Volume Title: The Economics of Artificial Intelligence: An Agenda

Volume Authors/Editors: Ajay Agrawal, Joshua Gans, and Avi Goldfarb, editors

Volume Publisher: University of Chicago Press

Volume ISBNs: 978-0-226-61333-8 (cloth); 978-0-226-61347-5 (electronic)

Volume URL: http://www.nber.org/books/agra-1

Conference Date: September 13–14, 2017

Publication Date: May 2019

Chapter Title: Artificial Intelligence, Labor, Productivity, and the Need for Firm-Level Data

Chapter Author(s): Manav Raj, Robert Seamans

Chapter URL: http://www.nber.org/chapters/c14037

Chapter pages in book: (p. 553 – 565)

22

Artificial Intelligence, Labor, Productivity, and the Need for Firm-Level Data

Manav Raj and Robert Seamans

22.1 Introduction

There have recently been dramatic increases in the technical capabilities of artificial intelligence (AI).[1] For example, in February 2016, Google's DeepMind used its AI to beat Korean Go master Lee Se-dol,[2] and in January 2017, an AI system called DeepStack beat humans at the complex poker game Texas Hold 'Em.[3] The Electronic Frontier Foundation (EFF) has tracked the rapid progress of AI in performing tasks at human-like levels of capability in domains including voice recognition, translation, visual image recognition, and others.[4] These advancements have led to both excitement about the capability of new technology to boost economic growth and concern about the fate of human workers in a world in which computer algorithms can perform many of the functions that a human can (e.g., Frey and Osborne 2017; Furman 2016b).

Indicative of this excitement and interest in the area, recent academic research, using national-level data on worldwide robotics shipments, suggests that robotics may have been responsible for about one-tenth of the

Manav Raj is a PhD student in the Management and Organizations Department at the Stern School of Business, New York University. Robert Seamans is associate professor of management and organizations at the Stern School of Business, New York University.

For acknowledgments, sources of research support, and disclosure of the authors' material financial relationships, if any, please see http://www.nber.org/chapters/c14037.ack.

1. Artificial intelligence is a loose term used to describe a range of advanced technologies that exhibit human-like intelligence, including machine learning, autonomous robotics and vehicles, computer vision, language processing, virtual agents, and neural networks.

2. https://www.nytimes.com/2016/03/10/world/asia/google-alphago-lee-se-dol.html.

3. https://www.scientificamerican.com/article/time-to-fold-humans-poker-playing-ai-beats-pros-at-texas-hold-rsquo-em/.

4. https://www.eff.org/ai/metrics.

increase in the gross domestic product (GDP) between 1993 and 2007 (Graetz and Michaels 2015). Moreover, according to the 2016 *Economic Report of the President*, worldwide demand for robotics has nearly doubled between 2010 and 2014, and the number and share of robotics-oriented patents have also increased (CEA 2016). Thus, robots may now be contributing even more to GDP growth than in the past.

However, even as these technologies may be contributing to GDP growth at a national level, we lack an understanding about how and when they contribute to firm-level productivity, what conditions they complement or substitute for labor, how they affect new firm formation, and how they shape regional economies. We lack an understanding of these issues because, to date, there is a lack of firm-level data on the use of robotics and AI. Such data will be important to collect to answer these questions and to inform policymakers about the role of these new technologies in our economy and society.

This chapter describes high-level findings about the effects of robotics on the economy while highlighting the few articles addressing the impact of AI, describes shortcomings of the existing data, and argues for more systematic data collection at the firm level. We echo a recent National Academies of Science Report (NAS 2017) calling for more data collection on the effects of automation, including both artificial intelligence and robotics, on the economy. More generally, collection of and access to granular data allows for better analysis of complex questions, and provides a "scientific safeguard" via replication work done by multiple sets of researchers (Lane 2003).

22.2 Existing Empirical Work

While there is little empirical work on the effects of either AI or robots, there are comparably more studies on robots, likely owing to their physical nature, which makes them easier to track over time and location. Initial studies of the effect of robots on productivity and labor provide a mixed view. Using robot shipment data at the country, industry, and year level from the International Federation of Robotics (IFR), Graetz and Michaels (2015) find large effects on productivity growth. Looking at national-level data on robot shipments across seventeen countries, Graetz and Michaels show that robots may be responsible for roughly one-tenth of the increases in the gross domestic product of these countries between 1993 and 2007 and may have increased productivity growth by more than 15 percent. This is a significant effect; according to the authors, it is comparable to the impact of the adoption of steam engines on British labor productivity in the nineteenth century. They also find evidence that, on average, wages increase with robot use, but hours worked drops for low-skilled and middle-skilled workers.

In another study using IFR data, Acemoglu and Restrepo (2017) examine the impact of the increase in industrial robot usage on regional US labor markets between 1990 and 2007. Using the distribution of robots at the industry level in other advanced countries as an instrument, the authors find that industrial robot adoption in the United States was negatively correlated with employment and wages during this time period. They estimate that each additional robot reduced employment by six workers and that one new robot per thousand workers reduced wages by 0.5 percent. The authors note that the effects are most pronounced in manufacturing, particularly in routine manual and blue-collar occupations, and for workers without a college degree. Further, they find no positive effects on employment due to the adoption of robotics in any industry.

The European Commission Report on Robotics and Employment (EC 2016) examined the use of industrial robots in Europe. The report relies on robotics data from the European Manufacturing Survey, a sample of 3,000 manufacturing firms in seven European countries, which has been periodically administered since 2001, most recently in 2012. Using this data, the authors find that the use of industrial robots is likelier in larger companies, firms utilizing batch production, and firms that are export oriented. The study finds no evidence that the use of industrial robots has any direct effect on employment, though firms utilizing robotics do have significantly higher levels of labor productivity.

More broadly, existing work on automation and employment has suggested that automation can either substitute for or complement labor. Frey and Osborne (2017) argue that almost half of the total US employment is at risk of being automated over the next two decades. Similarly, Brynjolfsson, and McAfee (2014) suggest that, due to the automation of cognitive tasks, new technologies may increasingly serve as substitutes rather than complements. On the other hand, other research has found that positive technology shocks have historically increased job opportunities and employment overall (e.g., Alexopoulos and Cohen 2016).

Regardless of the effect of automation on employment in the directly impacted industry, technology adoption may have positive upstream and downstream effects on labor. Autor and Salomons (2017) show that, while employment seems to fall within an industry as industry-specific productivity increases, positive spillovers to other sectors more than offset the negative own-industry employment effect. Further, Bessen (2017) finds that new technologies should have a positive effect on employment if they improve productivity in markets where there is a large amount of unmet demand. In the context of robotics and automation, Bessen suggests that new computer technology is associated with employment declines in manufacturing, where demand has generally been met, but is correlated with employment growth in less saturated, nonmanufacturing industries. Similarly, Mandel (2017), studying the effects of e-commerce, finds that job losses at brick-and-mortar department stores were more than made up for by new opportunities at fulfillment and call centers. Dauth et al. (2017) combines German labor market data with IFR robot shipment data and finds that, while each additional industrial robot leads to the loss of two manufacturing jobs, enough new jobs are created in the service industry to offset and in some cases overcompensate for the negative employment effect in manufacturing.

There has been less systematic work on the effect of AI on the economy. Two notable exceptions are studies by Frey and Osborne (2017) and the McKinsey Global Institute (MGI). Frey and Osborne (2017) attempt to determine what jobs may be particularly susceptible to automation and to provide an idea of how large an impact automation could have on the US labor force. The authors focus particularly on machine learning and its application to mobile robotics, and propose a model to predict the extent of computerization's impact on nonroutine tasks, noting potential engineering bottlenecks at tasks involving high levels of perception or manipulation, creative intelligence, and social intelligence. After categorizing tasks by their susceptibility to automation, Frey and Osborne map these tasks to the O*NET job survey, which provides open-ended descriptions of skills and responsibilities involved in an occupation over time. Integrating this data set with employment and wage data from the Bureau of Labor Statistics (BLS) allows the authors to propose certain subsets of the labor market that may be at high, medium, or low risk of automation. The study finds that 47 percent of US employment is at high risk of computerization. It should be noted that this study is at an aggregate level and does not examine how firms may react, any labor saving innovations that could arise, or potential productivity or economic growth.

Frey and Osborne's work has also been applied by researchers in other countries—mapping Frey and Osborne's occupation-level findings to German labor market data, Brzeski and Burk (2015) suggest that 59 percent of German jobs may be highly susceptible to automation, while conducting that same analysis in Finland, Pajarinen and Rouvinen (2014) suggest that 35.7 percent of Finnish jobs are at high risk to automation.

The Organisation for Economic Co-operation and Development (OECD) similarly set out to estimate the automatability of jobs across twenty-one OECD countries applying Frey and Osborne to a task-based approach. The OECD report argues that certain tasks will be displaced and that the extent that bundles of tasks differ within occupations and across countries may make certain occupations less prone to automation than Frey and Osborne predicted. Relying upon the task categorization done by Frey and Osborne, the authors map task susceptibility to automation to US data from the Programme for the International Assessment of Adult Competencies (PIAAC), a microlevel data source containing indicators on socioeconomic characteristics, skills, job-related information, job tasks, and competencies at the individual level. They then construct a model using the PIAAC to create a predicted susceptibility to automation based off of the observables in the PIAAC data to mirror the automatability score that Frey and Osborne created. This model is then applied at the worker level across all the PIAAC data to predict how susceptible occupations may be to automation. By conducting the analysis at the individual level, the OECD argues that it is better able to account for task variation between individuals within the same occupation. As a result, the report suggests that Frey and Osborne overestimated the extent to which occupations would be susceptible to automation. The OECD Report argues that only 9 percent of jobs in the United States and across OECD countries will be highly susceptible to automation. The report continues to discuss variations across OECD countries, suggesting that the percent can range from 6 percent (in Korea) up to 12 percent (in Austria).

Mann and Püttmann (2017) take a different approach to analyze the effects of automation on employment. In their study, the authors rely on information provided from granted patents. They apply a machine-learning algorithm to all US patents granted from 1976 to 2014 to identify patents related to automation (an automation patent is defined as a "device that operates independently from human intervention and fulfills a task with reasonable completion"). They then link the automation patents to the industries they are likely to be used in, and identify which areas in the United States that these industries are related in. By examining economic indicators in comparison to the density of automation patents used in an area, Mann and Puttman find that though automation causes manufacturing employment to fall, it increases employment in the service sector, and overall has a positive impact on employment.

In June 2017, the McKinsey Global Institute published an independent discussion paper examining trends in investment in artificial intelligence, the prevalence of AI adoption, and how AI is being deployed by companies that have started to use the technology (MGI Report 2017). For the purpose of their report, the authors adopted a fairly narrow definition of AI, focusing only on AI technology that is programmed to conduct one set task. The MGI report conducted their investigation with a multifaceted approach: it surveyed executives at over 3,000 international firms, interviewed industry experts, and analyzed investment flows using third-party venture capital, private equity, and mergers and acquisitions data. Using the data collected, the MGI report attempts to answer questions regarding adoption by sector, size, and geography; to look at performance implications of adoption; and to examine potential impacts to the labor market. Though the findings are presented at an aggregate level, much of the data, particularly the survey of executives, were collected at the firm level, allowing for further inquiry if one had access.

In addition to these published works, other researchers have begun to examine the effect of AI on occupations by looking at its impact on individual abilities and skills. Brynjolfsson, Mitchell, and Rock (forthcoming) apply a rubric from Brynjolfsson and Mitchell (2017) that evaluates the potential for applying machine learning to tasks to the set of work activities and tasks in the Bureau of Labor Statistics' O*NET occupational database. With this analysis, they create a "Suitability for Machine Learning" for labor inputs in the United States. Similar research by Felten, Raj, and Seamans (forthcoming) uses data-tracking progress in artificial intelligence aggregated by the Electronic Frontier Foundation (EFF) across a variety of different artificial intelligence metrics and the set of fifty-two abilities in the O*NET occupational database to identify the impact of artificial intelligence on each of the abilities, and create an occupation-level score measuring the potential impact of AI on the occupation. Because the data from the EFF is separated by AI metric, this work allows for the investigation and simulation of progress in different kinds of AI technology, such as image recognition, speech recognition, and ability to play abstract strategy games among others.

The current body of empirical literature surrounding robotics and AI adoption is growing, but is still thin, and despite often trying to answer similar questions, different studies have found disparate results. These discrepancies highlight the need for further inquiry, replication studies, and more complete and detailed data.

22.3 The Need for Firm-Level Data

While there is generally a paucity of data examining the adoption, use, and effects of both AI and robotics, there is currently less information available regarding AI. There are no public data sets on the utilization or adoption of AI at either the macro or micro level. The most complete source of information, the MGI study, is proprietary and inaccessible to the general public or the academic community.

The most comprehensive and widely used data set examining the diffusion of robotics is the International Federation of Robotics Robot Shipment Data. The IFR has been recording information regarding worldwide robot stock and shipment figures since 1993. The IFR collects this data from its members, who are typically large robot manufacturers such as FANUC, KUKA, and Yaskawa. The data are broken up by country, year, industry, and technological application, which allows for analysis of the industry-specific impacts of technology adoption. However, the IFR data set has shortcomings. The IFR defines an industrial robot as an "automatically controlled, reprogrammable, multipurpose manipulator, programmable in three or more axes, which can be either fixed in place or mobile for use in industrial automation applications."[5] This definition limits the set of industrial robots and ensures that the IFR does not collect any information on dedicated industrial robots that serve one purpose. Further, some of the robots are

5. https://ifr.org/standardization.

not classified by industry, detailed data is only available for industrial robots (and not robots in service, transportation, warehousing, or other sectors), and geographical information is often aggregated (e.g., data exist for North America as a category rather than the United States, or an individual state within the United States).

Another issue with the IFR data is the difficulty of integrating it with other data sources. The IFR utilizes its own industry classifications when organizing the data, rather than relying on broadly used identifiers such as the North American Industry Classification System (NAICS). Mapping IFR data to other data sets (such as BLS or census data) first requires cross-referencing IFR classifications to other identifiers. Industry-level data also cannot be used to answer micro-oriented questions about the impacts and reaction to technology adoption at the firm level.

While the IFR data are useful for some purposes, particularly examining the adoption of robotics by industry and country, its aggregated nature obscures differences occurring within industries and across regions, making it difficult to uncover when and how robots might serve as substitutes or complements to labor, and obscuring the differential effects of adoption within industries or countries. Additional data is needed to answer the issues raised above and to replicate existing studies. In particular, the National Academy of Sciences Report (NAS 2017) highlights the need for computer capital broken down at the firm and occupation level, skill changes over time by field, and data on organizational processes as they relate to technology adoption.

The European Manufacturing Survey (EMS) has been organized and executed periodically by a number of research organizations and universities across Europe since 2001, and is currently one of the only firm-level data sets examining the adoption of robotics. The overall objective of the EMS is to provide empirical evidence regarding the use and impact of technological innovation in manufacturing at the firm level. The EMS accomplishes this via a survey of a random sample of manufacturing firms with at least twenty employees across seven European countries (Austria, France, Germany, Spain, Sweden, Switzerland, and the Netherlands). While some aspects of the survey vary across countries, the core set of questions inquire about whether the firm uses robots, the intensity of robot usage, and reinvestment in new robot technology. Data currently exists for five survey rounds: 2001–2002, 2003–2004, 2006–2007, 2009–2010, and 2012–2013, and has been used in reports created by the European Commission to analyze the use of robotics and its impact on labor patterns, including wages, productivity, and offshoring.

As of now, the EMS appears to be one of the few data sources that are capturing the use of robots and automation at the firm level. This provides opportunities to analyze microeffects of robotics technology on firm productivity and labor, and to analyze firm decision-making following adoption. However, the EMS has its own limitations. The survey only considers industrial robots, and the core questionnaire only asks three questions regarding the use of robots in a factory setting. The survey is performed at the firm rather than establishment level, and the sample size of 3,000 is quite small. In contrast, the Census's Annual Survey of Manufacturers (ASM) surveys 50,000 establishments annually and 300,000 every five years.[6] Finally, similar to many other existing data sets, the EMS is purely focused on the manufacturing industry and does not address technology adoption at smaller firms with less than twenty employees.

22.4 Additional Firm-Level Research Questions

Firm-level data on the use of AI would allow researchers to address a host of questions including, but not limited to: the extent to which, and under what conditions, AI complement or substitute for labor; how AI affect firm- or establishment-level productivity; which types of firms are more or less likely to invest in AI; how market structure affects a firm's incentives to invest in AI; and how adoption is effecting firm strategies. As the nature of work itself changes with increased adoption, researchers can also investigate how firm management has been affected, particularly at the lower and middle level.

Additionally, there are many important policy questions that cannot be answered without disaggregated data. Some of these questions are related to the need to reevaluate how individuals are trained prior to entering the workforce. Without an understanding of the changes in worker experience resulting from technology adoption, it will be difficult to craft appropriate worker education, job training, and retraining programs. Further, issues related to inequality could be examined, particularly with relation to the "digital divide" and the effects of technology adoption on different demographics. There are also unanswered questions regarding the differential effects of adoption on regional economies. For example, the effects of AI on labor may be pronounced in some regions because industries, and even occupations within those industries, tend to be geographically clustered (Feldman and Kogler 2010). Thus, to the extent that AI or robots substitute for labor in certain industries or occupations, regions that rely heavily on those industries and occupations for jobs and local tax revenue may suffer. Moreover, following the recent financial crisis, unemployment insurance reserves in some states have been slow to recover (Furman 2016a). Data on the regional adoption of AI could be used to simulate the extent to which future adoption may increase unemployment and whether unemployment insurance reserves are adequately funded.

6. The census surveys all 300,000 manufacturing establishments every five years, and a rotating subsample of about 50,000 every year. See: https://www.census.gov/programs-surveys/asm/about.html.

Finally, these new technologies may have implications for entrepreneurs. Entrepreneurs may lack knowledge of how best to integrate robotics with a workforce and often face financing constraints that make it harder for them to adopt capital-intensive technologies. In the case of AI, entrepreneurs may lack data sets on customer behavior, which are needed to train AI systems. Firm-level surveys on the use of AI will help us develop a better understanding of these and related issues.

22.5 Strategies for Collecting More Data

Micro-level data regarding the adoption of AI, robots, and other types of automation can be created in a variety of ways, the most comprehensive of which would be via a census. Census data would provide information for the entire population of relevant establishments, and while the information provided would be narrow, quality is likely to be high. Additionally, data from the Census Bureau would be highly integrable with other government data sources, such as employment or labor statistics from the BLS. Data could be collected as a stand-alone inquiry, similar to the Management and Organizational Practices (MOPS) survey (see Bloom et al. 2017), or by adding questions to existing surveys, similar to work done by Brynjolfsson and McElheran (2016), which involved adding questions on data-driven decision-making to an existing census survey.

Data can also be created via a survey of firms. Survey data allows for more detailed inquiry than a census and can be carried out in a quicker and less expensive fashion. Further, a variety of organizations, both private and public, may have the interest and ability to conduct a survey regarding the adoption of AI or robotic technology. However, surveys introduce issues regarding sample selection and response rates, and depending on what organization is administering the survey, access to data can be limited or expensive.

Collecting survey data regarding the adoption of technology is not an entirely new concept. The Survey of Manufacturing Technology (SMT) was conducted by the Census Bureau in collaboration with the Department of Defense in 1988, 1991, and 1993 to measure the diffusion, use, and planned future use of new technologies in the manufacturing sector of the United States. The SMT surveyed 10,000 establishments to learn about plant characteristics and adoption of seventeen established technologies grouped into five categories: design and engineering, fabricated machining and assembly, automated material handling, automated sensors, and communication and control. Because the survey was administered by the Census Bureau, data from the SMT could easily be integrated with other firm-level data from the BLS or Census Bureau. The survey also allowed for panel analysis, as a subset of firms within the sample were respondents in multiple editions. Following the 1993 SMT, the Census Bureau discontinued the survey for funding reasons.

The Department of Defense used the SMT data to assess the diffusion of technology, and other federal agencies used the data to gauge competitiveness of the US manufacturing sector. The data were also used by the private sector in market analysis, competitiveness assessments, and planning. Multiple academic studies, including Dunne (1994), Mcguckin, Streitwieser, and Doms (1998), Doms, Dunne, and Troske (1997) and Lewis (2011) analyzed the SMT data to address questions related to productivity growth, skill-biased technical change, earnings, and capital-labor substitution, among others.

In many ways, the SMT could serve as a model for future inquiry into the adoption of robotics technology. It provided a broad look at the manufacturing industry in the United States and allowed for the examination of effects over time and for firm- and individual-level analysis when integrated with other data from the BLS or Census Bureau. However, any updated version of the SMT would need to redefine the relevant technologies, examine the intensity of use, and investigate what tasks different technologies are used for.

Private data collected at individual firms can also be a useful tool. Internal data from a firm exacerbates both the strengths and weaknesses of survey data. Data collected at a single establishment can provide an unmatched level of detail and richness compared to data created by either a census or a survey. For example, Cowgill (2016) uses detailed individual-level skill and performance data from a single establishment to assess the returns to machine-learning algorithms used in hiring decisions. However, with a sample size of one, selection on firm is a highly salient issue and generalizability may be low. Further, any data produced will almost certainly be proprietary and difficult to get access to by other researchers, making reproducibility difficult (Lane 2003).

22.6 Conclusion

The recent dramatic increases in technological capabilities we have seen in the fields of robotics and artificial intelligence provide society with a myriad of opportunities and challenges. To effectively take advantage of these technologies, we must have a complete and thorough understanding of the impacts of these technologies on growth, productivity, labor, and equality. Systematic data on the adoption and use of these technologies, particularly at the establishment level, is necessary to understand the effects of these technologies on the economy and society as a whole. The creation and aggregation of these data sets through the census, surveys conducted by public or private organizations, and internal data collected at individual firms, would provide researchers and policymakers with the tools needed to empirically investigate the impact of these technologies, and craft appropriate responses to this phenomenon.

Finally, the need for high-quality data in this area is also linked with national competitiveness, particularly in relation to crafting appropriate policy responses. Mitchell and Brynjolfsson (2017) argue that the lack of information on AI could cripple our ability to prepare for the effects of technological advancement, leading to missed opportunities and potentially disastrous consequences. For example, decisions regarding whether to tax or subsidize AI or robots rely on understanding whether or not the particular technology serves as a substitute or complement to labor. These decisions can affect adoption patterns, and if made with an incomplete understanding of the effect of these technologies on labor markets, can lead to lower economic growth, less hiring, and lower wages. In addition, data must also be utilized to properly respond to consequences stemming from technology adoption. Identifying which populations may be most vulnerable to job displacement and effectively structuring job retraining programs requires a comprehensive understanding of the microlevel impacts of adoption of these technologies.

References

Acemoglu, Daron, and Pascual Restrepo. 2017. "Robots and Jobs: Evidence from US Labor Markets." NBER Working Paper no. 23285, Cambridge, MA.

Alexopoulos, Michelle, and Jon Cohen. 2016. "The Medium Is the Measure: Technical Change and Employment, 1909–1949." *Review of Economics and Statistics* 98 (4): 792–810.

Autor, David, and Anna Salomons. 2017. "Robocalypse Now—Does Productivity Growth Threaten Employment?" Working paper, Massachusetts Institute of Technology.

Bessen, James. 2017. "Automation and Jobs: When Technology Boosts Employment." Law and Economics Paper no. 17-09, Boston University School of Law.

Bloom, Nicholas, Erik Brynjolfsson, Lucia Foster, Ron Jarmin, Megha Patnaik, Itay Saporta-Eksten, and John Van Reenen. 2017. "What Drives Differences in Management?" NBER Working Paper no. 23300, Cambridge, MA.

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

Brynjolfsson, Erik, and Kristina McElheran. 2016. "The Rapid Adoption of Data-Driven Decision-Making." *American Economic Review* 106 (5): 133–39.

Brynjolfsson, Erik, and Tom Mitchell. 2017. "What Can Machine Learning Do? Workforce Implications." *Science* 358 (6370): 1530–34.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. Forthcoming. "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?" *American Economic Association Papers and Proceedings*.

Brzeski, Carsten, and Inga Burk. 2015. "Die Roboter Kommen." ("The Robots Come.") ING DiBa Economic Research. https://www.ing-diba.de/binaries/content/assets/pdf/ueber-uns/presse/publikationen/ing-diba-economic-analysis_roboter-2.0.pdf.

Council of Economic Advisers (CEA). 2016. "Economic Report of the President." https://obamawhitehouse.archives.gov/administration/eop/cea/economic-report-of-the-President/2016.
Cowgill, Bo. 2016. "The Labor Market Effects of Hiring through Machine Learning." Working paper, Columbia University.
Dauth, Wolfgang, Sebastian Findeisen, Jens Südekum, and Nicole Wößner. 2017. "German Robots—The Impact of Industrial Robots on Workers." IAB Discussion Paper, Institut für Arbeitsmarkt- und Berufsforschung. https://www.iab.de/en/publikationen/discussionpaper.aspx.
Doms, Mark, Timothy Dunne, and Kenneth R. Troske. 1997. "Workers, Wages and Technology." *Quarterly Journal of Technology* 62 (1): 253–90.
Dunne, Timothy. 1994. "Plant Age and Technology Use in U.S. Manufacturing Industries." *RAND Journal of Economics* 25 (3): 488–99.
European Commission (EC). 2016. "Analysis of the Impact of Robotic Systems on Employment in the European Union—2012 Data Update."
Feldman, Maryann P., and Dieter F. Kogler. 2010. "Stylized Facts in the Geography of Innovation." *Handbook of the Economics of Innovation* 1:381–410.
Felten, Ed, Manav Raj, and Rob Seamans. Forthcoming. "Linking Advances in Artificial Intelligence to Skills, Occupations, and Industries." *American Economics Association Papers & Proceedings*.
Frey, Carl B., and Michael A. Osborne. 2017. "The Future of Employment: How Susceptible Are Jobs to Computerisation?" *Technological Forecasting and Social Change* 114:254–80.
Furman, Jason. 2016a. "The Economic Case for Strengthening Unemployment Insurance." Remarks at the Center for American Progress, Washington DC, July 11. https://obamawhitehouse.archives.gov/sites/default/files/page/files/20160711_furman_uireform_cea.pdf.
———. 2016b. "Is This Time Different? The Opportunities and Challenges of Artificial Intelligence." Remarks at AI Now: The Social and Economic Implications of Artificial Intelligence Technologies in the Near Term, New York University, July 7. https://obamawhitehouse.archives.gov/sites/default/files/page/files/20160707_cea_ai_furman.pdf.
Graetz, Georg, and Guy Michaels. 2015. "Robots at Work." CE P Discussion Paper no. 1335, Centre for Economic Performance.
Lane, Julia. 2003. "Uses of Microdata: Keynote Speech." In *Statistical Confidentiality and Access to Microdata: Proceedings of the Seminar Session of the 2003 Conference of European Statisticians*, 11–20. Geneva.
Lewis, Ethan. 2011. "Immigration, Skill Mix, and Capital Skill Complementarity." *Quarterly Journal of Economics* 126 (2): 1029–69.
Mandel, Michael. 2017. "How Ecommerce Creates Jobs and Reduces Income Inequality." Working paper, Progressive Policy Institute. http://www.progressivepolicy.org/wp-content/uploads/2017/09/PPI_ECommerceInequality-final.pdf.
Mann, Katja, and Lukas Püttmann. 2017. "Benign Effects of Automation: New Evidence from Patent Texts." Unpublished manuscript.
Mcguckin, Robert H., Mary L. Streitwieser, and Mark Doms. 1998. "The Effect of Technology Use on Productivity Growth." *Economics of Innovation and New Technology* 7 (1): 1–26.
McKinsey Global Institute (MGI). 2017. "Artificial Intelligence the Next Digital Frontier?" https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/how-artificial-intelligence-can-deliver-real-value-to-companies.
Mitchell, Tom, and Erik Brynjolfsson. 2017. "Track How Technology Is Transforming Work." *Nature* 544 (7650): 290–92.

National Academy of Sciences (NAS). 2017. "Information Technology and the U.S. Workforce: Where Are We and Where Do We Go from Here?" https://www.nap.edu/catalog/24649/information-technology-and-the-us-workforce-where-are-we-and.
Pajarinen, Mike, and Petri Rouvinen. 2014. "Computerization Threatens One Third of Finnish Employment." ETLA Brief no. 22, Research Institute of the Finnish Economy.