Additionally, the ABS measures whether each new technology a business has adopted has increased, decreased, or did not change the number of workers, skill level of workers, and/or the STEM skill of workers employed by the business. Furthermore, each adopter was asked whether each new technology increased, decreased, or did not change the number of production workers, nonproduction workers, supervisory workers, and nonsupervisory workers. This information provides key information about how the diffusion of technology has impacted employment and wages at the establishment level, as well as how it has influenced the employment of individuals with specific skills.

Finally, the ABS produces indicators on the factors that have adversely affected the adoption or utilization of specific technologies, such as AI and robotics, to produce goods or services. Businesses were provided with factors, such as the price of the technology, the availability of the required human capital and talent, and lack of access to capital. This information highlights potential barriers to the adoption of new technologies, such as lack of access to employees with the required skills.

²⁶ The ABS replaced three existing surveys: the five-year Survey of Business Owners (SBO) for employer businesses, the Annual Survey of Entrepreneurs (ASE), and the Business Research and Development and Innovation Survey for Microbusinesses (BRDI-M).

²⁷ The sample is stratified by state, frame, and industry and is systematically sampled within each stratum.

²⁸ "Annual Business Survey (ABS): Surveys & Instructions." Retrieved from: <https://www.census.gov/programs-surveys/abs/technical-documentation/surveys-instructions.html>.

In addition to these questions on the adoption of new technologies, the U.S. also measures a broader category of “innovation activities” at the business level on the ABS. The “Products and Processes” module on the ABS asks a series of questions about the introduction of new goods or services and R&D activity (discussed in Section 3.4.3). Most relevant to the study of new technologies and their impact on the labor market are several questions that ask whether the business has engaged in software development and database activities, as well as the acquisition of machinery, equipment, and other tangible assets.

While this measure of software development and database activities and the acquisition of machinery and equipment does not precisely capture the adoption of new technologies, it does provide data on expenditure totals for these activities, which can be linked to intangible inputs that impact productivity growth (see discussion of productivity measures in Section 3.4.1). Businesses report their spending on innovation activities over the past year. This provides an additional measure of technology adoption since it captures a dollar amount on expenditures on activities, such as software development, data analysis, and purchase of machinery, instead of a binary or categorical measure of usage.

These innovation measures also have the benefit of comparability to other countries. The innovation module on the ABS was previously included in the Business R&D and Innovation Survey (BRDIS),²⁹ which was collected from 2009 to 2016. The BRDIS, which was developed by the NSF and Census Bureau, was adapted from the EU’s Community Innovation Survey (CIS). The CIS is conducted every two years by EU member states and additional countries within the European Statistical System.³⁰ Statistical agencies outside of Europe, including those in Africa, Asia, North America, Oceania, and South America also conduct innovation surveys using OECD/Eurostat’s Oslo Manual as a common set of guidelines for collecting, reporting, and using innovation data (OECD/Eurostat 2018).

Given this international standardization, other countries do not collect data on innovation measures that are not currently measured by U.S. statistical agencies. For instance, the most recent data on innovation activities stem from a 2015 global innovation data collection effort, which includes data from 71 countries, including the U.S. (UNESCO 2017). Given the diverse set of countries involved in this effort, the methodology varies slightly between countries. This includes varying observation periods, statistical units (enterprise, enterprise group, establishment, kind of activity unit), sampling frame (national statistical business register, business association lists), survey methods (sample survey, census), and coverage of firms (micro, small, medium, and large firms). However, each country used the same standard questionnaire to ensure comparability of results.

While the U.S. adopts measures of innovation activities in line with other countries, the scope of measurement on new technologies is broader internationally, particularly in the EU. Eurostat mandates the collection of ICT usage indicators in member states. The survey, broadly known as the Community Survey on ICT Usage and E-Commerce in Enterprises, includes all enterprises with 10 or more employees classified in industries such as manufacturing, utilities, construction, wholesale and retail, transportation, lodging and food services, ICT, real estate, professional and technical activities, administrative and support service activities, and computer repair.

²⁹ “Business R&D and Innovation Survey (BRDIS).” Retrieved from: <https://www.census.gov/programs-surveys/brdis.html>.

³⁰ “Community Innovation Survey (CIS)” Retrieved from: <https://ec.europa.eu/eurostat/web/microdata/community-innovation-survey>.

The survey includes modules on big data analysis, employment of ICT specialists, the Internet of things (IoT), and use of robotics. The most relevant measures ask enterprises whether they performed big data analysis on the following data sources: data from smart devices or sensors (e.g., machine to machine communications, digital sensors, radio frequency identification tags), geolocation data from the use of portable devices (e.g., portable devices using mobile telephone networks, wireless connections or GPS), data generated from social media, and other big data sources (e.g., stock index data, transaction data, other open web data). The survey also measures whether the data were analyzed using machine learning (e.g., deep learning), which “involves ‘training’ a computer model to better perform an automated task” or natural language processing (NLP), “the ability for a computer program to understand human language as it is spoken, to convert data into natural language representation or to identify words and phrases in spoken language and convert them to a machine-readable format.”³¹

The measurement of big data analysis, particularly related to sensors, geolocated data, social media, and other sources, is not reflected in data collected by BLS or other U.S. statistical agencies. However, this represents a key activity related to AI, as businesses adopt machine learning techniques and NLP. However, existing measures in the U.S. focus on AI more generally, which does not provide detailed information on how specific AI applications may impact labor market outcomes.

Additionally, the ICT survey asks about robot usage, distinguishing between industrial robots and service robots (unlike the ABS, which focuses on adoption of robotics more generally).³² Industrial robots are conceptualized as “automatically controlled, reprogrammable, multipurpose manipulator programmable in three or more axes, which may be either fixed in place or mobile for use.” Service robots are defined as “machines that have a degree of autonomy that enables them to operate in a complex and dynamic environment that may require interaction with persons, objects, or other devices, excluding its use in industrial automation applications. They are designed to fit their tasks, working in the air (e.g., as a drone), underwater, or on land, using wheels or legs to achieve mobility with arms and end effectors to physically interact and are often used in inspection and maintenance tasks.”

A follow-up question asks about the tasks performed by the robots, including the following:

- Surveillance, security, or inspection tasks (e.g., use of autonomous airborne drones, etc.)
- Transportation of people or goods (e.g., use of automated guided vehicle, etc.)
- Cleaning or waste disposal tasks
- Warehouse management systems (e.g., palletizing, handling goods, etc.)
- Assembly works performed by service robots
- Robotic store clerk tasks
- Construction works or damage repair tasks

Collecting data on the tasks performed by the robots, instead of general adoption of robotics, would be particularly beneficial in the U.S. to evaluate how automation is impacting the labor market and, potentially, to measure reinstatement and displacement effects in the U.S. economy (by linking to BLS data on employment in these establishments).

³¹ “Community Survey on ICT Usage and E-Commerce in Enterprises, 2019, General outline of the survey.” Retrieved from: <https://circabc.europa.eu/sd/a/47b2dcfa-2eb9-4cc4-9e98-b93a85406d67/ICT-Entr%202020%20-%20Model%20Questionnaire%20V%201.2%20final.pdf>.

³² Outside of statistical agencies, other organizations collect establishment-level data on robotic systems. Fundación SEPI, a non-profit managed by a Spanish state holding company, maintains a panel dataset of Spanish manufacturing firms. The *Encuesta Sobre Estrategias Empresariales* (ESEE) has been collected annually over a 26-year period (1990-2016), covering approximately 1,900 Spanish manufacturing firms a year. The ESEE collects data on whether any of the following systems were used in the production process: computer-digital machine tools; robotics; computer-assisted design; combination of some of the above systems through a central computer (CAM, flexible manufacturing systems, etc.); and Local Area Network (LAN) in manufacturing activity” (Koch et al. 2019).

In addition to these survey measures, patent data represent another important aspect of technology diffusion and data are widely available using the OECD REGPAT Database.³³ These data include information on the country, subnational region, and patent details based on the address of patent applicants and inventors. This includes data on more than 2,000 regions within OECD countries, including the U.S.³⁴ However, a major limitation is a lack of access to industry data and micro-data at the firm level, which could be matched to other databases on revenue, employment, and wages.

Lastly, some have suggested using data on private equity investments in AI companies as a supplementary measure of commercial applications of AI (OECD 2018). Using proprietary data from Crunchbase on over 500,000 entities located in 199 countries, firms are classified by whether their primary activities list generic AI keywords (e.g., AI, machine learning), keywords pertaining to AI techniques (e.g., neural networks, deep learning), and keywords referring to AI applications (e.g., NLP, autonomous vehicles). This information provides granular information about each firm and whether its business activity is AI-focused. However, it still relies on self-identified business details and may omit the adoption of automation and AI beyond the firm's primary business activity.

Analysis of Data Gaps. Measurement of the adoption of new technologies is still in a nascent phase. Many new technologies have only had viable business applications in recent years and given the time needed to develop, administer, and process surveys on these applications, there is relatively sparse data on usage of AI and robotics collected by national statistical agencies.

In recent years, business surveys have included new measures of AI adoption at the enterprise level. However, these measures remain fairly general, often asking about the broad usage of AI or robotics. For instance, the Canadian Survey of Innovation and Business Strategy—using guidelines from the OECD's Oslo Manual for collecting innovation data—measures business usage of AI but does not ask about specific applications of AI.³⁵ Additionally, the ABS collects data on AI and robotics usage, as separate technologies, but does not classify these technologies at a more granular level.

Therefore, the primary data gap that BLS faces is the need for data that track the purchase and adoption of technologies over time, as well as to better differentiate between technologies. External surveys, such as Eurostat's Community Survey on ICT Usage and E-Commerce in Enterprises, address this gap at an international level. This survey, conducted in each EU member state, measures how businesses have adopted specific applications, such as machine learning, NLP, and different types of robots. Furthermore, it measures the tasks performed by industrial and service robots, which is necessary to fully understand how these technologies are displacing or reinstating human labor.

Since BLS does not currently capture diffusion of technology in its data products, it is reliant on other U.S. statistical agencies developing these more nuanced measures of technology adoption at the industry and national levels. Once linked with BLS data, namely employment and wage data from the OES program, it will then be better able to assess how the adoption of specific technologies, such as machine learning or robots performing different types of tasks, impact labor market outcomes.

3.4.3 Advanced Industries

Description of Measures. Advanced industries are those with high levels of R&D and high employment of workers in STEM occupations. This classification is not directly collected by national statistical agencies, but component measures can be used to capture these technology-producing industries. First, information on firm- and industry-level R&D are widely available, including binary measures of R&D activity, as well as spending on R&D. Additionally, worker characteristics, such as knowledge in the STEM domain, is captured by both individual- and firm-level surveys.

³³ https://www.oecd-ilibrary.org/science-and-technology/the-oecd-regpat-database_241437144144

³⁴ REGPAT uses information from patent applications to the European Patent Office, Patent Cooperation Treaty, and United States Patent and Trademark Office (USPTO). Therefore, it includes better coverage than using USPTO data alone.

³⁵ Retrieved from: <http://www23.statcan.gc.ca/imdb/p2SV.pl?Function=getSurvey&SDDS=5171>

Existing BLS Data Products. BLS does not currently collect data that directly classify advanced industries. However, OES data may be used to link employment and wage information about these technology-producing industries. When combined with R&D measures (described below), OES data can be used to evaluate the labor market dynamics within advanced industries.

Furthermore, BLS researchers have used OES data to categorize occupations into a STEM category (Fayer et al. 2017). Combining occupational categories ranging from computer, mathematical, and science occupations to STEM-related sales, Fayer et al. estimate that there were nearly 8.6 million STEM jobs in May 2015. Using this same methodology, OES data can link STEM occupations to specific industries and measure whether employment and wage outcomes differ between advanced industries and those that do not produce or use technology.

External Data Products. While advanced industries are not currently classified by U.S. statistical agencies, existing detailed R&D data can be used to classify whether an industry produces technology. The ABS and, prior to 2017, the BRDIS serve as primary data sources to characterize firms' engagement in technology production. The 2019 ABS measures whether a business uses new technologies (discussed in the previous section), as well as whether they sold technologies such as AI, cloud-based computing systems and applications, specialized software, robotics, and specialized equipment. Follow-up questions for technology-producing businesses focus on their motivation for producing these new technologies and whether production has increased, decreased, or had no impact on the number and skill of workers employed by the business.

Beyond the production of new technologies, the ABS includes an extensive section that captures R&D activities among firms. R&D is defined in the Frascati manual as "the creative and systematic work undertaken in order to increase the stock of knowledge and to devise new applications of available knowledge" (OECD 2015). Firms are first asked about various forms of R&D performed in the past year:

- Activities aimed at acquiring new knowledge or understanding without specific immediate commercial applications or uses
- Activities aimed at acquiring new knowledge for solving a specific problem or meeting a specific commercial objective
- Systematic work, drawing on research and practical experience and resulting in additional knowledge, which is directed to producing new products or processes or to improving existing products or processes
- Development and testing goods, services, or processes that were derived from scientific research or technical findings
- Development of software that advanced scientific or technological knowledge
- Production of findings that could be published in academic journals or presented at scientific conferences
- Application of scientific or technical knowledge in a way that has never been done before
- Creation of new scientific or technical solutions that can be generalized to other situations
- Work to discover previously unknown technological facts, structures, or relationships
- Work to extend the understanding of scientific facts, relationships, or principles in ways that could be useful to others

The ABS includes measures of both foreign and domestic R&D costs at the business level, as well as a breakdown in types of R&D costs (e.g., compensation, machinery, etc.). Additionally, firms provide the number of R&D employees, researchers, and funding sources for R&D activities. In combination, these measures provide a comprehensive view of a firm's R&D outlay and number of employees engaged in technology-producing tasks.

This approach to measuring R&D is consistent with the standardized methodology employed by other countries. In addition to the R&D data produced by Eurostat's CIS, country-level Business Enterprise Research and Development (BERD) surveys³⁶ provide key statistics on R&D using Frascati manual guidelines (OECD 2015). These are the same guidelines used by the U.S. The OECD has collated country-level data to produce a new taxonomy of industries per country according to their average level of R&D intensity, called the Analytical Business Enterprise Research and Development (ANBERD) database.³⁷ Additionally, UNESCO Institute for Statistics conducts an annual R&D survey,³⁸ producing cross-national data (including in developing countries) on R&D personnel, as well as total expenditure by sector.

In addition to R&D activity, the other key dimension of advanced industries is related to labor market outcomes. A key hypothesis related to advanced industries is that labor market demand will likely be different for technology-producing industries relative to technology-using industries. The ABS data enable a more direct test of this hypothesis since it measures whether the use of new technologies, such as AI or robotics, has impacted the number of workers overall, the number of skilled workers, and the number of workers with STEM skills employed by the business. Additionally, businesses report any changes in the employment of production, nonproduction, supervisory, and nonsupervisory workers, which enables a direct assessment of new technologies' impact on labor demand across different tasks and skills.

In addition to using OES data to construct a STEM occupational category, other researchers have used O*NET data to determine the STEM knowledge intensity of four-digit NAICS industries (Rothwell and Kulkarni 2015). This database was created by the DOL's Employment and Training Association and collects a variety of detailed data from workers on various aspects of their job and the job's requirements. In identifying STEM skills, the O*NET knowledge survey asks workers in specific occupations to rate the level of knowledge in STEM and other domains required to do their job. The level of knowledge required in each STEM domain can thus be quantified for every occupation. Using these measures of STEM-intensity per occupation, BLS industry-occupation matrixes can be used to determine the percentage of workers in STEM occupations for every four-digit NAICS industry at the national level.

Analysis of Data Gaps. There are no significant data gaps in categorizing advanced industries in the U.S., as long as existing micro-data across agencies can be linked. Specifically, micro-level data at the establishment level from non-BLS data sources are needed to assess whether the adoption or production of new technologies within establishments serves to displace or reinstate labor (using BLS data discussed above). R&D activity and expenditures do not suffer from some of the same data constraints as diffusion of technology. R&D has long been measured cross-nationally, and surveys on R&D output occur in both developed and developing countries.

Furthermore, OES data from BLS can be used to capture employment and wage information among STEM workers. These data can be used to determine the STEM occupational intensity of each industry, which can be further linked to R&D data to classify advanced industries. Linking OES data with external data products, such as ABS data, would provide a direct measure of employment and wage outcomes in advanced industries. Additionally, ABS includes a direct measure of whether technology producing industries have reported increasing or decreasing the number of skilled and STEM workers. This information can be used to assess the hypothesis that technology-producing firms have different labor demand patterns than technology-using firms.

³⁶ Examples: Mexico (http://en.www.inegi.org.mx/temas/ciencia/default.html#Informacion_general), New Zealand (http://datainfolplus.stats.govt.nz/Item/nz.govt.stats/4394653f-7947-487b-b0c7-dd48edb45822?_ga=2.205488272.184622000.1563650505-619231245.1563650505), and U.K. (<https://www.ons.gov.uk/economy/governmentpublicsectorandtaxes/researchanddevelopmentexpenditure/methodologies/ukbusinessenterpriseresearchanddevelopmentsurveyqmi>).

³⁷ "The OECD Analytical Business Enterprise Research And Development (ANBERD) Database." Retrieved from: http://www.oecd.org/sti/inno/ANBERD_full_documentation.pdf.

³⁸ "UIS Questionnaires." Retrieved from: <http://uis.unesco.org/uis-questionnaires>.

3.4.4 Skill and Labor Demand

Description of Measures. The primary measure of labor demand is employment growth overall and by occupation. As noted in the literature, skill demand is a more nuanced measure of labor demand that is necessary to assess the hypotheses related to skill-biased technological change and capital-skill complementarity. Key measures of skill demand include employment growth by skill and hours worked per week by skill.

Historically, skill has been operationalized as the level of education with college-educated (skilled) workers experiencing higher demand over time compared with those without a college degree. This measure, however, is an imperfect proxy since some workers without a college degree have valuable skills and the skill profile varies significantly even between those with a college degree. Notably, there are a number of skills proposed in the literature that may determine the impact of new technologies on the labor market. These skills can be broken down into different proficiencies, such as literacy, numeracy, and the ability to use computers, as well as foundational skills, such as the ability to collaborate, communicate, and solve problems.

Existing BLS Data Products. BLS collects extensive data on labor and skill demand in the U.S. Beyond the employment and wage data collected from establishments by the OES program, the Occupational Requirements Survey (ORS) collects establishment-level data on the physical demands, environmental conditions, education, training, and experience of jobs in the U.S. It also provides key data on the cognitive and mental requirements for specific occupations. The ORS program surveys approximately 25,300 establishments each year, focusing on the requirements of specific occupations from the employer's perspective. This data captures key measures of skill, such as education, pre-employment training (including certifications and licenses, prior work experience, and on-the-job-training).

BLS also collects individual-level data through the CPS. The CPS is conducted monthly by the Census Bureau for BLS. The CPS provides information about employment, unemployment, hours of work, earnings, and people not in the labor force. It produces 60,000 completes a month from eligible households using a combination of live telephone and in-person interviews with household respondents. The CPS dataset includes national, state, and local employment statistics on a monthly and annual basis. It also includes employment by industry and occupation, which can be used as measures of labor demand. Additionally, the CPS includes measures of skill supply, such as employment by level of education.

One gap in the CPS is that it is difficult to use these data to measure supply-demand mismatches by skill level. Researchers can observe variation in labor market statistics (such as the unemployment rate, annual job growth, or earnings growth) across educational or occupational categories, but cannot determine if employers are having difficulty filling positions by occupation or educational category, or how the supply of new graduates in a given field compares to demand.

Job vacancy data could play this role, particularly when coupled with more information on skills demanded for the entire stock of jobs. However, the BLS job vacancy product (JOLTS) only reports data by broad sectors and does not collect job vacancies by occupation. Moreover, the JOLTS survey does not collect any data on hiring difficulty such as the length of time an occupation has been posted or what percentage of vacant jobs have been filled. These are important measures of labor demand in the theoretical economics literature, but scholars have had to use commercial non-BLS data products to try to approximate these constructs (Rothwell 2014a; Hershbein and Kahn 2018).

Additionally, skill demand can be evaluated using BLS National Longitudinal Surveys (NLS). In particular, the NLSY97 tracks educational attainment, training activities, and employment outcomes among a nationally representative sample of approximately 9,000 individuals, who were 12 to 16 years old as of December 31, 1996. In the first round of data collection, each youth and one of that youth's parents received hour-long personal interviews. Each youth continues to be interviewed on an annual basis. These surveys include measures of educational attainment, as well as cognitive and non-cognitive tests, which generate valuable indicators of math, verbal, and other skills (using scores from the SAT, ACT, and ASVAB, and PIAT). This information from cognitive tests, in addition to educational attainment and training, can be used to evaluate the importance of individual skills on labor market outcomes over time.

Finally, the National Compensation Survey (NCS) produces indexes measuring change over time in labor costs and the level of average costs per hour worked. One concept captured in this data product is work level,³⁹ which categorizes certain aspects of a job to specific levels of work with assigned point values based on knowledge, job controls and complexity, contacts (nature and purpose), and physical environment. This information is linked to total compensation, wages and salaries, and details about benefits. Therefore, as a proxy measure of both skills and general tasks performed, job-level data can provide additional insights into the relationship of skills and tasks on labor market outcomes.

External Data Products. Beyond BLS measures of skill demand in the U.S., O*NET data, as noted in the previous section, can be merged with the occupational categories collected in the CPS, OES, and other BLS data products. O*NET data on skills, which includes basic skills, complex problem-solving skills, resource management skills, social skills, systems skills, and technical skills, can then be used to determine the demand for individuals with particular skills.

Section 3.4.5 includes a full discussion of how skill level has been linked to tasks in O*NET and other external datasets.

Additional measures of skill demand are provided at the international level by the OECD's PIAAC. This international survey is conducted in over 40 countries and sub-national units around the world,⁴⁰ including the U.S., where it is referred to as the International Survey of Adult Skills. The survey measures adults' proficiency in key information-processing skills—literacy, numeracy, and problem-solving—and gathers information and data on how adults use their skills at work.⁴¹ For further discussion of task information in this data product, see Section 3.4.5.

Analysis of Data Gaps. While data on labor demand is widely accessible at both granular geographic and temporal units, the availability of skill data remains the primary data need to be addressed before fully evaluating skill- and task-based models of technology's impact on labor market outcomes. Traditional proxies for skill, including education and income, lack the nuance needed to fully capture the role skill-levels play in labor markets, particularly given the growing number of tasks that can be completed through digitization, AI, and automation.

There are two primary data gaps related to skill demand. The first is occupational demand data and reporting that includes occupations and the ability to differentiate occupations by skill level. The CPS allows BLS to generate statistics on job growth and the unemployment rate by occupation, but not constructs more closely linked to demand flows, such as new hires or job vacancies. The JOLTS program does not currently include measures of occupation, skills, or tasks performed, which would be necessary to account for the impact of new technologies on labor market outcomes.

³⁹ Work level and critical jobs tasks are also captured by ORS.

⁴⁰ Australia, Austria, Belgium (Flanders), Canada, Chile, Cyprus, Czech Republic, Denmark, Ecuador, Estonia, Finland, France, Germany, Greece, Hungary, Indonesia, Ireland, Israel, Italy, Japan, Kazakhstan, Lithuania, Mexico, Netherlands, New Zealand, Norway, Peru, Poland, Russia, Singapore, Slovakia, South Korea, Spain, Sweden, Turkey, United Kingdom (England and Northern Ireland), and United States.

⁴¹ Administered by National Center for Education Statistics (NCES). The first round of data collection was conducted from April 2011 through 2012 with a nationally representative household sample of 5,000 adults between the ages of 16 and 65. Additional rounds of supplemental data collection were administered in 2014 and 2017, including U.S. PIAAC Prison Study of 1,270 adult inmates in federal and state prisons.

The second gap is having expanded access to data on skills to distinguish workers by occupation in job-specific skills (e.g., knowledge, ability) and/or general skills (e.g., cognitive and non-cognitive ability). In particular, data are needed on both technical proficiencies (e.g., computer skills, coding, and statistical skills), as well as harder-to-quantify skills such as the ability to collaborate, provide leadership, solve problems, and communicate effectively.

Some sources of job-specific skills (e.g., O*NET) lack the documentation and consistent production schedule to track occupational changes in skill over time. Furthermore, recent efforts to collect more granular information on cognitive and non-cognitive ability (including both technical proficiencies and interpersonal skills) are useful, but lack the timeliness needed to measure the rapid changes caused by these new technologies. For instance, PIAAC has only completed one cycle of data collection (2012-2017), and the second cycle is scheduled for 2021-2022. Therefore, expanded measures of skill and the task composition of jobs by BLS are needed on more frequently conducted surveys with clear production schedules in order to provide the necessary data to evaluate technology's impact on labor market outcomes.

3.4.5 Allocation of Tasks

Description of Measures. Task-based frameworks have been included in this analysis because some have argued that labor market outcomes are influenced not by an individual's skill, but by the tasks he or she completes during work (i.e., routine manual and cognitive tasks). Given the interrelatedness of tasks (roles, responsibilities, and activities performed by a worker) and skills (the ability to complete tasks), measures of tasks and individual skills are often discussed in tandem. However, in this section, we specifically address the allocation of tasks performed by workers in specific occupations.

Existing BLS Data Products. As described in Section 3.4.4, the ORS collects key data on the physical conditions of specific occupations, as reported by employers. The survey provides information on physical tasks performed at the occupational level, such as reaching, lifting, kneeling, walking, and using a keyboard. While this dataset provides useful information on physical occupational requirements, some of which could be replaced by robots, it does not publish cognitive tasks, which may also be performed by automation, AI, and digitization. ORS currently collects open-ended text on tasks performed, which may capture both physical and cognitive tasks, but it is not currently published and would require additional text analysis to provide reliable information at the occupational level.

External Data Products. O*NET provides a task profile for each occupation, as well as the frequency of tasks performed for each occupation. To develop these profiles, job incumbents and occupational experts are provided a list of tasks and asked to rate each one on the importance and frequency in their job. Once this list was developed, new tasks may be recorded and added to the list, but the methodology can be slow to document emerging tasks, leading to a lag in data.

Work activities or tasks in O*NET are categorized as information input, interacting with others, mental processes, and work output. For instance, under work output, possible tasks include controlling machines and processes, documenting/recording information, interacting with computers, operating vehicles, and repairing mechanical equipment. This information can be analyzed to determine if occupations, composed of particular tasks (e.g., manual labor), are subject to increased or decreased labor demand. This approach has been widely adopted in the academic literature. However, the process for producing O*NET data is not well documented and does not have a clear production schedule. Therefore, users are unable to assess how occupations change over time in their task composition and individuals learn new skills over the course of their careers. Since O*NET does not document these changes over time, this data product has limitations in its ability to link skill level and tasks performed to labor market outcomes documented in BLS data products.

As noted in the previous section, the OECD's PIAAC provides data needed to measure skill demand by tasks performed at work. While the survey measures adults' proficiency in key information-processing skills, it also asks questions about how often an individual's job usually involves particular tasks. These tasks range from interpersonal tasks, such as negotiating and influencing people, to information communications technology (ICT) tasks, such as working in spreadsheets or using a programming language. This information is then used to measure individual-level skills, namely numeracy, literacy, and problem-solving in technology-rich environments. Each dataset also includes detailed employment data, which can be used to model cross-national labor demand across different skillsets and tasks areas.

Statistical agencies, such as the UK's Office for National Statistics⁴², and researchers (Arntz et al. 2016) have used the PIAAC data on the task composition of jobs to estimate the likelihood of automation. These approaches differ from widely cited research, such as Frey and Osborne (2017), since they are able to use the task composition of jobs to relax the assumption that whole occupations rather than job-tasks are automated by technology. These research efforts using PIAAC have been used to estimate job automatability for 21 OECD countries, as well as the risk of automation by place of work and occupation in England.

Analysis of Data Gaps. Task accounting is one of the largest data collection challenges, given the potentially large number of tasks performed in specific occupations. However, current BLS products track what people do, such as the American Time Use Survey, and this expertise can be leveraged to focus specifically on how people spend their time during work activities. Such information would better allow researchers to track and analyze the task content of workers overall and by occupation. For example, it would be useful to know how much time administrative assistants spend on scheduling meetings and observe how the proliferation of meeting scheduling software might correlate with this task. Likewise, if the share of time devoted to one type of task decreases, it would be useful to know what new tasks workers might be spending more time on.⁴³

Existing data sources, such as the open-ended ORS task data (see Section 4.2.2A for a full description of this data), may be leveraged, but would require coding and analysis to be usable. These ORS data may be particularly useful since existing endeavors, such as O*NET, do not frequently update their task profiles and may not capturing emerging tasks. As already noted, PIAAC lacks the timeliness needed to measure the rapid changes in task profiles caused by these new technologies. Not only are there large gaps in the data collection schedule, preventing an up-to-date view of how individuals with particular skills performing a given task are affected by new technologies, the relevant tasks included on the instrument may not be current or comprehensive. Therefore, measures of task composition of jobs by BLS are needed on more frequently conducted surveys with clear production schedules. Additionally, the use of self-reported tasks performed from the ORS may identify emergent tasks not captured by existing external products that use fixed task profiles.

3.5 Summary of Major Gaps in Existing BLS Data

Given this review of gaps in current BLS data products and external data collection efforts, some primary opportunities have been identified to help BLS establish a more thorough understanding of how new technologies impact labor markets. These major gaps can be summarized into four main categories.

The first major data need centers on the demand for skills. BLS currently has limited information on real-time demand for work beyond macro-economic measures, such as total employment and employment by sector. This gap may be addressed by revising JOLTS to include occupations, as well as to potentially add or link data on skills, tasks, or hiring difficulties to better measure the heterogenous effects of new technologies on labor market outcomes.

⁴² "The probability of automation in England: 2011 and 2017." Retrieved from: <https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/articles/theprobabilityofautomationinengland/2011and2017>.

⁴³ Bureau of Labor Statistics, American Time Use Survey, <https://www.bls.gov/tus/>

Assessing the Impact of New Technologies on the Labor Market

Second, BLS and other statistical agencies provide extensive employment and wage data across education, age, and occupation categories that can be used to estimate skill demand among different categories of workers. However, there is no current data collection that can link labor market outcomes with cognitive ability, non-cognitive ability, knowledge, or job-specific skills, except for the NLS. Furthermore, there is no standardized, timely accounting of occupational tasks, which remains a significant gap when trying to better understand the impact of new technologies on labor market outcomes. Additionally, there is no longitudinal data on how the tasks performed by individual workers change over their careers. Potential solutions include adding task-related questions to the annual NLSY97 survey. Additionally, a selection of skills and task-related questions from PIAAC may be added to existing surveys, or the current PIAAC may be conducted more frequently in the U.S. This has the added benefit of being comparable to other countries. Further, ORS data may be leveraged, and O*NET may be revised to include a more transparent methodology and regular production schedule, which would enable better use of the data to track the evolution of skills and tasks performed over time, thus enabling a clearer linkage to labor market outcomes.

Third, current business surveys lack the granularity needed to determine the impact of new technologies on labor market outcomes. Surveys in the U.S., such as the ABS, ask general questions about AI and robotics, but do not provide detailed information about the types of technology used or the tasks performed by individual technologies. This makes it difficult to determine whether these technologies are reinstating or displacing labor. Adding questions to the ABS to better categorize adopted and purchased technologies, and making the microdata more easily available to BLS, would help address this key data gap. This is particularly important because current standalone data products do not allow for a full assessment of new technologies' impact on labor outcomes at the occupational level. By linking establishment-level data, particularly between Census and BLS data products, researchers and policymakers will be able to directly link technology adoption at the establishment level to the demand for skills and tasks performed by occupation.

Lastly, BLS does not currently produce any public statistics that estimate particular occupations' vulnerability to automation. Given the central importance of this outcome and the extensive academic research on this topic, standardized estimates of this nature would be valuable information for businesses, workers, researchers, and policymakers. Statistical agencies in other countries, such as ONS in the U.K., have adapted existing estimation techniques from the academic literature (Frey and Osborne 2017). However, there are many limitations to this approach and more sophisticated techniques developed by BLS, using direct measures of tasks performed by humans and machines within individual businesses, would improve the quality of estimates related to the potential displacement or reinstatement of labor generated by increased automation.

4. Identifying and Filling the Gaps in BLS Data Products for Measuring the Effects of Technology on the Labor Market

4.1 Introduction

In the previous section, we reviewed what international and domestic statistical agencies are doing to measure the key theoretical constructs linked to the labor market effects of technology.⁴⁴ In that review, we described several areas where data gaps are large relative to the theoretical importance of the constructs.

That section identified several major data gaps:

1. *The BLS has limited information that can be used to track demand by occupation and task. The Current Population Survey (CPS) allows BLS to generate statistics on job growth and the unemployment rate by occupation, but not measures more closely linked to demand flows, such as new hires or job vacancies.*
2. *Moreover, beyond occupations, BLS does not report data that allow for analysis of labor market outcomes by cognitive and non-cognitive skills or tasks performed. Scholars who have performed these analyses are generally forced to rely on other data sources, which do not follow BLS' more rigorous methodological requirements.*
3. *Third, business surveys from BLS and other agencies neither collect nor report detailed information about the types of technology used or the tasks performed by individual technologies, complicating efforts to determine whether the technology is displacing or spurring demand for labor.*

This section examines those constructs with a view to filling those data gaps. The goal is to develop data collection strategies that the BLS could consider to help further the understanding of how technology is affecting and is likely to affect the labor market. Many of the constructs mentioned above are closely related to data that the BLS already collects, and so this section attempts to capitalize upon existing data and sample sources as much as possible. Quantifying the costs of proposed changes in data collection is outside the scope of this project, but we draw upon publicly available data on the costs of operating various BLS surveys to provide a qualitative assessment of how our proposed strategy for data collection will affect BLS costs.⁴⁵

4.2 Topic Areas With Identified Gaps

This paper condenses the gaps identified from the previous section into two domains:

1. the labor market for occupations, including the tasks they involve and skills of workers who perform them
2. the cataloging and measuring of tasks performed by humans and machines

For each of these topics, we present a data collection strategy that considers a qualitative assessment of the trade-off between expected collection costs and the quality and depth of information collected. We also evaluate alternative data sources and collection efforts, including partnering with other agencies.

4.2.1 The Labor Market for Skills, Occupations, and Tasks

This section identifies three main gaps: 1) Reliable trend measures of occupational wages and employment; 2) Labor demand flows by occupation; 3) Comprehensive measures of skills that can be linked to workers in each occupational category and the types of tasks that workers in those occupations perform.

⁴⁴ "Measuring the Impact of New Technologies on the Labor Market: Technical Paper" Submitted to U.S. Department of Labor, Bureau of Labor Statistics on October 16, 2019 by Gallup, Inc.

⁴⁵ "FY 2020 Congressional Budget Justification Bureau of Labor Statistics,"

<https://www.dol.gov/sites/dolgov/files/general/budget/2020/CBJ-2020-V3-01.pdf> (accessed October 17, 2019).

Reliable trend measures of occupational wages and employment.

Since the second half of the 20th century, a widely held view in the economics literature is that workers with higher levels of education have disproportionately benefited from the introduction and diffusion of new technologies, particularly those related to information and communication. Partly as a result of this and other aspects connected to skills, many organizations and government agencies have an interest in understanding the relative supply and demand for skills in the labor market, and concerns over the effects of new technology elevate that interest still further. In order for governments, educators, policy organizations, and employers to know whether there is a shortage or surplus of workers with different skills sets and respond appropriately, they must have data related to the supply and demand of workers with various skills.

One relevant metric is the unemployment rate, which is available through the CPS for the nation and individual states for all workers and by education and by occupation.⁴⁶ Workers who are unemployed are asked to list their last occupation. These data provide some indication of the short-term problems that workers with different skills face in finding jobs.⁴⁷ They also offer insight into longer-term issues with respect to how different occupations and industries face persistent shocks to supply and demand, through trade, technology, immigration, the aging of the workforce, or other factors. If there are occupational groups with consistently high unemployment rates, it may indicate that workers in those fields need to be retrained for occupations with lower unemployment rates.

However, unemployment rates are insufficient measures of shortage or surplus and do not provide adequate information to guide workforce development investments. Research on the macroeconomics of job searching finds that both the number of vacancies and unemployed workers is relevant to understanding the labor market and the wage-setting process (Blanchard and Diamond 1989). An increase in the unemployment rate for a given occupation may be the result of a drop in the number of vacancies, a spike in the number of job searchers (due to new graduates or mass layoffs), or a change in how well workers are matched to vacancies (e.g., geographic or skill mismatch). Likewise, a low unemployment rate for a given occupation does not necessarily imply that universities or training institutions should increase the number of students or trainees for the relevant field, for there may be few vacancies and the low unemployment rate could result from the small number of job searchers previously employed in that field.

When considering potential mismatches in labor supply and demand, relative compensation growth and net changes in employment by occupation have also been used by economists, as noted long ago in Blank and Stigler (1957).

Employment and compensation data are available by occupation through three BLS products, but research on occupational changes is nonetheless limited. The CPS, the National Compensation Survey (NCS), and the Occupational Employment Statistics (OES) collect occupational data, but they each have strengths and limitations that make comprehensive trend analysis difficult.

⁴⁶ At the national level and many lower levels of geography, the Census Bureau's American Community Survey is available to measure data at the occupational level, such as income, unemployment, and labor market participation. Its large sample size make county and metropolitan statistical area measures feasible at the occupational level, depending on the number of people working in that occupational category.

⁴⁷ Labor Force Statistics from the Current Population Survey, "A-30. Unemployed persons by occupation and sex," available <https://www.bls.gov/web/empsit/cpseea30.htm> (accessed September 11, 2019).

Of these, the CPS has the most detailed respondent information but because it is designed to be representative of households rather than occupations it does not have enough responses to report reliable summary data for relatively rare occupations. Moreover, since occupations are self-reported by workers, there are concerns about the accuracy of occupational categories, relative to the OES data, which are reported by employers. The OES, by contrast, provides comprehensive salary and employment-level information for every occupation in the SOC system (over 800), but because it surveys only currently employed workers, one cannot use it to understand unemployment outcomes or job search behavior by occupation. The NCS provides useful information on the value of benefits, job controls and complexity, knowledge, supervision, the work environment, and other job characteristics, but is not meant to be used to track trends in labor market demand by occupation. Several features of the OES data collection and reporting process limit its utility for calculating growth in wages or employment by occupation, as the OES website states:

Challenges in using OES data as a time series include changes in the occupational, industrial, and geographical classification systems, changes in the way data are collected, changes in the survey reference period, and changes in mean wage estimation methodology, as well as permanent features of the methodology.

Our discussions with BLS staff indicate that these issues could be at least partly resolved with investments in changes to sample design and estimation methodology, along with new computer software and data processing. These issues have also been discussed by BLS staff (Dey, Miller, and Piccone 2014; Dey Piccone, and Miller 2019). The result would be a nearly comprehensive database of occupational changes. The major remaining limitation is that the OES excludes the self-employed, owners, and partners in unincorporated firms, and household workers. Data from CPS could be used to fill this gap.

4.2.1A. Summary recommendation. We recommend that BLS redesign data collection software for the OES or make any other design changes to support time series estimates of occupational employment and compensation, while drawing from any other relevant data sources (such as the CPS) as needed.

4.2.1A Qualitative assessment of tradeoff between costs and data quality. This proposal—to implement redesigns to the OES and create occupational time series information—requires no additional data collection. However, it would require additional labor hours from BLS staff. Overall, we regard this as a low cost way of providing higher-quality data to the public.

Labor demand flows by occupation.

Even with complete data on occupational changes in employment levels and wages, BLS would still be without a strong measure for demand flows by occupation. Job hires and job vacancies by occupation provide demand flow measures that could readily be aligned with data on supply flows, such as training by field of study and demand for foreign workers by occupation. These supply data are available outside of BLS and can be merged using SOC codes to BLS products.⁴⁸

The economics literature provides a strong theoretical rationale for analyzing vacancy and hiring data when assessing supply, demand, and matching of workers to jobs (Chavrid and Kuptzin 1966; Blanchard and Diamond 1989; Abraham 1987). Along these lines, BLS economists have written that job openings data “can help policymakers assess the state of the labor market and determine imbalances between the supply of and demand for labor” and that these data are needed to determine if unemployment is structural or cyclical (Mueller and Wolford 2008).

⁴⁸ U.S. Citizenship and Immigration Services, H-1B Employer Data Hub Files, <https://www.uscis.gov/tools/reports-studies/h-1b-employer-data-hub-files> (accessed October 16, 2019); National Center for Education Statistics, Classification of Instructional Programs, <https://nces.ed.gov/ipeds/cipcode/resources.aspx?y=55> (accessed October 16, 2019).

As further evidence that occupational-level vacancy data are valuable and feasibly collected, it is worth noting that the Workforce Investment Act of 1998 describes how states must use federal funding for displaced workers, and it stipulates that states must provide job vacancy data and information on the local occupations in demand as a core service.⁴⁹ Some states have met this requirement by conducting their own job vacancy surveys, while others purchase data from private providers, which use web-scraping techniques (e.g., The Conference Board).⁵⁰

The relevant state agencies of Maine, Utah, and Oregon have recently conducted their own job vacancy surveys. They directly surveyed employers and asked them to list information by job title on the number of positions for which they are actively recruiting, the duration of the vacancies, experience, and educational requirements, whether the job is permanent or temporary, whether the position is difficult-to-fill, and, if difficult, why.⁵¹ They release summary data at the occupational level. The BLS could potentially work with these and other state agencies to develop new occupational hiring and vacancy metrics in partnership with them, but presently, states use highly varied methods, so we regard coordinating the methods and funding across states for this purpose to be impractical in comparison to BLS collecting the data itself.

The JOLTS provides data on vacancies. However, its small sample frame (approximately 16,000 businesses) and methodology do not allow for analysis by occupation or industry.⁵² Thus, it is not possible to construct data on the number of vacancies per unemployed worker by occupation—a key theoretical construct in the job-matching literature that could guide workforce development policy better than only using unemployment or national measures of vacancies per unemployed worker. Below, we propose changes to JOLTS to measure hiring and job vacancies by occupation.

To acquire the necessary coverage by occupation, we propose that JOLTS include a supplemental module once per year to collect occupational data from respondents. Annual reporting would allow policymakers (e.g., Congress and members of the Federal Reserve Board) to monitor the labor market to observe how macroeconomic conditions are affecting demand for workers across occupational classifications. This additional module to JOLTS could oversample employers with a high likelihood of having workers in smaller occupational categories. Revising this sampling frame would require substantial effort and costs on the part of the BLS staff. The final occupational sample size should be large enough to be able to report all minor (three-digit) SOC groups, though BLS should consider what the costs would be for more detailed reporting.

It is worth noting that providing occupational details would pose a moderate additional burden on employers. However, we believe employers will not find this to be an unreasonable request. State governments frequently require employers to provide data on job hires and vacancies, as noted above. Moreover, employers are well aware of the roles for which they are recruiting or have recently hired and certainly capable of providing a brief description of the role, so the burden should be regarded as light compared to requirements for providing detailed financial information.

⁴⁹ Workforce Investment Act of 1998, 29 USC 2864, available <https://www.govinfo.gov/content/pkg/PLAW-105publ220/pdf/PLAW-105publ220.pdf>. The relevant section refers to the use of funds for employment and training activities.

⁵⁰ California and Alabama, for example, provide job vacancy data from The Conference Board, see [https://www.labormarketinfo.edd.ca.gov/data/help-wanted-online\(hwol\)/online-job-ads-data.html](https://www.labormarketinfo.edd.ca.gov/data/help-wanted-online(hwol)/online-job-ads-data.html) (accessed October 7, 2019); <http://www2.labor.alabama.gov/WorkforceDev/Default.aspx#HWOL> (accessed October 7, 2019).

⁵¹ Maine Center for Workforce Research and Information, <https://www.maine.gov/labor/cwri/jvs/>; State of Oregon, Employment Department, “Oregon’s Current Workforce Gaps,” <https://www.qualityinfo.org/documents/10182/13336/Oregon%27s+Current+Workforce+Gaps?version=1.1>; Mark Knold, “Difficult-to-Fill Jobs Study: Brief Report” Utah Department of Workforce Services, available at <http://www.dws.state.ut.us/wi/pubs/reports/difficulttofill.pdf>

⁵² Bureau of Labor Statistics, <https://www.bls.gov/jlt/jltsampl.htm> (accessed September 11, 2019).

It would significantly increase the cost associated with the JOLTS project because of the need to redesign the existing survey and sampling frame, the larger sample size requirements, and backend data coding and management. Overall, these changes would yield results that were comparable to the existing JOLTS data and to state vacancy surveys and fill an important gap in federal data provision. These data would provide the best measures available on demand flows at the occupational level and the only national publicly available data on the topic. Doing so would enrich and inform the analysis of supply-demand mismatch, job training needs, and workforce development planning.

To further assess the costs associated with this proposal, we examined the various methods that BLS uses to collect occupational data from businesses. The JOLTS survey collects data from 16,000 businesses, the NCS collects data from 11,400 establishments, the Occupational Requirements Survey (ORS) surveys approximately 25,300 establishments, and the OES program surveys 180,000 establishments. Aside from JOLTS, these three other surveys report occupational summary data. The largest, the OES, reports summary data for over 800 occupations at national and metropolitan levels, which is far greater detail than we would propose for a supplement to JOLTS. OES is also required to produce occupational data by industry. The NCS and ORS are able to achieve occupational coverage by oversampling large establishments. JOLTS does not use this method, but it could be adopted for the JOLTS supplement.

An alternative or complementary way to collect job vacancy data would be to follow private sector companies and scrape them from the internet (including job boards and company websites).

There are several advantages of this approach: 1) It would eliminate the respondent burden, 2) the marginal costs of additional data collection would be smaller than survey-based approaches, and 3) meta-data on skills, certifications, job tasks, and other job requirements listed on vacancies could enrich understanding and context.

One challenge on the cost side is that web-scraping a diverse and dynamic set of internet websites requires substantial upfront research and development costs, as well as ongoing maintenance and quality checks. To reduce this burden, BLS could outsource data collection and maintenance to an external partner. The resulting database could be used for internal R&D, but there would be significant challenges in presenting the data to the public.

Challenges with reported job vacancies from web-scraping include the fact that the universe of online job vacancies does not match the universe of all vacancies and is likely biased toward coverage of professional roles in large companies. Occupations in construction and restaurants (e.g., brick masons, general laborers, food service, and cooks) are under-posted online relative to what state labor agencies estimate through traditional survey collection methods (Rothwell 2014a). There may also be concerns about the accuracy of online postings, in terms of whether they represent a legitimate vacancy or whether a skill is required or preferred, or a duplicate posting, as different companies may have different policies related to postings. BLS research shows that larger and older companies are more likely to post their job vacancies online (Dalton, Khan, and Mueller 2019), and other evidence shows that the skill requirements of online job vacancies may be sensitive to the business cycle (Hershbein and Khan 2018). Finally, web-scraping provides no measure of hires, which is an important theoretical concept. The gap between hiring and vacancies is an important measure of hiring difficulty (or the efficiency of matching).

On the other hand, BLS and academic economists alike have matched public administration and survey data to web-scraped data, yielding valuable insights about hiring difficulty (Dalton, Khan, and Mueller 2019). These techniques could also be used to construct sample weights that could be used to more closely align internet-based data to administration or survey-based distributions.

We believe BLS should periodically revisit the merits of web-scraping as a form of data collection, as the prevalence of internet advertising increases and the technologies used to collect the information improve, but for now, we recommend that BLS focus collection efforts using more traditional methods (surveys). This will provide consistent data on hiring and vacancies and, given the sampling biases of internet advertising, result in higher data accuracy than web-scraping methods, and the agency can continue to purchase provided data to experiment with and use to inform its traditional efforts in data classification and collection.

4.2.1B Summary recommendation. We recommend that BLS expand the JOLTS to include occupational details annually. This could be thought of as an annual supplement to the existing JOLTS survey. In order to get occupational coverage without greatly expanding the number of businesses surveyed, we recommend that the JOLTS supplement use design methods closer to the NCS and ORS and oversample large establishments.

4.2.1B Qualitative assessment of tradeoff between costs and data quality. Occupational data on hiring and job vacancies are routinely collected by state governments and the BLS could collect high-quality data from employers that would be highly valuable to understanding skill- and occupational-based patterns in demand. The expanded sample size and redesigned methods would incur substantial costs to BLS, expanding the budget for the overall JOLTS project by a significant amount.

Comprehensive measures of skills that can be linked to workers in each occupational category and the types of tasks that workers in those occupations perform.

The final gap in this section pertains to the lack of comprehensive skills measures that can be linked to occupational analysis and labor market outcomes. Presently, occupational-based data on skills in BLS products is limited to a handful of measures that do not adequately capture the constructs related to skill.

Education and occupations are insufficient measures of skills to inform workforce investment strategies. There is tremendous heterogeneity in labor market conditions and job functions for workers with the same level of education. Occupations, as predictable bundles of tasks that draw upon various skills, provide more detailed information and often imply a required level of formal education, but even workers in the same occupation may perform different tasks at different skill levels—and are rewarded accordingly (Autor and Handel 2013). Likewise, job vacancies are harder to fill when they mention skills associated with higher-salary offers (Rothwell 2014a), and differences in demand for those skills by geography and establishment explain, at least in part, why software developers in Silicon Valley are paid more than software developers in New York City (Rothwell 2014b).

Aside from compensation, meta-analysis from the industrial psychology literature documents a number of important facets of skills that predict job performance for a given occupation, including general cognitive ability, non-cognitive ability, job-specific knowledge (or expertise), and experience (Schmidt, Oh, and Shaffer 2016). This is consistent with rich evidence from the economics literature that both cognitive skills and non-cognitive skills play important roles in the labor market and can be measured reliably (Heckman and Kautz 2012).

With occupation-level summary data on the typical skills or range of skills (e.g., at the 25th and 75th percentiles of workers), educators, trainers, and workforce development officials could more easily fulfill their mission to guide workers into appropriate and realistic career paths. With occupational data available on cognitive ability, workers displaced by technology could take an exam that gauges how their skills compare to existing workers' skills, help them decide if they should receive additional education, and if so, for how long and at what level. In the absence of such data, it is difficult to know if a displaced production worker or administrative assistant has the appropriate skillset to pursue an alternate line of work. The occupational summary statistics (on skills and unemployment) produced from this research could be merged with vacancy data (from JOLTS or online data) to produce measures of vacancy-to-unemployment ratios for high- and low-skill positions along each dimension of skill (e.g., numeracy, literacy, technological sophistication, conscientiousness, emotional stability, etc.)

Across BLS and other publicly available data products, we find that the skill-related elements are well captured by the National Longitudinal Survey of Youth (NLSY) and a non-BLS product: the Programme for the International Assessment of Adult Competencies (PIAAC). Both of these surveys contain valuable measures of cognitive and non-cognitive skills, as well as, in the case of PIAAC, measures relating to tasks performed on the job. Both also classify respondent work by occupation. The key weakness of these surveys is that they are not large enough to produce viable summary statistics by occupation—at least beyond two-digit major occupation categories. The other BLS data products lack important data on skills. Since BLS controls the NLSY, we propose that the BLS collect data from a new cohort of youth that is large enough to calculate summary statistics by occupation as they age and enter the workforce.

These data could also inform other BLS data projects in future work. The ORS, for example, collects limited data on cognitive skills (e.g., problem solving, literacy, and “people skills”), but these items do not fully capture cognitive ability, nor the non-cognitive skills associated with employee performance and productivity.⁵³

The NLSY97 cohort already reports data on occupation, industry, and other job characteristics of respondents. We propose the addition of one more module to the NLSY questionnaire and the creation of a new cohort with expanded sample size.

The NLSY97 tracks 9,000 individuals who were aged 12 to 16 in 1996. We believe that is too small to provide meaningful occupational coverage for the future working arrangements of adults who are tracked over time. If a meaningful reporting threshold is 100 workers for a given occupation, we estimate that a sample size of 9,000 will yield an insufficient number of observations for several major (two-digit) occupational groups, including life, physical, and social science occupations, and most minor (three-digit) occupational groups, and only 3% of detailed (six-digit occupations). A new cohort sample size of 30,000 would allow complete coverage at the two-digit level and would cover the 56 largest minor occupational groups out of 100 and 17% of the detailed occupations. A cohort sample size of 50,000 would reach an additional 10 minor occupational groups (66 out of 100) and 23% of detailed occupations.⁵⁴ As in previous cohorts, the new cohort could be designed to survey individuals who are aged 12 to 16 in a future baseline year.

However, we would also recommend that formal assessments of cognitive and non-cognitive skills be collected again (once per decade but not every wave) as adults. As with the PIAAC, having detailed skill data for adults allows researchers to directly measure the cognitive and non-cognitive skills connected to jobs. This would allow for the likelihood that these skills change over time and are themselves influenced by tasks, training, and subsequent education attained beyond secondary school. That would be consistent with research showing the IQ tends to increase as a result of further education (Ritchie and Tucker-Drob 2018). Thus, responses to cognitive and non-cognitive assessment could be recollected in the NLSY97 and NLSY79 cohorts, which would provide rich information on how these traits evolve over one’s lifetime and how work context and tasks affect absolute and relative changes in skills.

⁵³ Bureau of Labor Statistics, Occupational Requirements Survey (ORS) Collection Manual, available at <https://www.bls.gov/ors/information-for-survey-participants/pdf/occupational-requirements-survey-collection-manual-082019.pdf> (accessed October 7, 2019).

⁵⁴ These calculations were made using the 2017 American Community Survey with data from IPUMS USA. Steven Ruggles, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek. IPUMS USA: Version 9.0 [2017 American Community Survey]. Minneapolis, MN: IPUMS, 2019. <https://doi.org/10.18128/D010.V9.0>. We limited the sample population to people between the ages of 25 to 45 for the purposes of calculating the percentage of individuals who work across occupational groups.

Additionally, for each of the cohorts, the BLS can continue to use a new module on the tasks performed by workers. The BLS recently added items from the Princeton Data Improvement Initiative (PDII) to the NLSY79 (wave 27) and NLSY97 cohorts (wave 18). This information helps fill in previously missing information on the tasks performed by workers (see Autor and Handel 2013).⁵⁵ We would recommend continuing to collect these data across a range of tasks and modifying the tasks module after reviewing the validity and utility of the data. Broadly, the collection strategy should include the task categories that have been informed by empirical and theoretical economics literature (such as routine, non-routine and cognitive, non-cognitive) as well as other categories of tasks that may or may not matter in the future, as new technologies affect the labor market.

The resulting database is filling a significant gap. Until recently, no data were available from the BLS or other U.S. agencies that provide researchers with detailed information on the skill requirements or task content of occupational families and specific occupations. The new task module allows for scholars to observe how tasks change over time for the same individual within or across occupational groups. It will further allow observations on the relationship between skills expressed in youth and tasks and occupations performed in adulthood. These changes enrich understanding about occupations, how skills change over time, which tasks change over time, and the range of skills that tasks and occupations require.

4.2.1C Summary recommendation. First, we propose that BLS launch a new NLSY20 cohort with a larger sample size (30,000 instead of 10,000) to achieve complete two-digit occupational coverage. Second, we propose that BLS include a new small module of items every year on the tasks performed by workers for the NLSY79 and NLSY97 and a new NLSY cohort, as respondents age—drawing on the new PDII module that was launched for recent waves of data collection. Third, every 10 years, the adults assessed in the original cohort as teenagers should be reassessed for cognitive and non-cognitive skills.

By cognitive skills, we are referring to validated psychometric measures of intelligence or cognitive ability. Examples include the PIAAC, the SAT, the ACT, the Armed Force Qualification Test, the Stanford-Binet, and the Wechsler Intelligence Scale for Children. By non-cognitive skills, we are referring to any other enduring traits that are not captured by cognitive ability, including elements of the Big Five personality traits of conscientiousness, emotional stability, and extroversion, as well as concepts like grit, self-control, or self-management. The concept of “soft skills” is a subset of non-cognitive skills. Reviews of the relevant scientific literature as it related to economic outcomes are discussed in Heckman and Kautz (2012), Borghans et al. (2008), and Schanzenbach et al. (2016). California school districts have begun to explicitly measure self-management skills and whether students have a growth mindset (West 2016).

In measuring cognitive and non-cognitive skills in the NLSY, we recommend that BLS prioritize instruments that have been validated in academic research, have been measured in previous waves of the NLSY, and that can be benchmarked against other studies or populations. There may be some tradeoffs in satisfying each of these goals. The PIACC, for example, is attractive because its results can be readily compared to test takers in other countries but relying on the PIACC instrument would eliminate the possibility of trend analysis with previous NLSY cohorts and make it impossible to observe changes across generations.

4.2.1C Qualitative assessment of tradeoff between costs and data quality tradeoff. This proposal would draw on the existing infrastructure and methods of the NLSY, while making minor changes to the questionnaire. The cost of expanding the sample size for NLSY20 would be substantial relative to the baseline costs but highly valuable in yielding actionable insights about the relationship between fundamental cognitive and non-cognitive skills and work. When combined, these data will show how skills and tasks co-evolve over Americans’ working careers and provide data at the occupational level that could be linked to other data sources. This would provide a richer understanding of the skill and task requirements of occupations—and ultimately how technology affects the supply and demand for skills and the tasks performed by workers of various skill levels.

⁵⁵ National Longitudinal Survey of Youth 1979, <https://www.nlsinfo.org/sites/nlsinfo.org/files/attachments/18108/7.Employer%20Supplement.html> (accessed November 22, 2019)

4.2.2 Cataloging Tasks and Who or What Performs Them

This section identifies two main gaps: 1) The need for a classification system of tasks that can describe tasks performed by humans and machines; 2) An inventory of tasks performed by humans and machines (or capital). Before elaborating on these proposals, we start with a general discussion of the broader topic.

One lesson gleaned from the recent literature review on technology and the labor market is that jobs are best thought of as bundles of tasks. In the theoretical model of Acemoglu and Restrepo (2019), tasks performed by machines are automated and technology creates a displacement effect when it takes over the performance of a task. Crucially, technology also creates a reinstatement effect, potentially, by increasing the number of tasks performed by labor, and a productivity effect by creating value if a task is performed more efficiently by a machine.

Given this framework, it is arguably impossible to fully understand how technology affects the labor market without a comprehensive catalog of tasks performed in the economy and who or what performs them. It is also important to understand the value of those tasks as they relate to establishment revenue, otherwise it is not possible to understand whether or not a machine or human is more productive at a given task or how the productivity of machines or humans is changing at the task or aggregate level. At this point, neither BLS nor any other entity publishes available data that aim to comprehensively document all the tasks performed in the economy or whether they are performed by humans or machines.

The closest thing to an inventory of tasks is published by O*NET. The Department of Labor's Employment and Training Administration publishes O*NET, which seeks to "populate and maintain a current database on the detailed characteristics of workers, occupations, and skills" (O*NET Resource Center 2018, A-1). O*NET's database includes information on the skills needed to enter an occupation and descriptions of the tasks performed by workers in those occupations. Specifically, O*NET provides "task statements" for 967 distinct occupations. The number of task statements per occupation ranges from four to 40, with a mean and median of 20 tasks.

The O*NET database provides a catalog of tasks per occupation and, in principle, allows researchers to track changes in the number and types of tasks performed by occupations over time. However, one challenge with O*NET is that its sampling and survey methodology complicates interpretation of the data. From 2008 to 2017, roughly 100 occupations out of nearly 1,000 were updated each year. This means, some occupations are always more up-to-date than others (O*NET Resource Center 2018). In fact, an analysis of the O*NET version 24.0 meta-data by occupation reveals a wide range of dates for each occupation, with respect to when the occupation was last updated. Some occupations were updated in 2019, while others have not been updated since 2006. It is unclear which occupations are selected for updates and when, but in an economy with rapidly changing technology, much may have changed since 2006 with respect to the tasks performed by an occupation.

It is also not entirely clear from the provided documentation how O*NET analysts identify tasks, though once identified, a random sample of workers in each occupation is asked to rate the relevance, importance, and frequency of performance for each task statement connected to that occupation.

We believe the O*NET task database provides a potentially useful starting point for creating a more comprehensive inventory of all tasks performed in the economy. The inclusion of importance, frequency, and relevance fields provides additional value for potential analysis. The O*NET process for updating tasks for an occupation is not entirely transparent, but documentation suggests that it entails detailed internet searches of job vacancies and search engines to identify the tasks associated with occupations as they emerge (Dierdorff and Norton, 2011). However, we believe the limitations of O*NET warrant going beyond its methods and processes to create a new database.

The need for a classification system of tasks that can describe tasks performed by humans and machines.

The first element of our strategy for this topic is to create a classification system of tasks.

The O*NET collection of tasks is useful but is nothing like a classification system in its present form. Tasks are defined by O*NET as occupation-specific, so there are no tasks identified across multiple occupations. For example, food preparation and service workers do not share any common tasks with loan clerks, even though the database lists both as having a task statement that includes “accepting payment.” The task statement of accepting payment is coded differently and includes different words beyond “accept payment” across the two occupations.

A task database must have tasks (rather than task-occupation combinations) as the fundamental unit. We propose that BLS staff draw upon a number of data sources to create a classification system with tasks as the unit of analysis.

1. O*NET task statements
2. ORS task statements collected by BLS staff in the field
3. Patent records
4. Job vacancy descriptions, using proprietary data from a private provider (e.g., Burning Glass, Conference Board, etc.) or data scraped by BLS staff
5. Summary statements and task descriptions from resumes or digital repositories of quasi-resumes (e.g., LinkedIn)
6. The O*NET Tools and Technology Database
7. Websites, marketing materials, and catalogs from industry associations that specialize in technology, machines, and software
8. Technical and popular journals covering new developments and advances in technology
9. Wikipedia and other decentralized user-generated web pages on technical subjects
10. Data on equipment purchases from existing and new Census Survey products
11. Producer Price Index (PPI)

Starting with O*NET task statements, BLS staff could analyze the text from these descriptive statements to classify O*NET task statements into tasks with common attributes, such that accepting payment from a client or customer is regarded as one common task. One method for doing this would be to follow Webb (2019) and use verb-object pairs, which he defines as “capabilities” in the context of text from patents.

This method could be used across any database that describes the tasks performed by workers or machines. Webb (2019) uses this method to describe the tasks of technology (using patent data) and workers (using O*NET task statements). O*NET’s task statements have unclear origins and are likely not comprehensive, and patent data are limited to a small subset of machines and technologies. For this reason, our strategy entails the collection using Webb’s method (or similar approach) across the databases noted above.

We regard these sources as mostly self-explanatory but several warrant additional comments.

The O*NET Tools and Technology database aspires to document every tool or technology that is essential to the performance of an occupation (Dierdorff, Drewes, and Norton 2006). It does not contain task statements for tools or technologies, but something like task statements could potentially be appended to it using the other sources described (including web-scraping).⁵⁶

⁵⁶ The Tools and Technology Survey has several attractive features but ultimately does not capture most automation technologies. The technologies themselves are classified using the United Nations Standard Products and Services Code, which imposes a rigorous structure to the specific technologies. Moreover, the procedures for creating and updating this database involve extensive internet searches for each occupation and review of job postings information and review by subject matter experts. A serious limitation of the database is that it is meant only to capture technologies used by workers to perform their jobs. Thus, it makes for a poor list of automation technologies. For example, under the occupation “cashiers,” the database does not mention self-checkout machines, of the type used in grocery or convenience stores. These machines directly compete with cashiers, but cashiers don’t necessarily need to use them. In fact, kiosks and self-checkout machines are not listed anywhere in the Tools and Technology database.

The Census Bureau's Annual Capital Expenditure Survey (ACES) in its current form or in a modified version could be a vehicle for collecting the information required to build a classification of machines.⁵⁷ In the ACES, companies are asked to list aggregated investments in topic areas such as machines and robots. The survey lists common tasks performed by robots, including assembly, cleaning, delivery, and inspection. A modified version of this survey could require a subsample of firms to list more detailed information about their specific equipment and the tasks it performs. Just as the ORS survey samples only a subset of an establishment's occupations for detailed data collection, an expanded ACES could do the same for equipment to minimize respondent burden. For example, some establishments could be asked to list all of their computer equipment. Other establishments could be asked to list all of their industrial robots, while others could be asked to list all of their software applications. Ideally, these data would be shared across agencies, so that BLS could fully utilize it. We will return to the issue of inter-agency data sharing below. Moreover, anticipating the additional data needs described below, we also recommend that an adapted version of this database be used to collect information on the values and tasks performed by each technology.

The PPI is a BLS product that entails the survey of 25,000 business establishments, receiving price quotes for 100,000 products. The product list would be a potentially useful source for building a more comprehensive database of machines and software (or automation technologies). One advantage of using the PPI is that it also contains detailed cost and product characteristic information. Internet searches could potentially help analysts fill in missing details about the tasks performed by the technologies. We consider expanding and repurposing the PPI to systematically collect data on the tasks of automation to be a relatively low-cost data effort to generate data on automation technologies that are available for sale. The chief issue, however, is that data from internet searchers, wholesalers, retailers, and market-based products, no matter how comprehensive, cannot fully create an inventory of automation technologies. Many technologies used in factories or businesses will be heavily customized, produced in-house, or otherwise not available for sale in the market. That is why the PPI could only be one of several sources used to build this database.

Across these sources, task elements (i.e., capabilities, verb-object pairs) could be derived from each of these sources and, ideally, matched to a relevant higher-level categorization (i.e., an occupation, machine, or software program). From there, BLS could contract with experts on tasks from industrial organizational psychology, industrial engineering, or other disciplines to devise a categorization scheme. For example, such a scheme might group tasks broadly into cognitive and non-cognitive elements, and within cognitive, distinguish between analytic reasoning (applying formal rules and theory) and non-analytic reasoning (memory, processing information), speech, persuasion, reading comprehension, and other elements.

The final product would be a system for categorizing every task (verb-object pair) performed in the economy. The classification system would need to be updated annually by reviewing some or more of the lists mentioned above to attempt to capture new sources of information.

4.2.2A Summary recommendation. We recommend that BLS create a standardized task classification system that would, in principle, allow for the categorization of any economically meaningful activity performed by a human, machine, or technology. Recognizing that the Census Bureau has more experience and detailed data collection history with machines, we believe the BLS and Census Bureau should partner under a formal inter-agency agreement to help build this database, with Census providing information from adapted versions of its capital and equipment surveys. Given security and confidentiality rules, as well as budget appropriation constraints, this agreement and partnership would likely require Congressional legislation.

⁵⁷ U.S. Census Bureau, 2018 Annual Capital Expenditures Survey [https://www2.census.gov/programs-surveys/aces/technical-documentation/questionnaires/2018/ace-1\(m\)_worksheet.pdf](https://www2.census.gov/programs-surveys/aces/technical-documentation/questionnaires/2018/ace-1(m)_worksheet.pdf)? (accessed September 3, 2019).

4.2.2A Qualitative assessment of tradeoff between costs and data quality. This would be an extensive research and development project, but one of great value to the U.S. government and its people. It would have the potential to influence statistical offices around the world and be the first step to developing a comprehensive catalog of tasks, which is our next proposal.

An inventory of tasks performed by humans and machines.

With a task classification system in place, the next step in our proposed data collection strategy is to measure the prevalence and value of tasks performed by humans and by machines.

The first goal involves detailed data collection at the occupational level. Fortunately, BLS already has several surveys that do this. Additionally, the ORS also collects data on the tasks performed by workers. At specific business establishments, BLS staff conduct structured interviews with managers and other relevant employees to collect information about the tasks performed by workers in preselected occupational groups within that establishment. BLS should examine the feasibility of classifying these task statements into the task classification system described above. If information is lacking, BLS could change its procedures to more purposefully collect verb-object descriptions.

We recommend that this research take two phases.

In phase one, BLS should develop a bottom-up list of tasks commonly performed by each occupation, by asking open-ended questions, as the ORS currently does. Currently, the ORS collects information on job tasks that are deemed critical tasks or that use at least 10% of a job's time. A critical task is defined in ORS documentation as "An activity workers must perform to carry out their critical job function(s). A task is critical when it is a required component of the critical job function(s)," where a critical job function is defined as the main purpose of the job.⁵⁸ ORS also collects information on the duration of performance for physical demands, with levels ranging from seldom (up to 2% of the workday), to occasionally (2% and up to one-third of the workday) to frequently (one-third to up to two-thirds), and constantly (two-thirds or more). ORS does not collect these duration data for tasks, however. Collecting duration data for tasks would increase the cost of but yield more comprehensive data, which would be necessary to understand how technology affects worker tasks. One scenario is that a new technology does not displace a critical job function (or one that takes up most of a worker's time) but does displace a secondary function. This may make workers more productive in the tasks remaining. This analysis requires data on duration. BLS should consider whether the 10% and/or critical-task threshold yields a sufficiently complete list of the relevant tasks performed by workers.

In phase two, BLS should use the results of phase one research to create a top-down list of tasks commonly performed by each occupation and present it to subjects in the ORS study for them to fill out. This part of the study would follow O*NET-like methods in that it would ask respondents to state whether the task is performed by that worker and what percentage of time the worker spends on each task. It would also allow open-ended responses, so that new tasks could be added, as they emerge in workforce.

To identify the value of these tasks, the most straightforward method would be to use the ORS to also record the hourly compensation (wage plus benefits) at the worker level. The value of each task could be the product of the duration (percentage of hours performed) and the hourly compensation. This form of data collection is our recommended approach because it will yield the most accurate way of linking value to tasks, without having to value each task subjectively, which we do not believe is feasible. A second-best approach would be to append occupational wage from the OES or CPS to occupations, but this would assume that all workers in the same occupation perform the same tasks, which is not the case (Autor and Handel 2013).

⁵⁸ U.S. Bureau of Labor Statistics, Occupational Requirements Survey, <https://www.bls.gov/opub/hom/ors/pdf/ors.pdf> (accessed October 21, 2019).

We regard the second half of this effort—creating an inventory of the tasks performed by machines—as the most challenging of all our proposals, yet one that holds tremendous promise and value for a better understanding of the labor market.

We believe that the Census Bureau is the U.S. agency best positioned to collect these data because of its history, expertise, and infrastructure regarding the measurement of equipment and capital expenditures, especially through the ACES, which collects detailed expenditure data every five years from approximately 46,000 companies with at least one employee and 30,000 companies with no employees.⁵⁹

As described previously, a survey, call it the “Tasks of Technology Survey,” could use the same sampling frame as the ACES, but limit the respondent burden by narrowing the data collection effort for each company. The ACES survey is estimated to take 2.57 hours to complete and is mandated by U.S. law. The burden involves “time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information.” The need for detailed technical data makes this a burdensome survey, so we propose breaking up components of it.

One component could be dedicated to robots. Presently, the form has a section called “Capital Expenditures for Robotic Equipment.” This section lists 18 tasks performed by robots but excludes automated teller machines [automated teller machines, CNC (or computer numerically controlled) machining equipment, and kiosks (defined as “stationary, consumer-oriented machines with a graphic interface and no visible moving parts”)]. Firms randomly assigned to this condition would list every robot, the value of each robot, and the tasks each robot performs. This would be a relatively minor increase in the respondent burden because these firms would likely have to access the same information to answer the current ACES; the only difference is that instead of reporting only aggregated totals, the firm would have to report line items by robot. A second group could be asked to list detailed expenditures for computer software, whereas other groups could be asked to list other automation technologies, drawing from the classification system created in the above proposal.

The ACES survey takes place at the company level, which may be necessary to acquire the most accurate cost and inventory details from senior officials; local plant managers may not have access to this information. On the other hand, local managers may have a better understanding of which technologies are being used and what tasks they perform. We propose that Census consider both approaches to collection.

Another challenge with these data would be in describing the value. We believe the costs of creating these technologies is the most practical way to understand the value and could best be captured in expenditures on the piece of equipment for each year, including research and development expenses, contracting expenses, monitoring expenses, maintenance, and repair. Data on depreciation should also be gathered, as it is currently on a more aggregate scale.

Ideally, economists could determine ways to translate these total data into something like a hypothetical hourly rental price for the technology that could be compared to hourly worker compensation. For many products, there will not be a market rental price that could be collected for this exercise, and such a price would be artificially low, in any case, because it would exclude maintenance, repair, and monitoring. In limited cases, data on the productivity of machines could be used to compare to worker productivity at the same task, but we believe this is unlikely to be fruitful because machines and workers will likely be doing slightly different tasks and machines will likely be working with workers, making it impossible to allocate the value-added contributions of each. Still, the expenditure data and wage data can be used for a number of meaningful analytic exercises to better understand how technologies and workers interact, compete, and complement each other across tasks and industries.

⁵⁹ Annual Capital Expenditure Survey, <https://www.census.gov/econ/overview/mu2200.html>

A Task Inventory Database—which would combine the Tasks of Technology with the ORS data on the tasks of humans—would be highly valuable in its own right as a snapshot of the economic activity machines and humans perform. When tracked over time, these data would allow BLS to report the productivity effect described in Acemoglu and Restrepo (2019), as the contribution of automation to the value created by the task, which they say is proportional to the cost savings of automation.

For the U.S. economy, the productivity effect of automating a specific task could be calculated (at least roughly) as the difference between the cost of a human (h) performing the task (T) in the previous period minus the cost of a machine (m) performing the task presently, summed across the number of times the task is performed in a given period (i). If the costs of human performance are larger, this could be thought of as the cost savings that accrues to the national economy from automation.

$$Productivity\ Effect_T = \sum_{n=1}^i (C_{T,h,t} - C_{T,m,t-1})$$

For this to be a meaningful measure, the inventory database would have to be nationally representative of the number of tasks performed, how frequently they are performed, and how expensive it is for workers and machines to perform each task. To give a stylized example, imagine an automated kiosk can perform an average of 15 checkouts per hour at a grocery store at the cost of \$1 per hour, which is the rental price per hour for using the kiosk, inclusive of monitoring costs, maintenance, and repair. A worker could perform that same number of checkouts per hour, but at a cost of \$15 per hour. The savings per year would be roughly \$41,000, if we assume that both the kiosk and worker operate during the same number of days (e.g., the kiosk is non-operational for maintenance and repair during part of the year). If this example held true at scale, the productivity effect of introducing 1,000 kiosks nationwide, would be roughly \$41 million.

This productivity effect could be translated into demand for labor in various ways. For this hypothetical grocery store, some of the savings from automation could be used to hire other workers (e.g., more florists, butchers, or pharmacists), another portion of the savings could go to executive compensation and shareholders, another portion could go to suppliers (e.g., organic dairy farmers), and another portion could go to customers in the form of price reductions.

Economists could work with the task inventory data and other government statistics (like National Accounts and Input-Output tables from the Bureau of Economic Analysis) to estimate how the productivity effect translates into labor demand at the firm level, by industry, and for the national economy. Data from the BLS Consumer Expenditure Survey could be used to estimate the macroeconomic effects of consumer surplus.

We believe the BLS could create a tasks of technology database without the Census Bureau, but the costs would be very high, and involve the following: the creation of an entirely new sampling frame, creating a new questionnaire, reassigning or hiring new staff, and fielding, coding, and maintenance of an entirely new survey product. By contrast, the Census Bureau's infrastructure and expertise is neatly aligned to perform this portion of the project, without having to radically change its current data products and methods.

Inter-agency cooperation and data access.

As described above, the BLS would create the task classification system and the human task database. However, the first would benefit from cooperation with the Census Bureau, and the complete task inventory would require the Census Bureau's participation, under our proposal. This creates two complications. The first is that BLS and Census have different pools of funding and different but partially overlapping missions. Moreover, confidentiality agreements with respondents and legal restrictions about disclosing data that could be linked to tax records may prohibit cooperation even when partnerships are considered mutually desirable between agency leaders. Therefore, new legislation would be necessary to fully fund these efforts as an inter-agency endeavor and to allow greater data access and sharing between the BLS and Census Bureau, consistent with privacy protections.

Interagency collaboration would also be of considerable value in allowing BLS economists to merge all elements of the task database, the Annual Social and Economic Supplement (ASEC), and other information from the Census Bureau with data from ORS, the OES, and other surveys. For example, BLS economists have published research showing that the OES' methods allow for a nationally representative sample of establishments to be tracked longitudinally (Day and Handwerker 2016). Thus, merging OES longitudinal data with the Tasks Inventory database would provide a highly valuable means of studying how investments in automation technologies in one period predict changes in the composition and compensation of occupations in the future.

4.2.2B Summary recommendation. We recommend that BLS create an inventory of the tasks performed by humans, using data and methods developed for the ORS but adapted to expand the range of tasks collected beyond those deemed critical to include those of secondary importance. We also recommend that the BLS partner with the Census Bureau to create the Tasks of Technology portion of this database.

4.2.2B Qualitative assessment of tradeoff between costs and data quality. Documenting the preponderance, duration, and value of all tasks performed in the economy would be a costly and complex endeavor for the BLS and the Census Bureau, if they were to cooperate as we recommend. To minimize these costs and the complexities involved, this proposal suggests a strategy that leverages two existing survey collection platforms to accomplish this (the ORS and ACES) in a relatively efficient manner that would limit methodological adjustments and minimize respondent burden, while insuring high-quality, nationally representative data.

The benefits of creating a Tasks Inventory would be enormous. It would go a long way toward providing the empirical data needed to understand the economy in a way that is consistent with leading theoretical models on the effects of automation and technology on the labor market. Beyond the effects of technology, it would also help clarify other important debates in the labor economics literature on the role of skills, training, education, global markets, and regulations in shaping the demand for tasks and how they are compensated. Without these data, it is hard to imagine how economists could confidently analyze how technology is affecting the labor market, or forecast the demand for tasks, or understand how the share of national income that is likely to go to labor would change under different automation scenarios.

One of the many important uses of these data will be to observe what percentage of an occupation's tasks are already performed by machines and how that percentage has changed in recent years. This will provide a more precise and objective measure of potential displacement than measures based on expert surveys (as in Frey and Osborne 2017). As mentioned, the combination of OES data and the task database (or other ways of integrating the tasks of machines data to firm-level data) could potentially also be used to estimate the effect of technological investments on occupational layoffs, hiring, or shifts in the types of tasks performed. These uses would produce valuable insights into the effects of technology on work.

For these reasons, we believe the benefits heavily outweigh the costs to the public and BLS' direct and indirect data users.

4.3 Conclusion

The current data available to BLS through its own survey and data collection methods, as well as those available from Census, the Bureau of Economic Analysis, and other sources is not adequate to fulfill its important mission to clarify how technology is affecting the labor market and guide public policy.

With no additional data collection, greater inter-agency cooperation would be useful in advancing the understanding of how technology affects the labor market. For example, if data from the ACES were linked to data from OES or ORS, economists could observe how investments in robots or equipment at the firm level predict changes in occupational compensation and distribution within those establishments. If OES occupational data were presented longitudinally (and supplemented with other sources to capture farmers, sole proprietors, domestic workers, and owners of unincorporated firms), analysts could observe a clean time series of occupational employment levels, growth rates, and compensation changes.

Assessing the Impact of New Technologies on the Labor Market

However, without additional data collection, economists and the public would still have many unanswered and likely unanswerable questions. Presently, there are no compelling estimates for demand flows by occupation (including vacancies, new hires, and the relationship between vacancies and hiring, or new supply flows and demand), only limited data on the skills required to enter and succeed in different occupations, and no data on how the tasks performed by workers change over time or in relation to enduring worker personality or skill characteristics. Moreover, despite the vital importance of understanding the economy through the lens of tasks, there is no task classification system and no data on the preponderance, duration, or value of tasks performed by workers or machines, when both are needed to understand the relationship between technology and labor.

Finally, while the economics literature estimates that some occupations are more at risk to automation than others, the evidentiary basis for this conclusion is thin. The most widely cited evidence relies on a non-representative and unreplicated survey of technology experts, with no consideration of the actual costs or value of automating the tasks performed by workers—nor any consideration that the tasks of occupations may evolve in response to technology or that the partial automation of an occupation's tasks may be complementary to labor demand, at least up to some threshold. New research uses patents to measure how the tasks performed by new technologies overlap with the tasks performed by workers in different occupations, but these results do not indicate whether the overlap between the tasks of occupations and technologies will be complementary or not. Patent data also lack information on the costs of new technologies and their adoption in the market.

As described above, we believe the BLS could fill these data gaps and facilitate the thorough investigation of these issues by adopting the strategy outlined above. These efforts would incur significant costs, but by leveraging existing data products and methods and partnering with the Census Bureau, the costs would be mitigated to some degree. Meanwhile, the quality and importance of the data would be very high, and the benefits to social science and public policy would be potentially very large. The effects of technology on the labor market is a much-debated topic politically and socially. With these new data products, citizens and organizations could act with a much richer grounding in factual information.

5. References

- Abraham, Katharine G. 1987. "Help-Wanted Advertising, Job Vacancies, and Unemployment." *Brookings Papers on Economic Activity* 1987 (1): 207-248.
- Acemoglu, Daron, and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics: Volume 4B*, edited by David Card and Orley Ashenfelter, 1043-171. Elsevier, North-Holland.
- Acemoglu, Daron, and Pascual Restrepo. 2017. "Robots and Jobs: Evidence from U.S. Labor Markets." National Bureau of Economic Research Working Paper 23285.
- Acemoglu, Daron, and Pascual Restrepo. 2019. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives* 33 (2): 3-30.
- Agrawal, Ajay, Joshua S. Gans, and Avi Goldfarb. 2019. "Artificial Intelligence: The Ambiguous Labor Market Impact of Automating Prediction." *The Journal of Economic Perspectives* 33 (2): 31-50.
- Arntz, Melanie, Terry Gregory, and Ulrich Zierahn. 2016. "The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis." OECD Social, Employment, and Migration Working Paper 189.
- Atalay, Engin, Phai Phongthientham, Sebastian Sotelo, and Daniel Tannenbaum. 2017. "The Evolving U.S. Occupational Structure." University of Wisconsin Madison Working Paper.
- Atkinson, Robert D., and J. John Wu. 2017. "False alarmism: Technological Disruption and the U.S. Labor Market, 1850–2015." In *ITIF@Work*, 1-28. Information Technology & Innovation Foundation, Washington, D.C.
- Attewell, Paul. 1992. "Skill and Occupational Changes in U.S. Manufacturing." In *Technology and the Future of Work*, edited by Paul S. Adler, 46-88. New York: Oxford University Press.
- Autor, David H., and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the U.S. Labor Market." *American Economic Review* 103 (5): 1553-97.
- Autor, David H., and Michael J. Handel. 2013. "Putting Tasks to the Test: Human Capital, Job Tasks, and Wages." *Journal of Labor Economics*, 31 (S1): S59-S96.
- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2006. "The Polarization of the U.S. Labor Market." *American Economic Review* 96 (2): 189-94.
- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2008. "Trends in U.S. Wage Inequality: Revising the Revisionists." *Review of Economics and Statistics* 90 (2): 300-23.
- Autor, David H., Lawrence F. Katz, and Alan B. Krueger. 1998. "Computing Inequality: Have Computers Changed the Labor Market?" *The Quarterly Journal of Economics* 113 (4): 1169-1213.
- Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics* 116 (4): 1279-1333.
- Autor, David H., and Brenden Price. 2013. "The Changing Task Composition of the U.S. Labor Market: An Update of Autor, Levy, and Murnane (2003)." MIT Working Paper.
- Autor, David H., and Anna Salomons. 2018. "Is Automation Labor-Displacing? Productivity Growth, Employment, and the Labor Share." National Bureau of Economics Research Working Paper 24871.
- Berman, Eli, Rohini Somanathan, and Hong W. Tan. 2003. "Is Skill-Biased Technological Change Here Yet? Evidence from Indian Manufacturing in the 1990s." World Bank Policy Research Working Paper 3761.
- Bessen, James. 2015. *Learning by Doing: The Real Connection Between Innovation, Wages, and Wealth*. New Haven and London: Yale University Press.
- Bessen, James E. 2011. "Was Mechanization De-skilling? The Origins of Task-Biased Technical Change." Boston University School of Law Working Paper Number 11-13.
- Blanchard, Olivier John, Peter Diamond, Robert E. Hall, and Janet Yellen. 1989. "The Beveridge Curve." *Brookings Papers on Economic Activity* 1989 (1): 1-76.
- Blank, David M., and George J. Stigler. 1957. *The Demand and Supply of Scientific Personnel*. New York: National Bureau of Economic Research.
- Borghans, Lex, Angela Lee Duckworth, James J. Heckman, and Bas ter Weel. 2008. "The Economics and Psychology of Personality Traits." *The Journal of Human Resources* 43 (4): 972-1059.
- Borjas, George J., and Richard B. Freeman. 2019. "From Immigrants to Robots: The Changing Locus of Substitutes for Workers." National Bureau of Economic Research Working Paper 25438.

- Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence." *The Quarterly Journal of Economics* 117 (1): 339-76.
- Brynjolfsson, Erik, and Andrew McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York and London: WW Norton & Company.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2018. "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies." National Bureau of Economic Research Working Paper 25148.
- Bughin, Jacques, Eric Hazan, Susan Lund, Peter Dahlstrom, Anna Wiesinger, and Amresh Subramanian. 2018. "Skill Shift: Automation and the Future of the Workforce." McKinsey Global Institute Discussion Paper.
- Bureau of Labor Statistics. 2007. "Technical Information About the BLS Multifactor Productivity Measures." Bureau of Labor Statistics, Multifactor Productivity, Washington, D.C.
- Bureau of Labor Statistics. 2008. "Technical Information About the BLS Major Sector Productivity and Costs Measures." Bureau of Labor Statistics, Major Sector Productivity and Costs, Washington, D.C.
- Bureau of Labor Statistics. 2019a. "Productivity Change in the Nonfarm Business Sector, 1947-2014." United States Department of Labor. <https://www.bls.gov/lpc/prodybar.htm> (accessed November 6, 2019).
- Bureau of Labor Statistics. 2019b. "Occupational Outlook Handbook: Fastest Growing Occupations." United States Department of Labor. <https://www.bls.gov/ooh/fastest-growing.html> (accessed June 12, 2019).
- Burnstein, Ariel, and Jonathan Vogel. 2010. "Globalization, Technology, and the Skill Premium: A Quantitative Analysis." National Bureau of Economic Research Working Paper 16459.
- Byrne, David M., John G. Fernald, and Marshall B. Reinsdorf. 2016. "Does the United States Have a Productivity Slowdown or a Measurement Problem?" *Brookings Papers on Economic Activity*: 109-82.
- Calvino, Flavio, Chiara Criscuolo, Luca Marcolin, and Mariagrazia Squicciarini. 2018. "A Taxonomy of Digital Intensive Sectors." OECD Science, Technology and Industry Working Papers, No. 2018/14.
- Card, David, and John E. DiNardo. 2002. "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles." *Journal of Labor Economics* 20 (4): 733-83.
- Caselli, Francesco, and Alan Manning. 2019. "Robot Arithmetic: New Technology and Wages." *American Economic Review: Insights* 1 (1): 1-12.
- Charbonneau, Karyne, Alexa Evans, Subrata Sarker, and Lena Suchanek. 2017. "Digitalization and Inflation: A Review of the Literature." Bank of Canada Staff Analytical Note.
- Chavrid, Vladimir, and Harold Kuptzin. 1966. "Employment Service Operating Data as a Measure of Job Vacancies." In *The Measurement and Interpretation of Job Vacancies*, 373-403. National Bureau of Economic Research.
- Clark, Gregory. 2005. "The Condition of the Working Class in England, 1209-2004." *Journal of Political Economy* 113 (6): 1307-40.
- Clark, Gregory. 2008. *A Farewell to Alms: A Brief Economic History of the World*. Princeton and Oxford, Princeton University Press.
- Conte, Andrea, and Marco Vivarelli. 2011. "Imported Skill-Biased Technological Change in Developing Countries." *The Developing Economies* 49 (1): 36-65.
- Cortes, Guido Matias, and Andrea Salvatori. 2015. "Task Specialization within Establishments and the Decline of Routine Employment." University of Manchester Working Paper.
- Couch, Kenneth A., and Dana W. Placzek. 2010. "Earnings Losses of Displaced Workers Revisited." *American Economic Review* 100 (1): 572-89.
- Cowen, Tyler. 2011. *The Great Stagnation: How America Ate all the Low-Hanging Fruit of Modern History, Got Sick, and Will (Eventually) Feel Better*. New York: Dutton.
- Cramer, Judd, and Alan B. Krueger. 2016. "Disruptive Change in the Taxi Business: The Case of Uber." *American Economic Review* 106 (5): 177-82.
- Dalton, Michael, Lisa Kahn, and Andreas Mueller. 2019. "The Role of Recruiting Intensity on Vacancy Yields: Evidence from a Large-Scale Merge of Job Postings and Survey Data." Working Paper Shared with Gallup, Inc. by Bureau of Labor Statistics.
- David, Paul. 1990. "The Dynamo and The Computer: An Historical Perspective on the Modern Productivity Paradox." *American Economic Review* 80 (2): 355-61.

- Deaton, Angus. 2016. *The Great Escape: Health, Wealth, and the Origins of Inequality*. Princeton, New Jersey: Princeton University Press.
- De Pleijt, Alexandra M., and Jacob L. Weisdorf. 2017. "Human Capital Formation from Occupations: The 'Deskilling Hypothesis' Revisited." *Cliometrica* 11 (1): 1-30.
- De Long, J. Bradford, and Lawrence H. Summers. 1991. "Equipment Investment and Economic Growth." *The Quarterly Journal of Economics* 106 (2): 445-502.
- Deming, David J. 2017. "The Growing Importance of Social Skills in the Labor Market." *The Quarterly Journal of Economics* 132 (4): 1593-1640.
- Dey, Matthew and Elizabeth Weber Handwerker. 2016. "Longitudinal Data from the Occupational Employment Statistics Survey." *Monthly Labor Review*: 370-73.
- Dey, Matthew, Steve Miller, and Dave Piccone. 2014. "OES Time Series Alternative Designs and Estimation Methods." Bureau of Labor Statistics, Washington, D.C.
https://www.bls.gov/advisory/tac/dey_oes_redesign.pdf (accessed November 8, 2019)
- Dey, Matthew, David S. Piccone Jr., & Stephen M. Miller. 2019. Model-based estimates for the Occupational Employment Statistics program. *Monthly Labor Review*.
- Dierdorff, Erich C., Donald W. Drewes, and Jennifer J. Norton. 2006. "O*NET® Tools and Technology: A Synopsis of Data Development Procedures."
<https://www.onetcenter.org/reports/T2Development.html> (accessed November 8, 2019).
- Dierdorff, Erich C., and Jennifer J. Norton. 2011. "Summary of Procedures for O*NET Task Updating and New Task Generation." National Center for O*NET Development
<https://www.onetcenter.org/reports/TaskUpdating.html> (accessed August 1, 2019).
- DiNardo, Josh E., and Jorn-Steffen Pischke. 1997. "The Returns to Computer Use Revisited: Have Pencils Changed the Wage Structure Too?" *The Quarterly Journal of Economics* 112 (1): 291-303.
- Eden, Maya, and Paul Gaggl. 2018. "On the Welfare Implications of Automation." *Review of Economic Dynamics* 29: 15-43.
- Elsby, Michael W. L., Bart Hobijn, and Ayşegül Şahin. 2014. "The Decline of the U.S. Labor Share." *Brookings Papers on Economic Activity*: 1-52.
- Fayer, Stella, Alan Lacey, and Audrey Watson. 2017. "STEM Occupations: Past, Present, and Future." Bureau of Labor Statistics, Spotlight on Statistics, Washington, D.C.
- Frey, Carl Benedikt, and Michael A. Osborne. 2017. "The Future of Employment: How Susceptible Are Jobs to Computerisation?" *Technological Forecasting and Social Change* 114: 254-80.
- Ford, Martin. 2015. *Rise of the Robots: Technology and the Threat of a Jobless Future*. New York: Basic Books.
- Fuchs, Victor R. 1980. "Economic Growth and the Rise of Service Employment." National Bureau of Economic Research Working Paper 486.
- Gaggl, Paul, and Greg C. Wright. 2017. "A Short-Run View of What Computers Do: Evidence From a U.K. Tax Incentive." *American Economic Journal: Applied Economics* 9 (3): 262-94.
- Goldfarb, Avi, and Catherine Tucker. 2019. "Digital Economics." *Journal of Economic Literature* 57 (1): 3-43.
- Goldin, Claudia, and Lawrence F. Katz. 1998. "The Origins of Technology-Skill Complementarity." *The Quarterly Journal of Economics* 113 (3): 693-732.
- Goldin, Claudia, and Lawrence F. Katz. 2010. *The Race Between Education and Technology*. Cambridge and London: The Belknap Press of Harvard University Press.
- Goldin, Claudia, and Kenneth Sokoloff. 1982. "Women, Children, and Industrialization in the Early Republic: Evidence from the Manufacturing Censuses." *The Journal of Economic History* 42 (4): 741-774.
- Goos, Maarten, Jozef Konings, and Emilie Rademakers. 2016. "Future of work in the digital age: Evidence from OECD countries." Utrecht University Working Paper.
- 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. 2009. "Job Polarization in Europe." *American Economic Review* 99 (2): 58-63.
- Goos, Maarten, Alan Manning, and Anna Salomons. 2014. "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring." *American Economic Review* 104 (8): 2509-26.
- Gordon, Robert J. 2017. *The Rise and Fall of American Growth: The U.S. Standard of Living Since the Civil War*. Princeton: Princeton University Press.

- Gorg, Holger, and Eric Strobl. 2002. "Relative Wages, Openness and Skill-Biased Technological Change." Institute for the Study of Labor IZA Discussion Paper 596.
- Graetz, Georg, and Guy Michaels. 2018. "Robots at Work." *Review of Economics and Statistics* 100 (5): 753-768.
- Harrigan, James, Ariell Reshef, and Farid Toubal. 2016. "The March of the Techies: Technology, Trade, and Job Polarization in France, 1994-2007." National Bureau of Economic Research Working Paper 22110.
- Heckman, James J., and Tim Kautz. 2012. "Hard Evidence on Soft Skills." *Labour Economics* 19(4): 451-464.
- Hershbein, Brad, and Lisa B. Kahn. 2018. "Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings." *American Economic Review* 108 (7): 1737-72.
- Holzer, Harry. 2015. *Job Market Polarization and U.S. Worker Skills: A Tale of Two Middles*. Washington, D.C.: The Brookings Institution.
- Humphries, Jane. 2013. "Childhood and Child Labour in the British Industrial Revolution." *The Economic History Review* 66 (2): 395-418.
- Jones, Charles I. 2016. "The Facts of Economic Growth." National Bureau of Economics Working Paper 21142.
- Katz, Lawrence F., and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963-1987: Supply and Demand Factors." *The Quarterly Journal of Economics* 107 (1): 35-78.
- Katz, Lawrence F., and Robert A. Margo. 2014. "Technical Change and the Relative Demand for Skilled Labor: The United States in Historical Perspective." In *Human Capital in History: The American Record*, edited by Leah Platt Boustan, Carola Frydman, and Robert A. Margo, 15-57. University of Chicago Press.
- Keynes, John Maynard. 1978. "Economic Possibilities for our Grandchildren (1930)." In *The Collected Writings of John Maynard Keynes: Essays in Persuasion, Volume 9*, edited by Elizabeth Johnson and Donald Moggridge, 321-32. London: Macmillan Press.
- Kezdi, Gabor. 2002. "Two Phases of Labor Market Transition in Hungary: Inter-Sectoral Reallocation and Skill-Biased Technological Change." Budapest Working Papers on the Labour Market BWP.
- Kianian, Babak, Sam Tavassoli, and Tobias C. Larsson. 2015. "The Role of Additive Manufacturing Technology in Job Creation: An Exploratory Case Study of Suppliers of Additive Manufacturing in Sweden." *Procedia CIRP* 26: 93-98.
- Koch, Michael, Ilya Manuylov, and Marcel Smolka. 2019. "Robots and firms." CESifo Working Paper, No. 7608, Center for Economic Studies and Ifo Institute (CESifo), Munich.
- Krigman, Eliza. 2014. "Two-for-One Deal: Earning College Credit for STEM in High School." *US News & World Report*. <https://www.usnews.com/news/stem-solutions/articles/2014/05/06/two-for-one-deal-earning-college-credit-for-stem-in-high-school> (accessed November 8, 2019).
- Krueger, Alan B. 1993. "How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984-1989." *The Quarterly Journal of Economics* 108 (1): 33-60.
- Law, Kenneth S., Chi-Sum Wong, and William H. Mobley. 1998. "Toward a Taxonomy of Multidimensional Constructs." *The Academy of Management Review* 23 (4): 741-755.
- Lemieux, Thomas. 2006. "Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?" *American Economic Review* 96 (3): 461-98.
- Lindert, Peter H. 1986. "Unequal English Wealth Since 1670." *Journal of Political Economy* 94 (6): 1127-62.
- Lindert, Peter H., and Jeffrey G Williamson. 2016. *Unequal Gains: American Growth and Inequality since 1700*. Princeton: Princeton University Press.
- Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. 2017. "Distributional National Accounts: Methods and Estimates for the United States." *The Quarterly Journal of Economics* 133 (2): 553-609.
- Maese, Ellyn. 2019. "BLS Top 20 Growing Jobs: O*Net Task Analysis." Gallup, Inc., unpublished data set (accessed April 12, 2019).
- MacCrory, Frank, George Westerman, Yousef Alhammadi, and Erik Brynjolfsson. 2014. "Racing With and Against the Machine: Changes in Occupational Skill Composition in an Era of Rapid Technological Advance." Proceedings of the Thirty-Fifth International Conference on Information Systems. Association for Information Systems.
- Manning, Alan. 2004. "We Can Work It Out: The Impact of Technological Change on the Demand for Low-Skill Workers." *Scottish Journal of Political Economy* 51 (5): 581-608.

- Marx, Karl. 1867. "Chapter 25: The General Law of Capitalist Accumulation." In *Capital: A Critique of Political Economy*, 762-872. London: Penguin Books.
- McCloskey, Deirdre N. 2016. *Bourgeois Equality: How Ideas, Not Capital or Institutions, Enriched the World*. Chicago and London: University of Chicago Press.
- Michaels, Guy, Ashwini Natraj, and John Van Reenen. 2014. "Has ICT Polarized Skill Demand? Evidence from Eleven Countries Over Twenty-Five Years." *Review of Economics and Statistics* 96 (1): 60-77.
- Millington, Kerry A. 2017. *How Changes in Technology and Automation Will Affect the Labour Market in Africa*. Brighton, U.K.: Institute of Development Studies.
- Mueller, Charlotte, and John Wohlford. 2008. "Developing a New Business Survey: Job Openings and Labor Turnover Survey at the Bureau of Labor Statistics." Bureau of Labor Statistics, Washington, D.C. <https://www.bls.gov/osmr/research-papers/2000/pdf/st000160.pdf> (accessed November 8, 2019).
- Muro, Mark, Jonathan Rothwell, Scott Andes, Kenan Fikri, and Siddharth Kulkarni. 2015. *America's Advanced Industries: What They Are, Where They Are, and Why They Matter*. Washington, D.C.: The Brookings Institution.
- Muro, Mark, Sifan Liu, Jacob Whiton, and Siddharth Kulkarni. 2017. "Digitalization and the American Workforce." Brookings Papers on Economic Activity Working Paper.
- Nedelkoska, Ljubica, and Glenda Quintini. 2018. "Automation, Skills Use and Training." OECD Social, Employment, and Migration Working Papers 202.
- Nuvolari, Alessandro. 2002. "The 'Machine Breakers' and the Industrial Revolution." *Journal of European Economic History* 31 (2): 393-426.
- O*NET Resource Center. 2018. "O*NET® Data Collection Program Office of Management and Budget Clearance Package Supporting Statement Part A: Justification." <https://www.onetcenter.org/reports/omb2018.html> (accessed September 1, 2018)
- OECD/Eurostat. 2018. *Oslo Manual 2018: Guidelines for Collecting, Reporting and Using Data on Innovation, 4th Edition*. The Measurement of Scientific, Technological and Innovation Activities, Paris/Luxembourg: OECD Publishing/Eurostat.
- OECD. 2015. *Frascati Manual 2015: Guidelines for Collecting and Reporting Data on Research and Experimental Development*. The Measurement of Scientific, Technological and Innovation Activities. Paris: OECD Publishing.
- OECD. 2018. "Private Equity Investment in Artificial Intelligence." OECD Going Digital Policy Note.
- OECD. 2019. *Compendium of Productivity Indicators 2019*. Paris: OECD Publishing.
- Pabilonia, Sabrina Wulff, Michael W. Jadoo, Bhavani Khandrika, Jennifer Price, and James D. Mildenberger. 2019. "BLS Publishes Experimental State-Level Labor Productivity Measures." *Monthly Labor Review*, Bureau of Labor Statistics. <https://doi.org/10.21916/mlr.2019.12> (accessed August 1, 2019).
- Park, Jaehyuk, Ian B. Wood, Elise Jing, Azadeh Nematzadeh, Souvik Ghosh, Michael D. Conover, and Yong-Yeol Ahn. 2019. "Global Labor Flow Network Reveals the Hierarchical Organization and Dynamics of Geo-Industrial Clusters." *Nature Communications* 10 (3449).
- Piketty, Thomas. 2012. "About Capital in the Twenty-First Century." *American Economic Review* 105 (5): 48-53.
- Pilot, Michael J. 1999. "Occupational Outlook Handbook: A Review of 50 Years of Change." *Monthly Labor Review* 129: 8-26.
- Pratt, Gill A. 2015. "Is a Cambrian Explosion Coming for Robotics?" *Journal of Economic Perspectives* 29 (3): 51-60.
- Ritchie, Stuart. J., and Elliot M. Tucker-Drob. 2018. "How Much Does Education Improve Intelligence? A Meta-Analysis." *Psychological Science* 29 (8): 1358-1369.
- Rothwell, Jonathan. 2013. *The Hidden STEM Economy*. Washington, D.C.: The Brookings Institution.
- Rothwell, Jonathan. 2014a. *Still Searching: Job Vacancies and STEM Skills*. Washington, D.C.: Metropolitan Policy Program at Brookings.
- Rothwell, Jonathan. 2014b. "The Silicon Valley Wage Premium." The Brookings Institution, Washington, D.C. <https://www.brookings.edu/blog/the-avenue/2014/08/06/the-silicon-valley-wage-premium/> (accessed November 9, 2019).
- Rothwell, Jonathan, and Siddharth Kulkarni. 2015. "Data and Methods Appendix for America's Advanced Industries: What They Are, Where They Are, Why They Matter." Metropolitan Policy Program at Brookings, Washington, D.C. <https://www.brookings.edu/wp-content/uploads/2015/02/Advanced-Industries-Data-and-Methods-Appendix.pdf> (accessed August 1, 2019).

- Schanzenbach, Diane Whitmore, Ryan Nunn, Lauren Bauer, Megan Mumford, and Audrey Breitwieser. 2016. *Seven Facts on Noncognitive Skills From Education to the Labor Market*. Washington, D.C.: The Hamilton Project, Brookings Institution.
- Schmidt, Frank L., In-Sue Oh, and Jonathan A. Shaffer. 2016. "The Validity and Utility of Selection Methods in Personnel Psychology: Practical and Theoretical Implications of 100 Years of Research Findings" Fox School of Business Research Paper.
- Schwab, Klaus. 2016. *The Fourth Industrial Revolution*. World Economic Forum: Geneva, Switzerland.
- Smith, Aaron, and Monica Anderson. 2017. "Automation in Everyday Life." Pew Research Center: Internet and Technology.
- Spitz-Oener, Alexandra. 2006. "Technical Change, Job Tasks, and Rising Educational Demands: Looking Outside the Wage Structure." *Journal of Labor Economics* 24 (2): 235-70.
- Stone, Peter, Rodney Brooks, Erik Brynjolfsson, Ryan Calo, Oren Etzioni, Greg Hager, Julia Hirschberg, Shivaram Kalyanakrishnan, Ece Kamar, Sarit Kraus, Kevin Leyton-Brown, David Parkes, William Press, AnnaLee Saxenian, Julie Shah, Milind Tambe, and Astro Teller. 2016. "Artificial Intelligence and Life in 2030." One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel. Stanford, CA Stanford University.
- UNESCO. 2017. "Summary Report of the 2015 UIS Innovation Data Collection." UNESCO Institute for Statistics Information Paper, No. 37, United Nations Educational, Scientific and Cultural Organization (UNESCO), Montreal, Canada.
- United States Department of Education. 2016. "Science, Technology, Engineering and Math: Education for Global Leadership." United States Department of Education. <https://www.ed.gov/Stem> (accessed November 8, 2019).
- Webb, Michael. 2019. "The Impact of Artificial Intelligence on the Labor Market." Paper presented at the SI 2019 Conference on Research in Income and Wealth, July 15, 2019 in Cambridge, MA. <https://conference.nber.org/sched/SI19PRCR> (accessed November 8, 2019).
- West, Martin R. 2016. "Should Non-Cognitive Skills Be Included in School Accountability Systems? Preliminary Evidence from California's CORE Districts." *Evidence Speaks Reports* 1 (13): 1-7.
- Wilson, H. James, Paul R. Daugherty, and Nicola Morini-Bianzino. 2017. "The Jobs That Artificial Intelligence Will Create." *MIT Sloan Management Review* 58: 13-16.
- Wisskirchen, Gerlind, Blandine Thibault Biacabe, Ulrich Bormann, Annemarie Muntz, Gunda Niehaus, Guillermo Jimenez Soler, Beatrice von Brauchitsch. 2017. "Artificial Intelligence and Robotics and Their Impact on the Workplace." IBA Global Employment Institute.
- World Bank Group. 2016. *World Development Report 2016: Digital Dividends*. Washington, D.C.: World Bank.
- World Inequality Database. 1980-2019. "World Inequality: Dataset." <https://wid.world/> (accessed November 8, 2019).
- Wyatt, Ian D. and Daniel E. Hecker. 2006. "Occupational changes during the 20th century." *Monthly Labor Review* 129 35-57.
- Xu, Min, Jeanne M. David, and Suk Hi Kim. 2018. "The Fourth Industrial Revolution: Opportunities and Challenges." *International Journal of Financial Research* 9 (2): 90-5.

Appendix

Appendix A: Acronyms and Abbreviations

ABS	Annual Business Survey
ACES	Annual Capital Expenditure Survey
AI	Artificial Intelligence
ASEC	Annual Social and Economic Supplement
BEA	Bureau of Economic Analysis
BLS	Bureau of Labor Statistics
BRDIS	Business R&D and Innovation Survey
CES	Current Employment Statistics
CIS	Community Innovation Survey (EU)
CPS	Current Population Survey
GDP	Gross Domestic Product
EU	European Union
ICT	Information and Communication Technology
IFR	International Federation of Robotics
ILO	International Labour Organization
JOLTS	Job Openings and Labor Turnover Survey
MFP	Multifactor Productivity
NAICS	North American Industry Classification System
NCS	National Compensation Survey
NLP	Natural Language Processing
NLS	National Longitudinal Surveys
NSF	National Science Foundation
OECD	Organisation for Economic Co-operation and Development
OES	Occupational Employment Statistics
ONS	Office for National Statistics (U.K.)
ORS	Occupational Requirements Survey
PIAAC	Program for the International Assessment of Adult Competencies

Assessing the Impact of New Technologies on the Labor Market

PPI	Producer Price Index
R&D	Research and development
SOC	Standard Occupational Classification
STEM	Science, technology, engineering, or mathematical
UN	United Nations

Appendix B: BLS Data Products

- Consumer Price Index (CPI)
- Producer Price Index (PPI)
- International Price Program (IPP)
- Current Population Survey (CPS)
- Current Employment Statistics (CES)
- Job Openings and Labor Turnover Survey (JOLTS)
- Quarterly Census of Employment & Wages (QCEW)
- Business Employment Dynamics (BDM)
- American Time Use Surveys
- Local Area Unemployment Statistics (LAUS)
- Occupational Employment Statistics
- Consumer Expenditure Survey (CE)
- National Compensation Survey (NCS)
- Occupational Employment Statistics (OES) Survey
- Labor Productivity and Costs (LPC)
- Major Sector Multifactor Productivity
- Industry Productivity
- Occupational Requirements Survey (ORS)

Appendix C: Datasets From U.S. Principal Statistical Agencies

Principal Statistical Agencies

Census Bureau

- 2010 Census
- American Community Survey
- Annual Survey of Manufacturers
- Economic Census
- Survey of Income and Program Participation

Bureau of Economic Analysis

- Gross Domestic Product
- GDP by Industry
- Corporate Profits
- Gross Domestic Income
- Gross Output by Industry
- Input-Output Accounts
- Employment by Industry
- Integrated Industry-Level Production Account (KLEMS)
- International Transactions
- International Trade in Goods and Services
- New Foreign Direct Investment in the United States
- Fixed Assets by Type
- Industry Fixed Assets
- Employment by State
- Employment by County, Metro, and Other Areas
- Digital Economy
- Integrated Macroeconomic Accounts for the United States
- Integrated BEA GDP-BLS Productivity Account
- Personal Income and Outlays
- International Services
- Activities of US Multinational Enterprises (MNEs)
- Activities of US Affiliates of Multinational Enterprises (MNEs)
- Gross Domestic Purchases Price Index

Bureau of Transportation Statistics

- Commodity Flow Survey Data
- Transportation Economic Trends

Economic Research Service

- Agricultural Productivity in the US
- Creative Class Employment
- Farm Income and Wealth Statistics
- International Agriculture Productivity
- International Macroeconomic Data Set
- Farm Labor
- Farm Production Expenditures Annual Summary

Assessing the Impact of New Technologies on the Labor Market

National Center for Education Statistics

- Educational Longitudinal Study (ELS)
- High School Longitudinal Study (HSLS)
- Baccalaureate and Beyond (B&B)
- Baccalaureate and Beyond, Graduate Studies (B&B:GR)
- Beginning Postsecondary Students (BPS)
- National Postsecondary Student Aid Study, Undergraduate (NPSAS:UG)
- National Postsecondary Student Aid Study, Graduate (NPSAS:GR)
- Adult Training and Education Survey (ATES)

National Center for Science and Engineering Statistics

- Business Research and Development and Innovation Survey (BRDIS)
- FFRDC Research and Development Survey
- Survey of Federal Funds for Research and Development
- Survey of State Government Research and Development
- Science and Engineering State Profiles
- Higher Education Research and Development Survey (HERD)
- Integrated Postsecondary Education Data System Completions Survey
- National Survey of College Graduates
- National Survey of Recent College Graduates
- Survey of Doctorate Recipients
- Survey of Earned Doctorates
- Survey of Federal Science and Engineering Support to Universities, Colleges, and Nonprofit Institutions
- Survey of Graduate Students and Postdoctorates in Science and Engineering
- Survey of Postdocs at Federally Funded Research and Development Centers
- Survey of Science and Engineering Research Facilities

Statistics of Income

- Integrated Business Data (IBD): Selected Financial Data on Business

Other Statistical Programs of Federal Agencies

Department of Labor/North Carolina Department of Commerce

- O*NET

Federal Reserve Board

- Survey of Consumer Finances
- Survey of Household Economics and Decisionmaking

Foreign Agricultural Service

- Export Sales Reporting (ESR)
- Global Agricultural Trade System (GATS)
- Production, Supply, and Distribution (PS&D)

International Trade Administration

- Global Patterns of US Merchandise Trade
- US Exporting Companies
- Employment and Trade: Jobs Supported by State Exports

Appendix D: International Statistical Sources

European Central Bank

International Labour Organization (ILO)

International Monetary Fund (IMF)

International Telecommunication Union (ITU)

Organization for Economic Cooperation and Development (OECD)⁶⁰

- Includes datasets from the following countries:
 - Australia Bureau of Statistics
 - Statistics Austria
 - Statistics Belgium
 - Statistics Canada
 - Instituto Nacional de Estadísticas (Chile)
 - Czech Statistical Office
 - Statistics Denmark
 - Statistics Estonia
 - Institut National de la Statistique et des Études Économiques (France)
 - Le Portail de la statistique publique française (France)
 - Federal Statistics Office (Germany)
 - National Statistical Service of Greece
 - Hungarian Central Statistical Office
 - Statistics Iceland Central Statistics Office Ireland
 - Istituto Nazionale di Statistica (Italy)
 - Statistics Bureau of Japan
 - Statistics Korea
 - Central Statistical Bureau of Latvia
 - Statistics Lithuania
 - STATEC (Luxembourg)
 - INEGI (Mexico)
 - Statistics Netherlands
 - Statistics New Zealand
 - Statistics Norway
 - Statistics Poland
 - Statistics Portugal
 - Statistical Office of the Slovak Republic
 - Statistical Office of the Republic of Slovenia
 - Instituto Nacional de Estadística (Spain)
 - Statistics Sweden
 - Swiss Federal Statistical Office
 - Turkish Statistical Institute
 - UK Statistics Authority
 - Office for National Statistics (UK)

Paris 21

Statistical Office of the European Union (EUROSTAT)

United Nations Conference on Trade and Development (UNCTAD)

United Nations Economic and Social Commission for Asia and the Pacific (UNESCAP)

United Nations Economic Commission for Africa (UNECA)

United Nations Economic Commission for Europe (UNECE)

United Nations Economic Commission for Latin America and the Caribbean (UNECLAC)

United Nations Economic and Social Commission for Western Asia (UNESCWA)

⁶⁰ Some OECD data products include data from non-OECD countries (i.e., accession countries, key partners, and some G20 countries). In addition to OECD countries, statistical agencies from non-OECD countries were reviewed if categorized as having full statistical capacity by the World Bank. These agencies include: Statistical Committee of the Republic of Armenia, National Bureau of Statistics of China, Ministry of Statistics and Programme Implementation (India), National Bureau of Statistics (Moldova), and The State Committee of Statistics of Ukraine.

Assessing the Impact of New Technologies on the Labor Market

United Nations Educational, Scientific and Cultural Organization (UNESCO)

United Nations Industrial Development Organization (UNIDO)

United Nations Statistics Division (UNSD)

World Bank (WB)

World Trade Organization (WTO)