

NBER WORKING PAPER SERIES

AI AND THE ECONOMY

Jason Furman
Robert Seamans

Working Paper 24689
<http://www.nber.org/papers/w24689>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2018

We thank Wilson Powell IV for excellent research assistance and Sam Himel, Josh Lerner, Kyle Myers, Manav Raj, Scott Stern and attendees at the 2018 NBER Innovation Policy and Economy conference for helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Jason Furman and Robert Seamans. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

AI and the Economy
Jason Furman and Robert Seamans
NBER Working Paper No. 24689
June 2018
JEL No. H23,J24,J65,L1,L4,L78,O3,O4

ABSTRACT

We review the evidence that artificial intelligence (AI) is having a large effect on the economy. Across a variety of statistics—including robotics shipments, AI startups, and patent counts—there is evidence of a large increase in AI-related activity. We also review recent research in this area which suggests that AI and robotics have the potential to increase productivity growth but may have mixed effects on labor, particularly in the short run. In particular, some occupations and industries may do well while others experience labor market upheaval. We then consider current and potential policies around AI that may help to boost productivity growth while also mitigating any labor market downsides including evaluating the pros and cons of an AI specific regulator, expanded antitrust enforcement, and alternative strategies for dealing with the labor-market impacts of AI, including universal basic income and guaranteed employment.

Jason Furman
Harvard Kennedy School
79 John F. Kennedy St. Box 82
Cambridge, MA 02138
jason_furman@hks.harvard.edu

Robert Seamans
NYU Stern School of Business
44 West 4th Street, KMC 7-58
New York, NY 10012
rseamans@stern.nyu.edu

1. Introduction

Artificial intelligence (AI) has been advancing rapidly in recent years, measured both in terms of the amount of resources devoted to it and also in terms of its outputs.² *The Economist* estimated that AI-related mergers and acquisitions were 26 times larger in 2017 than in 2015.³ Increased investment has been driven by and also contributed to rapid increases in the technical capabilities of artificial intelligence (AI). For example, according to the AI Index, error rates for image recognition has dropped from 29 percent to less than 3 percent between 2010 and 2017, surpassing human performance levels.⁴ These rapid advancements apply not just AI, but also to robotics, sensors, and the connection of them all via digitization (also known as “Industry 4.0”). These advancements have started to manifest themselves in a variety of applications, including AI beating humans at complex strategy games,⁵ the creation of chatbots and virtual assistants such as Alexa and Siri,⁶ and Amazon’s new cashier-less and cash-less grocery stores.⁷

This has led both to excitement about the capability of technology to boost economic growth and to 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 2016a). Some have taken more extreme views. For example, Elon Musk has stated his belief that “AI is a fundamental risk to the existence of human civilization.”⁸

Throughout history, there has been a concern that automation, including mechanization, computing, and more recently AI and robotics, would kill jobs and generate irreversible damage to the labor market. For example, Keynes (1930) described technological unemployment as “unemployment due to our discovery of means of economising the use of labour outrunning the

² 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.

³ <https://www.economist.com/news/leaders/21739658-artificial-intelligence-pushes-beyond-tech-industry-work-could-become-fairer-or-more>

⁴ AI Index, November 2017; available: <https://aiindex.org/2017-report.pdf>

⁵ For example, in February 2016, Google’s DeepMind used its AI to beat Korean Go master Lee Se-dol <https://www.nytimes.com/2016/03/10/world/asia/google-alphago-lee-se-dol.html> and in January 2017, an AI system called DeepStack beat humans at the complex poker game Texas Hold ‘Em <https://www.scientificamerican.com/article/time-to-fold-humans-poker-playing-ai-beats-pros-at-texas-hold-rsquo-em>

⁶ <https://www.theatlantic.com/technology/archive/2011/10/siri-the-perfect-robot-for-our-time/246516/>

⁷ <https://www.nytimes.com/2018/01/21/technology/inside-amazon-go-a-store-of-the-future.html>

⁸ <https://www.cnbc.com/2017/07/17/elon-musk-robots-will-be-able-to-do-everything-better-than-us.html>

pace at which we can find new uses for labour.” Similarly, Leontief (1983), observing the dramatic improvements in the processing power of computer chips, worried that people would be replaced by machines, just as horses were made obsolete by the invention of internal combustion engines. In the past, automation has often substituted for human labor in the short term, but has led to the creation of complementary jobs in the long term (Autor 2015). Historically, automation appears to have had different effects by occupation. For example, in the 1980s and 1990s middle-skill jobs were displaced by automation, leading to labor market polarization (Autor, Kearney and Katz, 2006), though there is some evidence that labor market polarization has not continued in the last decade or two (Schmitt, Shierholz, and Mishel, 2013).

Despite the complex effects of automation on labor, there is ample evidence that, historically, automation fosters productivity growth. For example, Crafts (2004) documents the effects of steam engine technology on productivity in the UK in the 19th Century. Rosenberg (1983) and Schurr (1983) document the effect of electrification on manufacturing productivity in the early 20th Century. More recently, information technology (IT) has been credited with broad positive impacts on productivity (e.g., Oliner, Sichel, and Stiroh, 2007; Jorgenson, Ho, and Stiroh, 2008). Bloom, Sadun and Van Reenen (2012) show that better management of IT explains part of the difference in productivity between US and UK firms.

Recent productivity and labor trends highlight the importance of understanding the effect of AI on the economy. Slowing economic growth over the past decade underscores the importance of AI to deliver on its potential productivity benefits. Furman (2017) reports that 36 of 37 advanced economies had slower productivity growth in 2006-2016 compared to 1996-2006. Across these economies, growth has slowed from a 2.7 percent average growth rate in the earlier decade to a 1.0 percent average annual growth rate in the past decade. In order to boost productivity growth, it will be important to ensure that there are policies in place supporting efficient AI development and use, by both incumbent firms and startups.

Another important trend is the long term decline in the male labor force participation rate, which has fallen from a high of 98 percent in the 1950s to 89 percent in 2016 (Council of Economic Advisers 2016). This decline is concentrated among men with a high school degree or less. The decline in participation is concerning because it suggests that individuals are experiencing difficulty learning new skills and transitioning from one occupation to another (indeed, there has

also been a long term decline in labor market transitions and geographic mobility (Molloy, Smith, Wozniak, 2014)). To the extent that AI innovations lead to changes in occupations, then it will be important for the workforce to learn new skills to enable continued employment or transition to new employment. One particular concern with AI is that the changes will happen so quickly that there will be sustained periods of time in which large segments of the population are not working (see Goolsbee (2017) for a discussion of speed of adoption and Acemoglu and Restrepo (2016) for a useful model). These rapid changes, and the potential disruption to the workforce, suggest it is important that there are policies in place to support workers and retraining.

The paper proceeds as follows. The first part of Section 2 provides a number of basic facts about AI investment, robotics shipments, and patents for robots and AI. We focus some of our attention on robotics as they are easier to measure and have some clear analogies to AI. The second part of Section 2 provides a brief overview of research on the links between AI (and robotics) and economic outcomes including labor and productivity. Sections 3 and 4 then discuss specific recent policy proposals, discussing the tradeoffs and issues raised by them, with Section 3 focusing on AI and competition policy and with Section 4 focusing on AI and the labor market. Section 5 discusses additional broader questions including whether a new AI-specific agency is needed. Section 6 concludes.

2. What Do We Know?

2.1. Basic Statistics

There are multiple metrics tracking the ability of AI to perform certain specific functions. For example, as reported in Felten, Raj and Seamans (2018) seven different metrics track AI's performance for image recognition. Across all these metrics, performance has increased dramatically over the past decade. Similar increases in performance are found across multiple categories, including real-time video games, abstract strategy games (e.g., Chess, Go), video recognition, reading comprehension, translation, and others.⁹ Many of these performance increases are due to breakthroughs in various machine learning techniques, and, as described

⁹ See AI Progress Measurement from Electronic Frontier Foundation for more details, available at <https://www.eff.org/ai/metrics>.

below, these scientific breakthroughs are starting to find their way to commercial applications. However, some have argued that while there has been rapid progress on the scientific front in the past decade, there may be limits to what current techniques can accomplish (Marcus 2018).

Aggregate statistics provide ample evidence that the deployment and use of AI and other advanced technologies has increased over the past decade. The AI Index, a non-profit project designed to track activity and progress in AI, provides a number of interesting facts designed to track the scientific progress in and impact of artificial intelligence and robotics.¹⁰ For example, academic papers focused on AI have increased 9 times since 1996; in comparison, computer science papers have increased 6 times since 1996. The number of students enrolled in artificial intelligence and machine learning courses at Stanford has increased 11 times since 1996; similar trends are observed at other universities including UC Berkeley, University of Illinois, Georgia Tech, and others. The share of jobs requiring AI skills has increased almost 5 times since 2013 (and growth is especially rapid in Canada and the UK). There appears to be particularly high demand for workers with machine learning or deep learning skills. The statistics collected and published by the AI Index provide a useful snapshot of progress in AI related research and its growing impact on society, particularly the workforce.

By many measures, investment in AI, both by established firms and by venture capitalists and startups, has increased. The McKinsey Global Institute (MGI Report 2017) estimates that established firms spent between \$18 and \$27 billion on internal corporate investment in AI-related projects in 2016. Such firms also spend money on AI-related investments in the form of acquisitions. Facebook, Google, Amazon and Apple have bought up hundreds of innovative startups over the past decade, including ones that focus on AI or AI-related technologies.¹¹ MGI also notes that established firms spent \$2 to \$3 billion on AI-related M&A in 2016 alone.


While less in dollar value, investment in AI-related startups has also been increasing. An analysis of Crunchbase data by Himel and Seamans (2016) indicates an increase in venture capital funding that begins in 2012 and then accelerates sharply in 2014 (Figure 1). This observation

¹⁰ AI Index, November 2017; available: <https://aiindex.org/2017-report.pdf>

¹¹ Tech Platforms Weekly: A Closer Look at Amazon's Conduct in the Book Market; More Claims of Search Bias; Facebook, Apple, and Net Neutrality Updates; The Myspace Myth, THE CAPITOL FORUM (Jan. 20, 2017) <http://thecapitolforum.com/t/ViewEmail/j/91CFEB1924D56C52/45A74A929A973E10E663AB054A538FB>

corroborates findings reported in the MGI Report (2017) that venture capital investment in AI startups grew by 40 percent between 2013 and 2016.


Figure 1



Robots are typically referred to as an “actuated mechanism programmable in two or more axes with a degree of autonomy, moving within its environment, to perform intended tasks [ISO].”¹² The International Federation of Robotics (IFR) provides annual, aggregated statistics of the number of robots shipped by country and by industry. Figure 2 provides estimated industrial robot shipments by year, 2004-2016. The figure indicates that annual shipments were relatively flat between 2004 and 2009 before starting to rapidly increase between 2010 and 2016. Worldwide robot shipments increased about 150 percent between 2010 and 2016. The increase in robot shipments to the United States was not quite as dramatic, increasing about 100 percent between 2010 and 2016.

¹² ISO 8373, 2012, available at <https://www.iso.org/obp/ui/#iso:std:iso:8373:ed-2:v1:en>

Figure 2




This rapid increase is likely due to a combination of factors including a decrease in robot prices, an increase in robot functionality and flexibility, an increase in ease of use and interface, growing awareness of the potential cost-saving and/or value-added benefits provided by robots, and an increase in number and skill of robot integrators. Graetz and Michaels (2015) estimate that robot prices decreased 50-80 percent between 1990 and 2005. According to Green Leigh and Kraft (2017) integrators—firms which specialize in designing and building automation solutions for manufacturers—have been growing in importance in the United States: they now out-employ, outsell and outnumber robot suppliers by a margin of two to one (Green Leigh and Kraft, 2017). Membership of integrators in the Robotics Industry Association (RIA), which runs a certification program for integrators, has increased over 300 percent over the past 10 years.

Figure 3 uses data from the IFR to provide an annual breakdown of robotics shipments (flows) into selected U.S. industry sectors. In 2016, approximately half of all robot shipments were into the U.S. automotive sector; this figure has been more or less constant over time. Acemoglu and Restrepo (2017) estimate that automotive purchasers account for about 39 percent of the *stock* of robots in the United States. Shipments into the automotive sector have increased about 90 percent above 2004 levels. In 2016, about 20 percent of robot shipments were into the consumer electronics sector. This is the fastest growing sector for robot shipments; shipments into the sector have increased almost 400 percent above 2004 levels.

Figure 3

U.S. Industrial Robot Shipments by Sector, 2004-2016




The Council of Economic Advisers (CEA) provides a breakdown of robots per worker in the 2016 *Economic Report of the President* (CEA, 2016). CEA’s analysis shows that in the U.S. automotive sector there were approximately 1,091 robots per 10,000 workers in 2012. In contrast, the average of all other industries was 76 robots per 10,000 workers. The intensity of robots per worker in the United States lagged that of Japan and Germany: in 2012, there were approximately 1,563 robots per worker in the Japanese automotive sector and approximately 1,133 robots per worker in the German automotive sector.

The 2016 *Economic Report of the President* (CEA, 2016) also reports that the number and share of robotics patents granted by the U.S. Patent and Trademark Office (USPTO) has increased dramatically since 2010.¹³ Counts of patent applications with the term “artificial intelligence” in its abstract have also increased dramatically; applications in 2016 and 2017 were roughly double the average applications in 2002-2015. Figure 4 provides annual counts of applications to the USPTO that include the term “artificial intelligence” in the title or abstract.

¹³ CEA (2016) counts as a robot patent any patent that received the patent subclass number 901 (robots)

Figure 4



2.2. AI and Productivity

The best collection of current research on the link between AI and the economy appears in *The Economics of Artificial Intelligence* (“EAI”), an NBER handbook edited by Ajay Agarwal, Joshua Gans, and Avi Goldfarb. A wide variety of topics are covered, including for example the effects of AI on competition policy (Varian 2017), on innovation (Cockburn, Henderson, and Stern, 2017), on international trade (Goldfarb and Trefler, 2017), on inequality (Sachs 2017) and on productivity growth (Brynjolfsson, Rock, and Syverson, 2017), among others. A notable characteristic of the *EAI* handbook chapters is the reliance on theoretical models and aggregate, national level statistics, rather than firm-level data. This is because there is currently a paucity of data about the use of AI, robotics, and other advanced technologies at the firm level. Mitchell and Brynjolfsson (2017) and Raj and Seamans (2017) argue that more granular data is needed to better understand the effects of these technologies on workers and firm level productivity.

Economists are generally enthusiastic about the prospects of AI on economic growth. Economic literature has linked innovation to economic growth (Romer 1990). Many believe that AI and other forms of advanced automation, including robots and sensors, can be thought of as a general purpose technology (GPT) that enable lots of follow-on innovation that ultimately leads to productivity growth (Cockburn, Henderson, and Stern, 2017). However, if this theory is true, then it begs the question why, despite recent rapid technological progress in AI, there are not (yet)

corresponding increases in productivity gain. In a recent paper, Brynjolfsson, Rock and Syverson (2017) explore this question and argue this is due to a notable lag between technological progress and the commercialization of new innovative ideas building on this progress which often rely on complementary investments. The authors argue that lags of this sort are particularly notable in the case of GPTs, citing historical examples of electrification and the integrated circuit. On the other hand, Robert Gordon (2014) reminds us that even though Moore's Law has led to exponential improvement in computing performance, there has been no such analogous improvement in productivity. Moreover, Bloom, Jones, Van Reenen, and Webb (2017) document the many domains in which ideas are getting harder to find—that is larger research inputs are needed to produce additional productivity outputs.

The case for productivity growth from AI can look to empirical research on robotics for support. According to Graetz and Michaels (2015), robotics added an estimated 0.4 percentage points of annual GDP growth between 1993 and 2007 on average for the 17 countries in their sample (accounting for about one-tenth of GDP growth during this time period). The authors note that these effects are of similar magnitude to the impact of steam engines on growth in the United Kingdom. Other studies have generally found a positive effect of robotics on productivity. For example, in what appears to be the first study of robots on firm-level productivity growth, The European Commission Report on Robotics and Employment (2016) finds evidence that the use of industrial robots is correlated with significantly higher levels of labor productivity among the 3,000 manufacturing firms they surveyed.

2.3 AI and Labor


To date AI has been too small a component of the overall economy to have a significant impact on labor markets. In fact, in the last decade job growth has generally outperformed expectations while GDP growth has fallen below expectations—precisely the opposite of what would be expected if automation were replacing significant amounts of labor. As just discussed, however, AI has been growing rapidly. To the degree this leads to increases in output per hour going forward, there is still a question of whether higher productivity will result in a change in

work hours. If hours are unaffected, output would rise, but it is also possible that hours could fall, leaving output unchanged.

Three different perspectives, a theoretical perspective, an empirical/historical perspective, and attempts to make granular predictions about nascent technologies, can each offer insight into the effects of AI on labor market. A reasonable inference from these three perspectives is that, to a first approximation, AI will not be labor displacing, but could still pose significant downsides and raise other concerns.

From a theoretical perspective, innovation has four effects on labor markets. The first is that automation can directly displace labor in the affected sector. The second is that automation can create new jobs in new areas. Mandel (2017), for example, finds that job losses at brick-and-mortar department stores were more than made up for by new opportunities at fulfillment and call centers. Third, higher incomes increase demand for jobs throughout the economy, including in ways that are not directly linked to technology. For example, the share of workers in leisure and hospitality in the United States has steadily trended upward as household incomes have risen as shown in Figure 5, enabling people to afford more restaurants and travel. Finally, technology may replace specific tasks rather than entire jobs—leaving substantial room for human employment in jobs that will be changed by worker’s having a new tool at their disposal.

Figure 5



Past experiences bear out these different channels. Bessen (2018) argues 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, non-manufacturing industries. If AI is similar to other types of automation, then one might expect similar positive job creation spillovers. Dauth, Findeisen, Südekum, and Wößner (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 over-compensate for the negative employment effect in manufacturing. Other evidence is more mixed. Graetz and Michaels (2015) find a noisy effect of robot adoption in an industry on employment in that industry, whereas Acemoglu and Restrepo (2017) find a significant negative effect of robot adoption in the U.S. automotive sector on employment in that sector.

A literature has also taken the task perspective, including applying it to AI specifically. Such an approach was taken by Autor, Levy and Murnane (2003) to study how computer use affects demand for occupational skills. For example, an OECD Report (Arntz, Gregory and Zierahn, 2016) argues that there may be task variation between individuals within the same occupation. For example, managers of different firms may treat shop-floor labor differently, depending on whether they view workers as partners in the production process or as inputs into a production function (Helper, Martins, and Seamans 2018). More generally, there is much evidence that management practices vary across firms (Bloom et al. 2017), and prior research has shown that the use of technology varies by management technique (Brynjolfsson and Hitt 2000).

A recent study by Felten, Raj, and Seamans (2018) links past advances in AI to occupational abilities, and finds some evidence that the Bureau of Labor Statistics was more likely to update the definitions of occupations that were more impacted by advances in AI. Felten, Raj and Seamans (2018) provide a method that could be used by other researchers and policymakers to identify which occupations will be most affected by advances in different aspects of AI. In a related paper, Brynjolfsson, Mitchell and Rock (2018) provide a rubric for calculating which tasks are most affected by machine learning.

Part of how all of these theoretical channels operate is through relative wages. For example, a technology that replaces unskilled workers and complements skilled workers would result in a relative wage decline for unskilled workers, maintaining employment in both sectors but at a different equilibrium price. In other words, employment may be preserved but at the cost of greater inequality.

From an empirical perspective, there is both cross-sectional and time series evidence on the impact of technology on labor markets. The cross-sectional evidence is that there is no relationship between the level of productivity and the rate of employment as shown in Figure 6. Luxembourg has much higher output per hour than Italy, but this does not manifest itself in differences in hours per person—instead it means that output in Luxembourg is higher. The historical evidence also shows that the unemployment rate has not exhibited an upward trend as technology has advanced and, in fact, in the United States it has cycled around 4 or 5 percent for more than a century. The historical evidence, however, is more nuanced when it comes to labor force participation which has been declining for prime age men in the United States since the 1950s as shown in Figure 7. Most of the advanced economies have experienced declines, but the large heterogeneity of those declines with no apparent relationship to automation again strongly suggests that they are a function of much more than simply the degree of automation.

Figure 6





Figure 7



Looking forward, a number of efforts have tried to isolate how AI itself will substitute for specific occupations or tasks. The question these efforts are asking is, in effect, will the experience with AI be different than with previous technologies. This could either happen if the pace of change is much more rapid than with previous technological changes, affecting employment across the entire economy at once. In this case, while the long-run equilibrium points about labor markets may operate in theory as discussed above, in practice it may take decades for the adjustment to happen—with substantial increases in unemployment in the interim. AI could also be different if it replaces a wider-range of what had previously been uniquely human skills and abilities.

For these scenarios to play out AI would need to develop along a very different trajectory than it has to date, turning into more of a “general AI” or “true AI” that can work across the economy. To date, AI has largely been about lowering the cost of prediction through machine learning (Agrawal, Gans and Goldfarb 2018 and forthcoming). This would suggest a trajectory more like what we have seen in the past when, for example, computation became much cheaper—which is to say, a sequence of sector-specific and skill-specific disruptions without an unprecedented economywide effect.


In either case, it is important to understand the types of workers that will likely be affected. Frey and Osborne (2017) use a panel of experts to categorize tasks by their susceptibility to automation, link these tasks to occupation, employment and wage data, and find that 47 percent of

U.S. employment is at high risk of automation. In contrast, the OECD Report described above uses individual level data to predict how susceptible occupations may be to automation, and finds that 9 percent of jobs in the U.S. and across OECD countries will be highly susceptible to automation. The MGI Report (2017) estimates that at least 30 percent of activities are automatable in about 60 percent of occupations. But, the MGI Report also cautions that such automation will not necessarily substitute for labor, and reports that less than a fifth of respondents said AI was being adopted to reduce labor costs. Rather, respondents report that AI is used to improve capital efficiency or enhance existing products.

In all of these cases, however, there is a strong relationship between the occupations or skills that can be automated and income or education. CEA (2016) used the Frey and Osborne characterizations and found that jobs making less than \$20 per hour had an 83 percent probability of automation while jobs making over \$40 per hour only had a 4 percent probability of automation, as shown in Figure 8. Although the levels are very different in the OECD study, the gradient is the same—with jobs that require a high school degree or less much more likely to be automatable than jobs with a college or graduate degree, as shown in Figure 9. This highlights that going forward it is reasonable to expect that to the degree that AI does not displace labor, part of that will be because relative wages adjust, in other words that inequality rises. In addition, the pressure on lower-skilled jobs risks the continuation of the same trend that has contributed to declining labor force participation for prime-age workers.

Figure 8


Probability of Automation by an Occupation's Median Hourly Wage



Source: Council of Economic Advisers (2016).

Figure 9

Share of Jobs with Highly Automatable Skills, by Education



Source: Amtz, Gregory, and Zierahn (2016) calculations based on the Survey of Adult Skills (PIAAC) 2012.

Given the potential for AI to boost productivity on the one hand, and the potential for large disruptions to the workforce on the other, it will be important to ensure that there are appropriate policies in place. We next consider current and potential new policies that aim to address these issues.

3. AI, Antitrust and Data Portability

A number of economists document that the U.S. economy has become highly concentrated in a number of sectors (e.g., De Loecker and Eeckhout, 2017; Gutiérrez and Philippon, 2017). AI and digitization more broadly has the potential to increase competition in many ways, but at the same time, changing technology will bring new sources of concentration, including powerful network effects (e.g., Khan 2017).

3.1 Impact of the Digital Economy on Concentration Broadly

So far, internet markets have tended to favor large digital platforms that hold high market shares, a characteristic that is traditionally associated with low competition in brick-and-mortar

markets. However, understanding the competitive implications of these new markets requires a closer analysis. The markets of the digital economy are in many ways different from “old economy” markets. Some of those differences are differences of degree—the internet lowers many costs for small businesses, increasing their ability to rapidly and inexpensively scale up, collect information on potential consumers, and create new products and ideas. These differences do not transform the structure of the market; instead, they merely lower the cost of doing business. Other differences, however, are differences of type: business models may be dramatically different due to digitization. These differences of type warrant closer consideration.

One type of business model that has flourished with digitization is the “platform” model, which relies heavily on direct and indirect network effects to grow. Direct network effects—whereby the value to a customer is increasing in the number of other customers using the same platform—are particularly important for social media platforms like Facebook, Twitter and LinkedIn because the primary benefit to any customer is access to other customers. Switching costs for customers are particularly high in these markets—no one wants to be the first and only user of a social media platform—and these direct network effects can act as a barrier to entry.

Indirect network effects—where the value to a customer is increasing in the number of customers on the other side of the platform—are also important. As a result of these indirect network effects, companies may subsidize one side of the market by profiting from the other side of the market (Rochet and Tirole, 2003; 2006). For example, social media sites often offer free services to users and charge for ads. This provides a challenge when trying to determine the optimal level of competition in these new markets. Usually, economists use prices as indicators of the level of competition, but price is insufficient in this case given the low prices on one side of the platform’s market. The lack of high prices for consumers does not mean that consumer harms or other risks could not occur. Industry watchers have raised concerns about whether the large companies that dominate search and social networking may be able to acquire inefficient power in ads or control people’s access to news. Another concern is that instead of raising prices or reducing quantity, these companies may reduce innovation. Firms holding quasi-monopolies may lose the incentive to keep improving the quality of their products (Arrow 1962).

Switching costs are traditionally an indicator of competition, and many may assume that switching costs in internet markets are virtually zero because competition is just a “click” away.

This may have been true in the early ages of the internet, but may be less true now. For example, the original search engines were merely directories of websites, and their quality did not depend on how many users they had. However, search engines today collect data on the behavior of their users and use it to improve their services and tailor those services to individual users. Thus, in order for other firms to be competitive, they need a large user base and the data that comes with it. Furthermore, for each individual user looking to switch services, the incumbent, with its existing knowledge of that user, has a significant advantage over a competitor that does not yet know the user and therefore cannot tailor services to him or her.

Lastly, digitization could bring a new level of opacity to businesses. Traditionally, price fixing and collusion could be detected in the communications between businesses. The task of detecting undesirable price behavior becomes more difficult with the use of increasingly complex algorithms for setting prices. This type of algorithmic price setting can lead to undesirable price behavior, sometimes even unintentionally. The use of advanced machine learning algorithms to set prices and adapt product functionality would further increase opacity.

3.2 Large Datasets as a Barrier to Entry

A related concern is the effect that concentration of digital assets among a few dominant platforms may have on AI-enabled startups. Large datasets are a critical input for firms that want to create or use AI systems. Even the best AI algorithms are useless without a large dataset because these datasets are needed for the initial training and fine-tuning of AI algorithms. This is a particular concern given the hope that development and commercialization of AI-related products and services will result in greater productivity growth. Absent competition from startups and other entrants, the economy may get less of this productivity growth.

Antitrust enforcement officials in the U.S. and Europe recognize the challenges that may arise when large technology firms control the vast majority of such data. For example, FTC Commissioner Terrell McSweeney says, “It may be that an incumbent has significant advantages over new entrants when a firm has a database that would be difficult, costly, or time consuming

for a new firm to match or replicate.”¹⁴ However, we still need more research to understand the conditions under which large amounts of data may provide advantages to incumbents over entrants. For example, using data from Amazon, Bajari et al (2018) provide evidence that data on larger number of products does not improve forecast errors, but data over more time periods for these products does (though at a diminishing rate). Similarly, Chiou and Tucker (2017) find little evidence that search engine’s use of longer data series allows for better search results. The implications of these studies are twofold: first, there is limited evidence of increasing returns to scale for data and, second, restrictions on data storage and use are not likely to harm large technology platforms. As Chiou and Tucker state: “Our results also suggest that limits on data retention may impose fewer costs in instances where overly long data retention leads to privacy concerns such as an individual’s ‘right to be forgotten.’”

To the extent that data is a critical resource for entrants, new enforcement policies and regulatory strategies such as data portability may be needed to ensure that both incumbent and potential entrant firms have access to the datasets they need to innovate in the AI domain. However, as pointed out by Himel and Seamans (2017), there are a range of existing policies and other approaches that may be useful in aiding entry of AI-enabled startups, including litigation strategies against large technology platforms alleging anticompetitive conduct or consumer harm. As describe above, a challenge in such cases is the two-sided nature of the platforms, on which the consumer side typically involves a low or “free” price. There is a need, therefore, for research that carefully assesses the benefits and drawbacks of using existing policies, regulations, and litigation strategies to address the myriad concerns that arise from the rise of AI.

Data portability allows customers to take their data from one provider to another. The concept is not unique to digital platforms; it could be used for banking data as well, for example. The idea is similar to telephone-number portability from the Telecom Act of 1996 which gave customers a greater ability to leave a telephone company for a rival. Guy Rolnick and Luigi Zingales from the University of Chicago argued in an op-ed in the *New York Times* for the

¹⁴ Commissioner Terrell McSweeney, Opening Remarks for a Panel Discussion, “Why Regulate Online Platforms?: Transparency, Fairness, Competition, or Innovation?” at the CRA Conference in Brussels, Belgium, at 5 (Dec. 9, 2015), https://www.ftc.gov/system/files/documents/public_statements/903953/mcsweeney_-_cra_conference_remarks_9-12-15.pdf.

portability of social graphs which would allow users to take all of the digital connections they create on a platform like Facebook to another rival platform.¹⁵

In principle, data portability helps increase competition between established firms in the market, because any potential customer could easily shift her data from one established firm to another. However, it is unlikely that data portability alone would increase competition from startups that may need access to large datasets to train their AI algorithms. Under a data portability model, startups would need to induce multiple individual users to port their data to the startup.

Another issue with data portability is where the customer's data would reside, which has implications for the data's security, and for the consumer's privacy. One issue around privacy is the extent of control that a consumer has not just over their own sensitive information, which they may choose or not to share with different companies, but over the inferences that an AI algorithm can make about the consumer by comparing patterns in the individual behavior to patterns seen across similar populations of individuals. Law and economics scholars have provided in depth treatments of privacy in a digital age (Tucker 2012; Calo 2017).

Trusted third parties can potentially play a role to safeguard consumer information while allowing conditional access to large datasets for AI-enabled startups. There may be other benefits to use of trusted third parties, including the creation of standardized training datasets.¹⁶ For example, Mitchell and Brynjolfsson (2017) argue that when AI-related data is collected from diverse sources and integrated together it may allow for the identification of bias or skew in the data. An open question is “who” exactly would play the role of the trusted third party. Potential solutions include a university consortium, an existing or new government agency—the running of which would be a potential role for an “AI-specific Agency” explored in a later section—or a public-private partnership or non-profit entity whose sole mission is the collection, curation, protection of large scale datasets. More generally, the competition policy issues described above would be an area that could potentially fall under the domain of an AI-specific Agency.

¹⁵ <https://www.nytimes.com/2017/06/30/opinion/social-data-google-facebook-europe.html>

¹⁶ This is one of several recommendations made in the AI Now 2017 Report: https://ainowinstitute.org/AI_Now_2017_Report.pdf

4. AI and Labor Markets: UBI, wage supplements and guaranteed employment

AI has the potential to continue or possibly even exacerbate trends towards declining labor force participation and increased inequality. Fears about these changes have helped to motivate and expand interest in the centuries old idea of a Universal Basic Income (UBI) that would partially or completely replace existing safety net programs with a single, unconditional cash transfer to every adult in the United States. Although less often framed as a response to AI, two other larger ideas that have grown in interest in recent years could serve as alternatives to UBI: the first is a large-scale system of wage subsidies to help create an incentive for work and increase the reward to work (e.g., Phelps 1997) and the second is guaranteed Federal employment (e.g., Mitchell 1996; Mosler 1998). This section reviews some of the pros and cons of these three ideas.

In understanding how policy should respond to labor market changes, it is important to understand whether artificial intelligence is more like a macroeconomic shock or a series of sector-specific shocks. The more it is like a series of sector-specific shocks, the more that a response should be targeted and focused, building on successful past efforts rather than an unprecedented approach.

4.1 Universal Basic Income

A universal basic income has three characteristics. First it is available nearly universally, limited only by criteria such as citizenship or possibly age, but generally not by income or by other factors, like disability or employment status. Second, it provides cash, not in-kind benefits like many existing programs that target food, heating oil, housing and the like. Third, it is unconditional so unlike many existing programs does not require the applicant to be working, looking for work, attending school, or other forms of conditioning. Most UBI proposals are also intended to be substantial enough to raise people to something like the poverty line.

For example, Andrew Stern and Lee Kravitz (2016) have proposed \$12,000 for every non-elderly adult U.S. citizen while Charles Murray has proposed \$10,000 for every U.S. citizen aged 21 and older (2006). Generally, UBI proposals have focused on replacing existing means-tested programs (e.g., food assistance, housing subsidies, and cash welfare) but differ in whether they

would also replace Social Security retirement income and health programs like Medicare and Medicaid. With these offsets a UBI would require about \$1 trillion in additional financing annually, for example through revenue increases.

Some of the arguments for UBI are unrelated to technological developments, for example the claim that a single streamlined program would be more efficient to administer than a variety of programs with different rules and conditions. In addition, some of the claims (e.g. Thomas Paine 1797) are more grounded in moral premises than specific efficiency statements. These arguments, however, have been advanced with greater force along with new arguments related to AI and the future of work. One argument is that UBI would be a solution to mass joblessness, helping to ensure a basic income floor for people. A related argument is that UBI is a way to share the benefits of the increased output associated with AI. In addition, another argument made by Mark Zuckerberg (and echoed by many others) is that UBI can help to “give everyone a cushion to try new things.” (Zuckerberg 2017).

UBI also has a number of drawbacks. The first is the cost of the program. An additional roughly \$1 trillion in annual financing would require approximately doubling existing payroll taxes or a roughly 50 percent increase in existing individual income taxes. Given the limited tax appetite, particularly in the United States, the feasibility of such a substantial program may be called into doubt.

A second question about UBI is whether or not it is the optimal way to handle the tradeoffs inherent in targeting. UBI is likely to result in smaller net transfers to households that get larger transfers today, either by virtue of their lower incomes, larger number of children, or situation—for example, disability status. This may raise equity concerns for some. In addition, it could also raise efficiency issues. For example, unemployment insurance receipt is conditional on being unemployed—which increases moral hazard but also helps households smooth their consumption. Shifting to unconditional transfers would reduce the moral hazard associated with substitution effects, but at the expense of the consumption smoothing currently provided by the existing safety net.

A third issue is with the argument that UBI may stimulate entrepreneurship and innovation. There are strong arguments for policies that stimulate entrepreneurship, given that rates of new

business entry have declined over the past three decades (Decker, Haltiwanger, Jarmin, and Miranda 2014). But there is little evidence of heightened entrepreneurship and innovation in regions with UBI-like programs, such as oil-rich areas which provide income to most residents, including in Alaska, Norway and some Gulf states. Moreover, the argument for increased entrepreneurship and innovation relies on the assumption that entrepreneurs and inventors will take more risk. But there may be a countervailing force, in that lenders, wary of increased risk-taking by entrepreneurs, may cut back on lending. There is some evidence of this type of behavior from the economics of bankruptcy literature (e.g., Berkowitz and White, 2004).

4.2 Employment Subsidies

An alternative to UBI is employment subsidies. These would also be provided in cash, but unlike UBI they would be not be universally available and would be conditional on work, and potentially other circumstances. Currently the largest such program is the Earned Income Tax Credit (EITC). For a head of household with two children, the EITC provides a subsidy of \$0.40 for every \$1 earned up to a maximum subsidy of \$5,616 in 2017. The subsidy starts phasing out at \$0.2106 for each \$1 earned above \$18,340 in 2017 and is eliminated when household income reaches \$45,007. These parameters are adjusted to result in a somewhat smaller subsidy for households with 1 child, a somewhat larger subsidy for households with 3 or more children, and a very small subsidy—with a maximum of \$510—for households without qualifying children.

There are two types of proposals to expand employment subsidies. The first would use the EITC-structure, administering the subsidy through the tax code, paying it directly to households, and making it contingent on household circumstances. One proposal, made by both House Speaker Paul Ryan and President Barack Obama, would have been to substantially increase the EITC for households without qualifying children. These proposals are grounded in the academic evidence that the EITC increases participation in the workforce (e.g., Eissa and Liebman 1996 and Hotz, Mullin, and Scholz 2006) as well as distributional concerns.

An alternative approach was put forward by Phelps (1997) and would be a subsidy for employers, a considerably larger version of the Work Opportunity Tax Credit (WOTC) in the law today. This employment subsidy could, for example, provide an additional \$7 per hour for

households making the minimum wage of \$7.25 per hour, phasing down the subsidy as earnings rose. All of the administration would be undertaken by the employer and the tax authority, with the employee just seeing a higher wage. The alternative approach is partly motivated by liquidity concerns and administrative concerns, ensuring that the worker would be paid regularly with no administrative effort on her part. In addition, it could address any stigma associated with taking advantage of a tax and transfer program. On the other hand, wage subsidies would lose the ability to target—for example, based on overall household income or circumstances.

Like UBI, employment subsidies would also have a fiscal cost that would need to be financed in some manner. Moreover, because employment subsidies are conditional, administrative costs would be higher relative to UBI and also increases the incentives for fraud, for example misreporting earnings or hours. Because the subsidies would be conditional on work and would phase out at higher incomes, the overall cost would be considerably smaller than UBI—with some proposals to expand the EITC for households without qualifying children costing about \$5 billion annually. Unlike UBI, however, employment subsidies would act as an incentive for employment and would send a signal that work was still central to social support.

4.3 Guaranteed Employment

Support is also growing for some sort of guaranteed employment, potentially a Federal backstop or guaranteed job that would provide payment but only in exchange for labor services. A few specific proposals or concepts have been put forward including by Center for American Progress and Jeff Spross (Paul, Darity, and Hamilton 2018; Spross 2017). These proposals have generally not been motivated by AI, and in fact use 1930s era programs as an inspiration, but they would functionally provide a third approach for addressing the concerns raised by AI on labor force participation.

Like UBI and wage subsidies, such a proposal could have a substantial fiscal cost. If the guaranteed job paid \$15 per hour and 10 million people took it up, then that would be about \$300 billion or more in annual costs, although it is possible that the associated labor would also yield some public benefits. Such an employment guarantee would have the advantage of being the most direct way to subsidize work, potentially keeping people in the labor force and improving

countercyclical fiscal responses. On the other hand, such a program would also have to grapple with substantial administrative complexity, the danger of trapping people in lower-wage jobs without career advancement prospects, and the distortions to labor markets.

Additional research and experimentation with UBI, employment subsidies, and guaranteed employment would all be worthwhile, especially in understanding how these, in some cases long-standing proposals, would interact with likely labor market changes going forward.

5. Is an AI-Specific Agency Needed?

As noted by Ryan Calo (2017), an overarching policy challenge is how to best introduce expertise about AI, robotics and other advanced technologies into all branches and levels of government to aid decision making. This need has led to calls for specific commissions on AI and robotics (e.g., Calo 2014) or even a new agency tasked with overseeing AI, either to maximize its benefits or minimize any associated harms.

Calo (2014) reminds us that the Federal government forms new departments and agencies from time to time, depending on circumstance. For example, the Department of Homeland Security (DHS) was formed in 2002 in response to the September 11, 2001 terrorist attacks, and several existing agencies such as Immigration and Naturalization Service were re-organized under DHS instead of the Department of Justice (DOJ). Likewise, the Department of Energy was formed by consolidating the Energy Research and Development Administration, Federal Power Commission, and Federal Energy Administration in 1977, in response to the oil crisis earlier in the 1970s.

There are multiple challenges with creating a new commission or agency, including defining the mission or scope of the new agency and reorganizing existing agencies. When considering the mission of the agency, it is useful to consider first whether the agency would have any enforcement authority, similar to Federal Trade Commission (FTC), DOJ or Securities and Exchange Commission (SEC), or whether the agency would instead provide more of an advisory function. In the first case, the relevant thought experiment is whether the existing enforcement agencies could perform a similar role or not; that is, is there something specific about robotics and AI that necessitates a dedicated enforcement authority. In the second case, the relevant thought

experiment is whether a standalone agency is needed to provide an advisory function, or whether each agency can address its own perceived needs via hiring of dedicated staff as the FTC did in creating the position of “Chief Technologist” and as the Obama White House did when it created the Chief Technology Officer position within the Office of Science and Technology Policy. In some cases this may require the creation of a “technology office” if one does not already exist. For example, there have been recent calls for Congress to revive the Office of Technology Assessment (Graves and Kosar, 2018). Regardless of the bureaucratic structure, there is a need for the Federal government to assess existing policy and evaluate proposals for new policy tools as we have done in the preceding sections.

Other key questions in evaluating this issue are whether AI should be thought of as a new area or instead as a tool that is used in a variety of areas. From the former perspective, one agency would have to distinguish between automation due to AI and automation due to other causes and apply its analysis in a disparate set of domains. From the latter perspective, the goal would be to get AI experts in key positions in, for example, the National Highway Traffic Safety Administration (NHTSA) and the SEC to look at automated driving systems and automated trading respectively—rather than to have one body that looks at automation applied to both areas.

6. Conclusion

Artificial intelligence has the potential to dramatically change the economy. On the one hand, the potential for increased productivity growth is welcome given the decades-long slowing in productivity growth in the United States and other advanced economies. On the other hand, the potential for AI-induced labor disruptions could potentially exacerbate existing problems in the labor force, including the decades-long decline in male labor force participation rate. Economic research has only started to assess these issues. Early research findings suggest that AI and robotics do indeed boost productivity growth, and that effects on labor are mixed. However, more empirical research is needed in order to confirm existing findings on the productivity benefits, better understand conditions under which AI and robotics substitute or complement for labor, and understand regional level outcomes. For these reasons, a number of others have called for systematic collection and dissemination of establishment level data (e.g., Raj and Seamans, 2017;

Mitchell and Brynjolfsson, 2017) to address the need for publicly available data on the deployment and use of robotics and AI in manufacturing and service establishments.

A variety of policy solutions—ranging from an AI-specific commission, to data portability, to UBI and other strategies—have been suggested to address actual and perceived issues arising from increased use of AI and robotics in the economy. Any assessment of these policies should compare how they might address potential AI-related issues relative to current policies. For universal basic income in particular, it would appear that there are a number of other proven and effective policies, such as an expansion of the Earned Income Tax Credit or the establishment of wage subsidies, might achieve the goals of increasing labor participation with fewer spillovers on other aspects of the economy. Also, while we have assessed data portability in light of the need to continue to spur growth and investment in AI, it is worth pointing out that it might address other competition policy issues more generally, such as the increase in concentration observed across a number of industries and markets.

When weighing the tradeoffs of various policy approaches, it will be useful to consider the speed with which AI may or may not affect the economy. As highlighted above, on one hand, AI's performance in certain limited areas, such as image recognition and abstract strategy games, has improved dramatically in recent years. On the other hand, AI may be hitting performance limits (Marcus 2018), and commercial applications have not yet had any dramatic impact on economic productivity (Brynjolfsson, Rock and Syverson 2017). Traditional safety net programs may therefore be well suited to address the transitory dislocations that may arise from AI in the short and medium term, and may be particularly appealing given scarce Federal resources.

References

- Acemoglu, Daron, and Pascual Restrepo. 2016. "The race between machine and man: Implications of technology for growth, factor shares and employment." National Bureau of Economic Research working paper No. w22252.
- _____. 2017. "Robots and Jobs: Evidence from U.S. Labor Markets." NBER Working Paper 23285.
- Agrawal, Ajay, Joshua S. Gans and Avi Goldfarb. 2018. *Prediction Machines: The Simple Economics of Artificial Intelligence*. Cambridge, MA: Harvard Business Review Press.
- Agrawal, Ajay, Joshua S. Gans and Avi Goldfarb. Forthcoming. "Introduction to 'The Economics of Artificial Intelligence: An Agenda'." in *The Economics of Artificial Intelligence*, eds. Ajay Agrawal, Joshua S. Gans and Avi Goldfarb, Chicago: University of Chicago Press (forthcoming).
- 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 Papers, No. 189.
- Arrow, Kenneth J. 1962. "Economic Welfare and the Allocation of Resources for Inventions." In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, edited by R.R. Nelson, 609-626. Princeton, NJ: Princeton University Press.
- Autor, David H. 2015. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives* 29(3): 3-30.
- Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The skill content of recent technological change: An empirical exploration." *The Quarterly Journal of Economics* 118(4): 1279-1333.
- Autor, David H, Lawrence F. Katz, and Melissa S. Kearney. 2006. "The Polarization of the U.S. Labor Market." *American Economic Review Papers and Proceedings* 96(2): 189-194.
- Bajari, Patrick, Victor Chernozhukov, Ali Hortaçsu, Junichi Suzuki. 2018. "The Impact of Big Data on Firm Performance: An Empirical Investigation." National Bureau of Economic Research working paper No. 24334.
- Berkowitz, Jeremy, and Michelle J. White. 2004. "Bankruptcy and small firms' access to credit." *RAND Journal of Economics* 35(1): 69-84.
- Bessen, James. 2018. "AI and Jobs: the Role of Demand." National Bureau of Economic Research working paper No. 24235.

- Bloom, Nicholas, Raffaella Sadun, and John Van Reenen. 2012. “Americans do I.T. Better: US Multinationals and the Productivity Miracle.” *American Economic Review* 102(1): 167–201.
- Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb. 2017. “Are Ideas Getting Harder to Find?” National Bureau of Economic Research working paper No. 23782.
- Bloom, Nicholas, et al. 2017. “What Drives Differences in Management.” 2017. National Bureau of Economic Research working paper No. 23300.
- Brynjolfsson, Erik, and Lorin M. Hitt. 2000. “Beyond computation: Information technology, organizational transformation and business performance.” *The Journal of Economic Perspectives* 14.4: 23-48.
- Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. 2018. “What Can Machines Learn, and What Does It Mean for the Occupations and Industries.” Working paper.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2017. “Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics.” in *The Economics of Artificial Intelligence*, eds. Ajay Agrawal, Joshua S. Gans and Avi Goldfarb, Chicago: University of Chicago Press (forthcoming).
- Bughin, Jacques, et al. (MGI Report). 2017. “Artificial Intelligence: The Next Digital Frontier?” McKinsey Global Inst., June 2017. Available: <http://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/how-artificial-intelligence-can-deliver-real-value-to-companies>.
- Calo, Ryan. 2014. The Case for a Federal Robotics Commission. Brookings Center for Technology Innovation. <https://www.brookings.edu/research/the-case-for-a-federal-robotics-commission/>
- _____. 2017. “Artificial Intelligence Policy: A Primer and Roadmap.” Available at SSRN: <http://dx.doi.org/10.2139/ssrn.3015350>
- Chiou, Lesley and Catherine E. Tucker. 2017. “Search Engines and Data Retention: Implications for Privacy and Antitrust.” NBER Working Paper No. w23815.
- Cockburn, Iain, Rebecca Henderson, and Scott Stern. 2017. “The Impact of Artificial Intelligence on Innovation.” in *The Economics of Artificial Intelligence*, eds., Ajay Agrawal, Joshua S. Gans and Avi Goldfarb, Chicago: University of Chicago Press (forthcoming).
- Council of Economic Advisers (CEA). 2016. *Economic Report of the President*. <https://obamawhitehouse.archives.gov/administration/eop/cea/economic-report-of-the-President/2016>
- Crafts, Nicholas .F.R. 2004. “Steam as a General Purpose Technology: A Growth Accounting Perspective.” *The Economic Journal* 114 (495): 338–351.

- Dauth, Wolfgang, Sebastian Findeisen, Jens Südekum, Nicole Wößner. 2017. “German Robots – The Impact of Industrial Robots on Workers.” IAB Discussion Paper.
- Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda. 2014. “The role of entrepreneurship in US job creation and economic dynamism.” *Journal of Economic Perspectives* 28(3): 3-24.
- De Loecker, Jan. and Jan Eeckhout. 2017. “The rise of market power and the macroeconomic implications.” National Bureau of Economic Research working paper No. 23687.
- Eissa, Nada and Jeffrey B. Liebman. 1996. “Labor Supply Response to the Earned Income Tax Credit.” *Quarterly Journal of Economics* 111(2): 605-37.
- European Commission (EC). 2016. “Analysis of the Impact of Robotic Systems on Employment in the European Union – 2012 Data Update.”
- Felten, Ed, Manav Raj, and Robert Seamans. 2018. “A Method to Link Advances in Artificial Intelligence to Occupational Abilities.” *American Economic Association Papers & Proceedings* 108: 54-57.
- Furman, Jason. 2016a. “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, 2016. Available at: https://obamawhitehouse.archives.gov/sites/default/files/page/files/20160707_cea_ai_furman.pdf
- _____. 2016b. “The Economic Case for Strengthening Unemployment Insurance.” Remarks at Center for American Progress, Washington DC, July 11, 2016. Available at: https://obamawhitehouse.archives.gov/sites/default/files/page/files/20160711_furman_uireform_cea.pdf
- _____. 2017. “Should We Be Reassured If Automation in the Future Looks Like Automation in the Past?” in *The Economics of Artificial Intelligence*, eds. Ajay Agrawal, Joshua S. Gans and Avi Goldfarb, Chicago: University of Chicago Press (forthcoming).
- Frey, Carl B. and Michael A. Osborne. 2017. “The Future of Employment: How Susceptible Are Jobs to Computerisation?” *Technological Forecasting and Social Change* 114: 254–280.
- Goldfarb, Avi and Daniel Trefler. 2017. “AI and International Trade.” in *The Economics of Artificial Intelligence*, eds., Ajay Agrawal, Joshua S. Gans and Avi Goldfarb, Chicago: University of Chicago Press (forthcoming).
- Goolsbee, Austan. 2017. “Public Policy in an AI Economy.” in *The Economics of Artificial Intelligence*, eds., Ajay Agrawal, Joshua S. Gans and Avi Goldfarb, Chicago: University of Chicago Press (forthcoming).

- Gordon, Robert J. 2014. "The demise of US economic growth: restatement, rebuttal, and reflections." National Bureau of Economic Research working paper No. 19895.
- Graetz, Georg and Guy Michaels. 2015. "Robots at Work." *Centre for Economic Performance Discussion Paper No. 1335*.
- Graves, Zach and Kevin Kosar. "Bring in the Nerds: Reviving the Office of Technology Assessment." *R Street Policy Study No.128*.
- Green Leigh, Nancey, and Benjamin R. Kraft. 2017. "Emerging robotic regions in the United States: insights for regional economic evolution." *Regional Studies*: 1-13.
- Gutiérrez, Germán. and Thomas Philippon. 2017. "Declining Competition and Investment in the US" National Bureau of Economic Research No. 23583.
- Helper, Susan, Raphael Martins, and Robert Seamans 2018. "Value Migration and Industry 4.0: Theory, Field Evidence, and Propositions." New York University working paper.
- Hilary Hoynes. "A Revolution in Poverty Policy: The Earned Income Tax Credit and the Well-Being of American Families." *Pathways* Summer 2014, pp. 23-27.
- Himel, Samuel and Robert Seamans. 2017. "Artificial Intelligence, Incentives to Innovate, and Competition Policy." *Antitrust Chronicle*. Fall 2017 Vol 1(3).
- Hotz, V. Joseph, Charles H. Mullin, and John Karl Scholz. 2006. "Examining the Effect of the Earned Income Tax Credit on the Labor Market Participation of Families on Welfare." National Bureau of Economic Research working paper No. 11968.
- Jorgenson, Dale W., Mun S. Ho, and Kevin J. Stiroh. 2008. "A Retrospective Look at the U.S. Productivity Growth Resurgence." *Journal of Economic Perspectives* 22(1): 3–24.
- Keynes, John Maynard. 1930. "Economic Possibilities for our Grandchildren." In *Essays in Persuasion (2010)*. London: Palgrave Macmillan.
- Khan, Lina M. 2017. "Amazon's Antitrust Paradox." *Yale Law Journal* 126: 710-722.
- Leontief, Wassily. 1983. "National Perspective: The Definition of Problems and Opportunities." In *The Long-term Impact of Technology on Employment and Unemployment: a National Academy of Engineering Symposium, June 30, 1983*. Washington, DC: National Academy Press.
- Mandel, Michael. 2017. "How Ecommerce Creates Jobs and Reduces Income Inequality." Progressive Policy Institute. Available at: http://www.progressivepolicy.org/wp-content/uploads/2017/09/PPI_ECommerceInequality-final.pdf
- Marcus, Gary. 2018. "Deep Learning: A Critical Appraisal." arXiv:1801.00631

- Mitchell, Tom, and Erik Brynjolfsson. 2017. "Track how technology is transforming work." *Nature* 544(7650): 290-292.
- Mitchell, William F. 1996. "Inflation and Unemployment: A Demand Story." Presented to the European Unemployment Conference, sponsored by the European Commission, at the European University Institute, Florence, November 1996.
- Molloy, Raven, Christopher L. Smith, and Abigail K. Wozniak. 2014. "Declining migration within the US: the role of the labor market." National Bureau of Economic Research working paper No. 20065.
- Mosler, Warren. 1998. "Full Employment and Price Stability." *Journal of Post Keynesian Economics* 20(2): 167-183.
- Murray, Charles A. 2006. *In Our Hands: a Plan to Replace the Welfare State*. Washington, D.C.: AEI Press.
- Oliner, Stephen D., Daniel E. Sichel, and Kevin J. Stiroh. 2007. "Explaining a Productive Decade." *Brookings Papers on Economic Activity* 1: 81–137.
- Paine, Thomas. 1797. "Agrarian Justice."
- Paul, Mark, William Darity, Jr., and Darrick Hamilton. 2018. "The Federal Job Guarantee – A Policy to Achieve Full Employment." Center on Budget and Policy Priorities. Report. March 9. Available at: <https://www.cbpp.org/research/full-employment/the-federal-job-guarantee-a-policy-to-achieve-permanent-full-employment>
- Phelps, Edmund S. 1997. *Rewarding Work: How to Restore Participation and Self-support to Free Enterprise*. Cambridge, MA: Harvard University Press.
- Raj, Manav, and Robert Seamans. 2017. "AI, Labor, Productivity and the Need for Firm-Level Data." in *The Economics of Artificial Intelligence*, eds., Ajay Agrawal, Joshua S. Gans and Avi Goldfarb, Chicago: University of Chicago Press (forthcoming).
- Rochet, Jean-Charles, and Jean Tirole. 2003. "Platform competition in two-sided markets." *Journal of the European Economic Association* 1(4): 990-1029.
- _____. 2006. "Two-sided markets: A progress report." *The RAND Journal of Economics* 37(3): 645-667.
- Romer, Paul M. 1990. "Endogenous Technological Change." *Journal of Political Economy* 98(5): S71-S102.
- Rosenberg, Nathan. 1983. "The Effects of Energy Supply Characteristics on Technology and Economic Growth." In *Energy, Productivity, and Economic Growth*, eds., Sam Schurr, Sidney

Sonenblum, and David O. Wood, Cambridge, MA: Oelgeschlager, Gunn, and Hain.

Sachs, Jeffrey. 2017. "R&D, Structural Transformation, and the Distribution of Income." in *The Economics of Artificial Intelligence*, eds., Ajay Agrawal, Joshua S. Gans and Avi Goldfarb, University of Chicago Press: Chicago (forthcoming).

Schmitt, John, Heidi Shierholz, and Lawrence Mishel. 2013. "Don't Blame the Robots: Assessing the Job Polarization Explanation of Growing Wage Inequality." EPI-CEPR Working Paper.

Schurr, Sam. 1983. "Energy Efficiency and Economic Efficiency." In *Energy, Productivity, and Economic Growth*, eds., Sam Schurr, Sidney Sonenblum, and David O. Wood, Cambridge, MA: Oelgeschlager, Gunn, and Hain.

Spross, Jeff. 2017. "You're Hired!" *Democracy* Spring 2017.

Stern, Andy, and Lee Kravitz. 2016. *Raising the Floor: How a Universal Basic Income Can Renew Our Economy and Rebuild the American Dream*. New York: Public Affairs.

Tucker, Catherine. 2012. "Privacy and Innovation." *Innovation Policy and the Economy*, Vol. 11, eds. Josh Lerner and Scott Stern, National Bureau of Economic Research.

Varian, Hal. 2017. "Artificial Intelligence, Economics, and Industrial Organization." in *The Economics of Artificial Intelligence*, eds., Ajay Agrawal, Joshua S. Gans and Avi Goldfarb, Chicago: University of Chicago Press (forthcoming).

Zuckerberg, Mark. 2017. "Mark Zuckerberg's Commencement Address at Harvard." May 25. Available at: <https://news.harvard.edu/gazette/story/2017/05/mark-zuckerbergs-speech-as-written-for-harvards-class-of-2017/>.