Is This Time Different? The Opportunities and Challenges of Artificial Intelligence

Jason Furman
Chairman, Council of Economic Advisers

Remarks at *AI Now: The Social and Economic Implications of Artificial Intelligence Technologies in the Near Term*

New York University
New York, NY
July 7, 2016

This is an expanded version of these remarks as prepared for delivery.

Much of the debate about the economic impact of artificial intelligence (AI) centers on the question of whether this time will be different. Some optimists argue that AI is no different than technologies that came before it and that centuries of fears that machines will replace human labor have proven unfounded, with machines instead creating previously unimagined jobs and raising incomes. Others argue that AI is different—by replacing cognitive tasks, it could render much of human employment redundant, leading to mass unemployment in the eyes of the pessimists or historically unparalleled freedom for leisure in the eyes of the optimists. As I will argue this morning, I see little reason to believe that the economic impact of AI will be very different from previous technological advances. But unlike many of the optimists, I do not find that similarity fully comforting, as technological advances in recent decades have brought tremendous benefits but have also contributed to increasing inequality and falling labor force participation. However, as I will emphasize this morning, the effects of technological change on the workforce are mediated by a wide set of institutions, and as such, policy choices will have a major impact on actual outcomes. AI does not call for a completely new paradigm for economic policy—for example, as advocated by proponents of replacing the existing social safety net with a universal basic income (UBI) —but instead reinforces many of the steps we should already be taking to make sure that growth is shared more broadly.

But before turning to concerns about some of the possible side effects from AI, I want to start with the biggest worry I have about it: that we do not have enough of AI. Our first, second and third reactions to just about any innovation should be to cheer it—and ask how we get more of it, the issue I will discuss first in my remarks. But I will then discuss the potential labor market downsides of AI. Finally, I will conclude with the role of public policy in addressing these issues—both helping to advance AI while ensuring that more people share in the benefits of it, two goals that are ultimately complementary.

Why We Need More Artificial Intelligence

I sometimes experience whiplash moving back and forth between conversations of economists who are worried about the lack of measured productivity growth and technologists who see transformative change all around us. It may not shock you that I am with the economists on this

one. Measured productivity growth has slowed in 30 of the 31 advanced economies, slowing from a 2 percent average annual growth rate from 1994 to 2004 to a 1 percent average annual growth rate from 2004 to 2014. Notably, the United States still has the fastest productivity growth of any G-7 country, with annual productivity growth of 1.1 percent from 2004 to 2014 as compared to 2.3 percent from 1994 to 2004, as shown in Figure 1.

Figure 1
Labor Productivity Growth, G-7 Countries

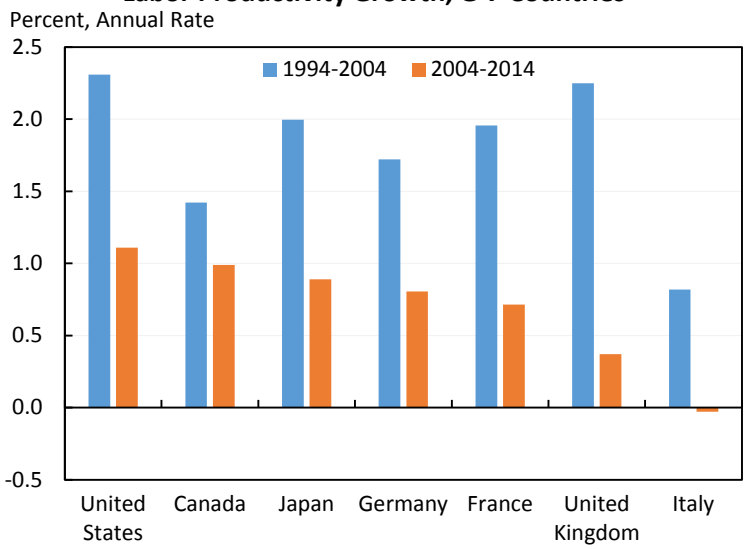

There are many reasons to believe that the official statistics fail to capture the full range of productivity improvements, so the 1.1-percent estimate likely understates U.S. productivity growth from 2004 to 2014. But so, too, does the 2.3-percent figure understate productivity growth from 1994 to 2004, a period that witnessed the *de facto* invention of the World Wide Web and its associated uses for search, ecommerce, email and much more—not to mention the widespread adoption of cellphones and invention of mobile email. Recent research has confirmed that there is little reason to doubt the magnitude of the reduction in productivity growth, including pointing out that the slowdown has also occurred in well-measured industries (Byrne, Fernald, and Reinsdorft 2016; Syverson 2016).

The slowdown in productivity growth has had profound consequences, contributing to slower growth in real wages and increasing our long-run fiscal challenges. It is outside the scope of these remarks to discuss all of the causes of the productivity slowdown, but suffice it to say that weak business investment has been part of the issue but that the pace of innovation also seems to have slowed.[1] This may seem counterintuitive given all the excitement around new innovations—including in robotics, AI, and automation more generally—but as exciting as these innovations may be, they still represent only a tiny fraction of our lives when compared to other sectors of the economy like housing, retail, education, and health.

[1] The productivity slowdown is discussed more fully in Furman (2016b) and Chapter 2 of the 2016 *Economic Report of the President*.

That said, the technology sector of our economy is making important contributions to productivity growth. A 2015 study of robots in seventeen countries found that they added an estimated 0.4 percentage point on average to those countries' annual GDP growth between 1993 and 2007, accounting for a bit more than one-tenth of those countries' overall GDP growth during that time (Graetz and Michaels 2015). Moreover, since 2010, worldwide shipments of industrial robots have increased dramatically, as shown in Figure 2, potentially signaling even more productivity growth in the future.

Figure 2

Estimated Worldwide Annual Supply of Industrial Robots, 2004-2014

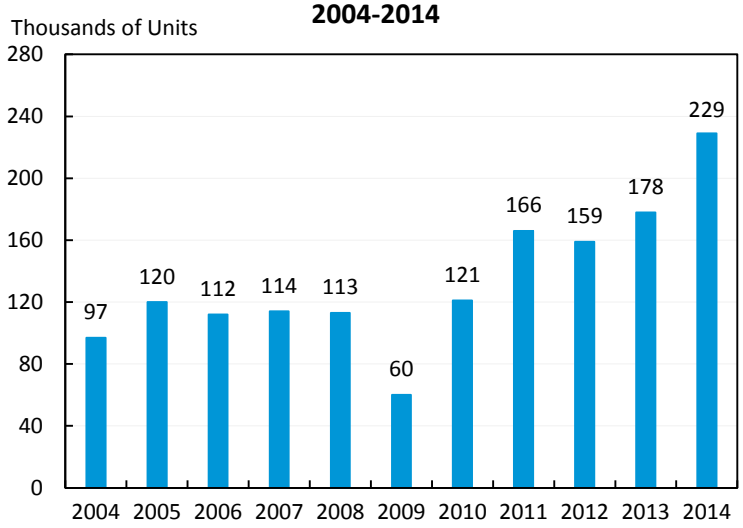

Relatedly, there has been dramatic progress in recent years in AI and its application in a diverse set of areas. For example, companies are using AI to analyze online customer transactions in order to detect and stop fraud, and, similarly, social networking sites are using it to detect when an account may have been hijacked. Thanks to AI, web search applications are now more accurate—for example, by correcting for manual entry error—thereby reducing costs associated with search. In radiology, where doctors must be able to examine radiological images for irregularities, AI's superior image processing techniques may soon be able to provide more accurate image analysis, expanding the potential for earlier detection of harmful abnormalities and reducing false positives, ultimately leading to better care.

AI is also making inroads in the public sector as well. For example, predictive analytics has great potential to improve criminal justice procedures, although it must be used responsibly to avoid bias. Charlotte-Mecklenburg, North Carolina is using predictive analytics to help inform pre-trial release decisions, something supported by the President's Data Driven Justice Initiative as a way to reduce both recidivism rates and the jail population (White House 2016a). Though some believe that AI removes bias from decision-making, we have to remember that imperfect human beings, with their own biases, write the algorithms and have collected and analyzed the data over time. As noted in the Executive Office of the President (2016) report *Big Risks, Big Opportunities: the Intersection of Big Data and Civil Rights*, it is important for us to be cognizant of and correct for the ways past bias can impact the future of AI.

However, while AI research has been underway for decades, recent advances are still very new, and, as a result, AI has not had a large macroeconomic impact, at least not yet. The most recent major progress in AI has been in deep learning, a powerful method but one that must be applied in a customized way for each application. To foreshadow a point I will make below, it is notable that the recent advances in deep learning built on research on neural nets by university labs which was largely funded by the Defense Advanced Research Projects Agency (DARPA) and other government agencies in the 1980s and 1990s. Even though we have not made as much progress recently on other areas of AI, such as logical reasoning, the advancements in deep learning techniques may ultimately act as at least a partial substitute for these other areas.

While AI has an advantage over humans in many respects, humans still maintain a substantial advantage over AI for tasks that involve social intelligence, creativity, and general intelligence. For example, AI today can do decent translations but cannot come close to what a human can do with his or her knowledge of both languages, social and cultural context, and sense of the author's argument, emotional states, and intentions. As it stands, even the most popular machine translator still fails to reach the accuracy of a human translator.

We have had substantial innovation in robotics, AI, and other areas in the last decade. But we will need a much faster pace of innovation in these areas to really move the dial on productivity growth going forward. I do not share Robert Gordon's (2016) confidently pessimistic predictions or Erik Brynjolfsson and Andrew Mcafee's (2014) confidently optimistic ones because past productivity growth has been so difficult to predict. I take some solace from the evidence that major new inventions like electricity have manifested themselves in the past in successive waves of added productivity growth, a pattern that could repeat itself in the future (Syverson 2013). More importantly, as interesting as the endless debate over the future of productivity growth is, it is considerably less interesting and important than the question of what we can do about it, a question to which I will return in the final part of this talk.

Past Innovations Have Sometimes Increased Inequality—and the Indications Suggest AI Could Be More of the Same

Advanced economies have seen vast amounts of innovation in the last three centuries. Most of the kinds of jobs that existed in the 1700s do not exist today, but jobs no one could have imagined then have taken their place. As a result, over long periods of time it has generally been the case that about 95 percent of the people in the United States who want a job at a given point in time can find one—despite massive changes in technology.

Although labor markets do not function like the stylized models for a commodities like wheat that populate economics textbooks, within broad parameters the basic operation of supply and demand is the mechanism that makes sure that just about everyone who wants a job can find one. For this to happen, however, wages need to adjust to make supply equal to demand. In recent decades, much of that adjustment in wages has been in the form of a large decline in wages for low-skilled workers relative to high-skilled workers. From 1975 until 2014, those with a high school degree watched their relative wages fall from over 80 percent of the amount earned by full-time, full-year workers with at least a college degree to less than 60 percent (CEA 2016b).

My worry is not that this time could be different when it comes to AI, but that this time could be the same as what we have experienced over the past several decades. The traditional argument that we do not need to worry about the robots taking our jobs still leaves us with the worry that the only reason we will still have our jobs is because we are willing to do them for lower wages.

One indication of the impact of automation on inequality comes from the work of Carl Frey and Michael Osbourne (2013). Frey and Osbourne's headline result is that approximately 50 percent of U.S. jobs are at risk of being replaced by automation. This estimate has launched a debate, with, for example, researchers at the OECD (Arntz and Zierahn 2016) estimating that 9 percent of jobs are at risk of being replaced by automation. But let us ignore the debate around the headline results and instead treat Frey and Osbourne as a plausible but admittedly highly uncertain guess of what occupations will come under pressure from technology. At the Council of Economic Advisers, we ranked these occupations by wages and found that, according to the Frey and Osbourne analysis, 83 percent of jobs making less than $20 per hour would come under pressure from automation, as compared to 31 percent of jobs making between $20 and $40 per hour and 4 percent of jobs making above $40 per hour (see Figure 3).

Figure 3
Probability of Automation by an Occupation's Median Hourly Wage

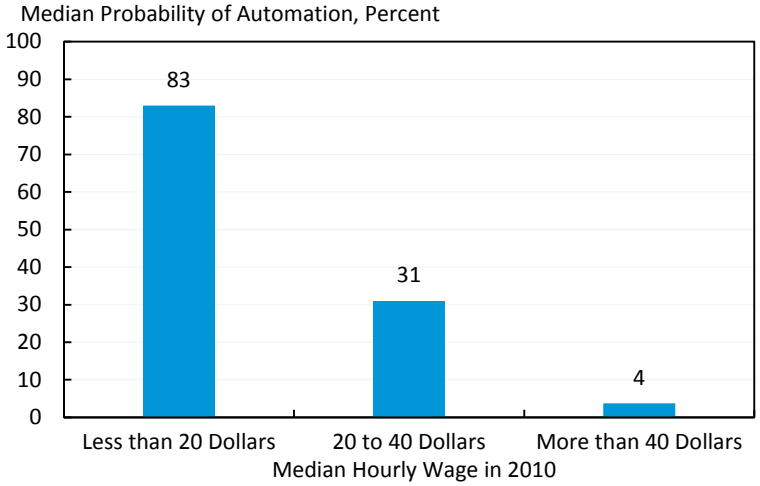

Even if the exact numbers are somewhat off, the relative magnitudes are massive. To the degree that wages and skills are correlated, this means a large decline in the demand for less-skilled jobs and little decline in the demand for higher-skilled jobs. This result points to a shift in the impact of automation on the labor market. At points in the past, automation led to a so-called polarization of the labor market because jobs requiring a moderate skill level—which historically included bookkeepers, clerks, and certain assembly-line workers—were easier to routinize, although more recently that process of polarization appears to have stopped (Autor 2014; Schmitt, Schierholz, and Mishel 2013). Conversely, higher-skill jobs that use problem-solving capabilities, intuition, and creativity, as well as lower-skill jobs that require situational adaptability and in-person interactions, were less easy to routinize. If anything, the new trends could put more pressure on earnings inequality. We are already seeing some of this play out—for

example, when we go shopping and take our groceries to a kiosk instead of a cashier, or when we call a customer service help line and interact with an automated customer service representative.

It would be wrong, however, to believe that inequality is purely a function of technology. Relative wages do depend in part on the demand for labor, which is partially a function of technology. However, they also depend on the supply of different levels of skill—in other words, the distribution of educational attainment (Goldin and Katz 2008)—and also on institutional arrangements that affect wage setting, such as collective bargaining (Western and Rosenfeld 2011).

Technology, in other words, is *not* destiny. Many countries have experienced similar technological change as the United States, yet over the last four decades the United States has seen both a greater increase in income inequality and higher overall levels of inequality than other major advanced economies, as shown in Figure 4. When it comes to inequality—and, as I will note in a moment, to the labor market more broadly—institutions and policies can help determine whether and to what extent changes in technology shape economic outcomes.

Figure 4
Share of Income Earned by Top 1 Percent, 1975-2015

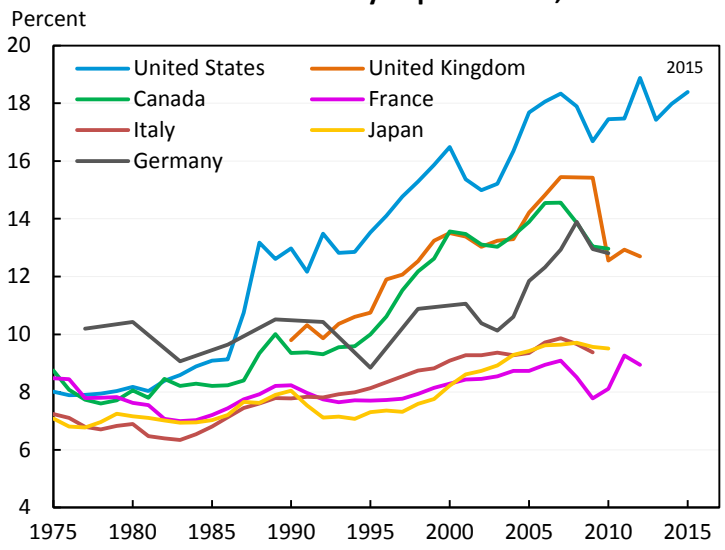

The Long-term Decline in the Labor Force Participation Rate Raises Other Concerns About the Potential Impact of AI

Moreover, the experience of the U.S. labor market over the last half century raises questions around even this (relatively) optimistic view that we can avoid large-scale job losses at the expense of greater inequality. The fact that the labor force participation rate for men between the ages of 25 and 54 has declined steadily from a high of 98 percent in the 1950s to 88 percent today raises important doubts about the complacency about full employment as a general state of the economy. As discussed in detail in a recent report by the Council of Economic Advisers (2016b), the decline in the labor force participation rate has been concentrated among men with a high school degree or less and has coincided with a decline in their relative wages. This decline

suggests that decreasing labor force participation among this group is a manifestation of reduced labor demand, resulting in both fewer employment opportunities and lower wages for less-skilled men. Technological advances, including the increasing use of automation, may partly account for this decline in demand for less-skilled labor, with globalization likely contributing as well.

(I focus on prime-age men here because I believe their experience over the past six decades to be the best historical parallel for future effects of technological change on participation in the workforce for both men *and* women. In the second half of the 20th century, prime-age women's participation rose sharply, as social and cultural changes in the decades following World War II swamped any negative effects on participation due to technological change. It is important to note, however, that prime-age women's participation has fallen in the last decade and a half, paralleling the earlier experience of prime-age men.)

The concern is not that robots will take human jobs and render humans unemployable. The traditional economic arguments against that are borne out by centuries of experience. Instead, the concern is that the process of turnover, in which workers displaced by technology find new jobs as technology gives rise to new consumer demands and thus new jobs, could lead to sustained periods of time with a large fraction of people not working. The traditional economic view is largely a statement about long-run equilibrium, not about what happens in the short-to-medium term. The fall in the labor force participation rate suggests that we must also think carefully about short-run dynamics as the economy moves towards this long-run equilibrium. In the short run, not all workers will have the training or ability to find the new jobs created by AI. Moreover, this “short run” (which is a description of where the economy is in relation to equilibrium, not a description of a definite length of time) could last for decades and, in fact, the economy could be in a series of “short runs” for even longer.

As a result, AI has the potential—just like other innovations we have seen in past decades—to contribute to further erosion in both the labor force participation rate and the employment rate. This does not mean that we will necessarily see a dramatically large share of jobs replaced by robots, but even continuing on the past trend of a nearly 0.2-percentage-point annual decline in the labor force participation rate for prime-age men would pose substantial problems for millions of people and for the economy as a whole.

As in the case of inequality, however, we should not interpret this as technological determinism. While most other advanced economies have seen declines in prime-age male labor force participation, the decline in the United States has been steeper than in almost every other advanced economy, as shown in Figure 5. Part of the reason may be that U.S. labor market institutions are less supportive of participation in the workforce than other countries’ (CEA 2016b).

Figure 5
Prime-Age Male Labor Force Participation Rates Across the OECD

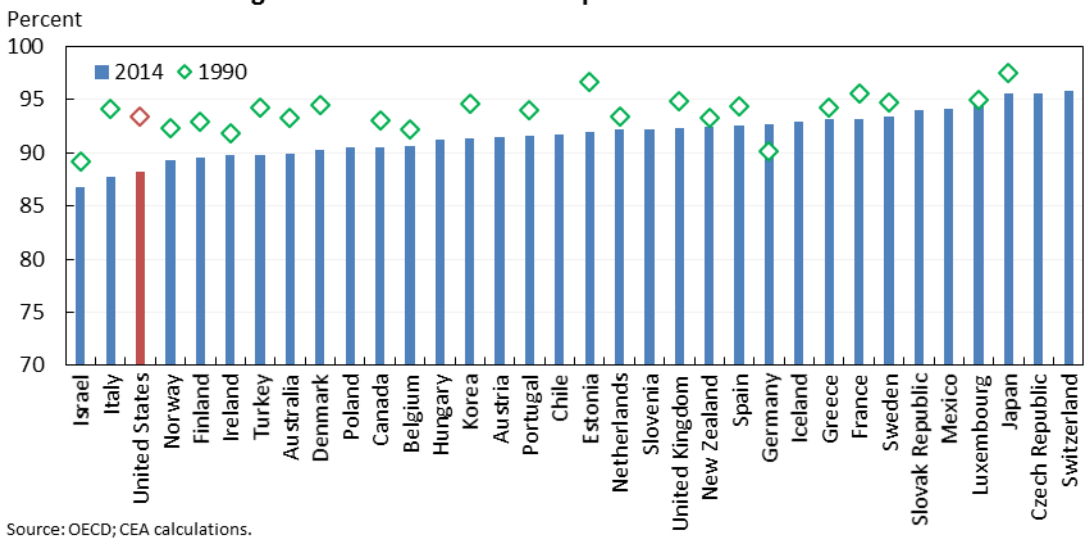

Source: OECD; CEA calculations.

There is no reason the economy cannot generate substantial levels of employment at much higher levels of technology and productivity than we have today. What matters, however, is how our labor market institutions cope with these changes, help support the creation of new jobs, and successfully match workers to them. Some of the policies along these lines proposed by the President were discussed extensively in the same recent CEA report, and include expanding aggregate demand, increasing connective tissue in labor markets, reforming taxes to encourage work, and creating more flexibility for workers (CEA 2016b). Other policy responses include expanding education and training so more people have skills that complement and benefit from innovations, increasing the progressivity of the tax system to make sure that everyone shares in the overall benefits of the economy, and expanding institutional support for higher wages, including a higher minimum wage and stronger collective bargaining and other forms of worker voice (Furman 2016a).

Replacing the Current Safety Net with a Universal Basic Income Could Be Counterproductive

Fears of mass job displacement as a result of automation and AI, among other motivations, have led some to propose deep changes to the structure of government assistance. One of the more common proposals has been to replace some or all of the current social safety net with a universal basic income (UBI): providing a regular, unconditional cash grant to every man, woman, and child in the United States, instead of, say, Temporary Assistance to Needy Families (TANF), the Supplemental Nutrition Assistance Program (SNAP), or Medicaid.

While the exact contours of various UBI proposals differ, the idea has been put forward from the right by Charles Murray (2006), the left by Andy Stern and Lee Kravitz (2016), and has been a staple of some technologists' policy vision for the future (Rhodes, Krisiloff, and Altman 2016). The different proposals have different motivations, including real and perceived deficiencies in the current social safety net, the belief in a simpler and more efficient system, and also the

premise that we need to change our policies to deal with the changes that will be unleashed by AI and automation more broadly.

The issue is not that automation will render the vast majority of the population unemployable. Instead, it is that workers will either lack the skills or the ability to successfully match with the good, high paying jobs created by automation. While a market economy will do much of the work to match workers with new job opportunities, it does not always do so successfully, as we have seen in the past half-century. We should not advance a policy that is premised on giving up on the possibility of workers' remaining employed. Instead. our goal should be first and foremost to foster the skills, training, job search assistance, and other labor market institutions to make sure people can get into jobs, which would much more directly address the employment issues raised by AI than would UBI.

Even with these changes, however, new technologies can increase inequality and potentially even poverty through changes in the distribution of wages. Nevertheless, replacing our current antipoverty programs with UBI would in any realistic design make the distribution of income worse, not better. Our tax and transfer system is largely targeted towards those in the lower half of the income distribution, which means that it works to reduce both poverty and income inequality. Replacing part or all of that system with a universal cash grant, which would go to all Americans *regardless of income*, would mean that relatively less of the system was targeted towards those at the bottom—increasing, not decreasing, income inequality. Unless one was willing to take in a much larger share of the economy in tax revenues than at present, it would be difficult both to provide a common amount to all individuals and to make sure that amount was sufficient to cover the needs of the poorest households. And for any additional investments in the safety net that one would want to make—and the President has proposed numerous such investments—one must confront the same targeting question.

Finally, some of the motivation for UBI has nothing to do with future technological developments. Instead, some UBI proponents have put forward the argument that it would be simpler, fairer and less distortionary than the social assistance system we have today. This is not the space to go into great detail on this, but suffice it to say that today's system could be improved, and the President has proposed a number of improvements to social assistance programs (OMB 2016). But at the same time, a wave of recent research has found that many of the common criticism of these programs—for example, that they discourage work, or that they do little to reduce poverty—have been greatly overstated, and a number of programs—including nutritional assistance, Medicaid, and the Earned Income Tax Credit (EITC)—have important benefits for the long-run earnings, health and educational attainment of children who grow up in recipient households.[2]

This is not to say that we should not make the tax-and-transfer system more progressive—just that we need to match our ambitions to the revenue available and to build on what is already successful in our social safety net.

[2] See CEA (2014) for a survey of research on anti-poverty programs and Furman and Ruffini (2015) for a summary of research the long-term benefits of anti-poverty programs.

What Government Can Do to Advance AI and to Ensure that All Americans Share its Benefits

With or without AI, we would have a lot to do to address high levels of inequality and the falling labor force participation rate—which is the motivation for several of the President's proposals that I referenced earlier. To the degree that we are optimistic about AI, that should increase our motivation to undertake these changes. But there is little basis for believing that AI should dramatically change the overall direction or goals of our current policies.

Despite the labor market challenges we may need to navigate, my bigger worry is that we will not invest enough in AI. And public policy can play a role here. To be sure, the private sector will be the main engine of progress on AI. In 2015 the private sector invested $2.4 billion on AI, as compared to the approximately $200 million invested by the National Science Foundation (NSF).[3] The government's role should include policies that support research, foster the AI workforce, promote competition, safeguard consumer privacy, and enhance cybersecurity. As I will point out in turn, these policies will not only help consumers, but would also benefit firms and ultimately lead to stronger economic growth.

Basic Research

In 2015, American businesses devoted almost 1.8 percent of GDP to research and development, the highest share on record. But government investments in R&D have fallen steadily as a share of the economy since the 1960s. While business investment is critical, it is not sufficient. Basic research discoveries often have great social value because of their broad applicability, but there tends to be underinvestment in basic research by private firms because it is difficult for a private firm to appropriate the gains from such research. In fact, while the private sector accounts for roughly two-thirds of all spending on R&D, it is important to keep in mind that it largely invests in applied research while the Federal government provides 60 percent of the funding for basic research.

As a result, in the absence of public investment, aggregate R&D investment (not only basic research but also applied research and experimental development) is bound to fall short of what is socially optimal (Nelson 1959). In fact, recent analysis suggests that the socially optimal level of R&D investment—the amount that would produce the greatest rate of economic growth—is two to four times greater than actual spending (Jones and Williams 1998; Bloom, Schankerman, and Van Reenen 2013; Akcigit, Hanley, and Serrano-Velarde 2013). This gap is particularly large for basic research, since its role as the "seed corn" of future innovations means that it generates the largest spillovers.

In the past, government-funded research has been the catalyst for many of the AI technologies that we know today. For example, starting in 2004 DARPA has held several long-distance,

[3] For private funding see https://www.cbinsights.com/blog/artificial-intelligence-funding-trends/#funding. For public funding see http://www.nsf.gov/about/budget/fy2017/pdf/18_fy2017.pdf. According to the NSF, in 2015 there was $194.58 million in funding for the NSF Directorate for Computer and Information Science and Engineering's Division of Information and Intelligent Systems (IIS), much of which is invested in research on AI. These figures do not include investment by other agencies, including Department of Defense.

driverless car competitions. These Grand Challenges have awarded cash prizes for innovations for autonomous cars and have been widely credited with dramatically accelerating their progress. More recently, the government has been investing in public-private partnerships such as the Smart Manufacturing Innovation Institute that is working to build U.S. leadership in smart collaborative robots (White House 2016b).

Looking forward, we will need to ensure that government continues to provide this same sort of support to emerging AI technologies. The President's Fiscal Year 2017 budget request includes a 4-percent increase in overall R&D funding from 2016, with proposals to increase both discretionary and mandatory funding. In addition, the President has proposed to expand and simplify the Research and Experimentation tax credit, an effective incentive for businesses to increase their often more applied research.

A Workforce to Develop AI

It may be computers that ultimately beat the world champions won at *Jeopardy*, chess, and *go*—but these computers were built and programmed by teams of humans. And human learning and skills will continued to be critical if we are to have more AI. This, in turn, depends on our policy choices. One of those is education. Government support for basic research does not just contribute to research, it also makes universities a more attractive place to conduct that research—helping to stem some of the drain from universities to the private sector by researchers looking not just for higher salaries but also for greater ability to undertake research. This, in turn, not only helps produce more research but also enables universities to continue training students to make advances in areas like deep learning.

Another way high-skill workers can enter the labor market is through immigration, and recent evidence shows that the contribution of skilled migration to innovation has been substantial. For example, Peri, Shih, and Sparber (2014) find that inflows of foreign science, technology, engineering, and mathematics (STEM) workers explain between 30 and 50 percent of the aggregate productivity growth that took place in the United States between 1990 and 2010. This suggests that we should increase the number of visas—which is currently capped by legislation—to allow more high-skilled workers to come into the country.

Finally, the policy steps that I referenced when discussing inequality and labor force participation would not just help address some of the potentially problematic side effects of AI—they could also contribute to AI's progress itself. Advancing AI will require as many brilliant ideas and innovations as possible, and when a fraction of the population is disadvantaged from contributing based on their income or perceptions about their gender or race and ethnicity—issues that can be especially important in software—it reduces both the pool of available talent and the pace of innovation. Moreover, a workforce that is better trained to take advantage of AI in their jobs will also increase the demand for the development of AI—and accelerate its progress. Some of the Administration's recent programs will help. For example, the bipartisan Workforce Innovation and Opportunity Act, which President Obama signed into law in July 2014, consolidates existing funding initiatives, helps retrain workers in skills for which employers are looking, and matches those workers to employers. In addition, in March 2015, the Administration launched the TechHire initiative, which aims to equip 17-to-29 year-olds with

skills necessary for jobs in information technology fields, including software development, network administration, and cybersecurity.

Competition

Competition from new and existing firms has always played an important role in the creation and adoption of new technologies and innovations, and this is no different in the case of AI. Startups are a critical pathway for the commercialization of innovative new ideas and products. Startups, or the possibility of entry by a startup, also incentivize established firms to innovate and reduce costs. More than 50 years ago, the Nobel Prize-winning economist Kenneth Arrow (1962) argued that a monopolist may have relatively weak incentives to innovate, since its innovations do not allow it to "steal" business from competitors. On the other hand, competition pushes firms to invest in new technologies that help to lower costs, and also to invest in innovations that can lead to improvements in the quality of existing products.

The rapid evolution of technology can pose challenges for developing sound pro-competition policies, both in terms of defining the scope of the market and assessing the degree of contestability or the possibilities for disruption. For example, while it is probably too early to assess the role of AI in competition policy, one might imagine that when a large incumbent has access to most of the customer data in the market, it is able to use AI to refine its products better than any potential entrant could hope, and can thereby effectively foreclose entry. Treating consumer data as a critical resource may therefore be an effective remedy, though making consumer data "open" also requires that robust privacy protections are in place. But we should not overstate these challenges, as they have been present in traditional markets. Moreover, some past "disruptions" were themselves at least in part enabled by government policies aimed at promoting competition.

As a result, antitrust authorities such as the Department of Justice and Federal Trade Commission (FTC) continue to play a vital role in enforcement of antitrust policy. Moreover, U.S. antitrust agencies (including sector-specific regulators, such as the Federal Communications Commission and the Department of Transportation), have recently pursued creative approaches to merger review, whereby mergers have been able to proceed conditional on ensuring equitable access to critical resources to compete in given markets. These actions—which span the wireless mobile service, wireline broadband, technology manufacturing, and aviation industries, to name a few—represent a relatively novel and effective supplement to takeover reviews.

Privacy

As society becomes more and more reliant on technology, and on AI in particular, consumers have begun to share more and more data, and companies are collecting more data about users' activities both on- and offline. While consumers may be willing to part with personal information as part of the price of engaging in a digital transaction, it is not clear that they always have the relevant information about all of the associated tradeoffs involved in a digital transaction. Even if consumers understand the various tradeoffs involved when personal information is shared in order to access a tailored product or service at a point in time, it could be unclear how that data will subsequently be used by the firm or other firms. This is made even

more difficult because, thanks to AI, data is being created, analyzed, and used in new ways not previously imagined by consumers, firms and governments.

As highlighted in a 2014 FTC report, there is generally a lack of transparency about what firms do with their customers' data, and about how much control customers may or may not have over this data.[4] Thus, one role that governments may increasingly have to consider embracing is in setting transparency rules surrounding firms' collection and use of consumer data. If structured correctly, these rules could allow consumers to stand a better chance of making informed choices about surrendering their data beyond the user agreements that appear to be the industry standard today. Transparency may not be enough in some cases, and the government may also need to consider whether there are uses of the data that should be restricted or prohibited.

The rapidly evolving landscape surrounding privacy of firms' use of individuals' data highlights the importance of the President's robust privacy agenda. The Administration has been very active on consumer privacy issues over the past seven years, working in a range of ways, from multi-stakeholder efforts to the Consumer Privacy Bill of Rights, the 90 Day Big Data and Privacy Review, and 2016's *Big Data: A Report on Algorithmic Systems, Opportunity, and Civil Rights*.

C*ybersecurity*

Finally, some of the biggest risks facing consumers, firms, and governments alike are identity theft and data breaches. AI has the potential to help detect fraud and combat cyber intrusions, but it also could be used for malign purposes—and its widespread adoption could raise the severity of the cyberattacks that do end up succeeding.

Enhanced information security can help reduce identity theft, but the tradeoff between security and efficiency may result in fewer security protections being put in place then are expected or are necessary. Security is costly, and at some point the benefits from additional security outweigh the costs. A firm's decision about the point at which the benefits outweigh the costs could very well be different than the social optimum. One reason for this discrepancy is that an individual firm does not bear all the costs of a security breach, as many of these costs are distributed across an entire network of individuals. Another reason for this discrepancy is asymmetric information, whereby it is hard for the buyer to evaluate the security of the products sold by the seller.

To help address these challenges, the Administration has devised an approach that leverages the cybersecurity capabilities of the Government, businesses, and individuals. Earlier this year, the President announced his Cybersecurity National Action Plan, which brings together over seven years' worth of efforts across the public and private sector on this complex issue. As part of this plan, the President established a bipartisan Commission on Enhancing National Cybersecurity, comprised of top thinkers in this area from outside the government, and a public outreach campaign in partnership with the private sector to empower Americans to take additional steps necessary to secure their online accounts and data. Critical to this effort is $19 billion in cybersecurity funding requested by the President, a more than 35-percent increase over current

[4] For example, in many cases, consumers are not allowed to update or even correct inaccurate data (FTC 2014, p. vii).

funding levels. In addition, the President's BuySecure Initiative provides consumers with more tools to secure their financial future by assisting victims of identity theft, improves the Government's payment security as a customer and a provider, and accelerates the transition to stronger security technologies and the development of next-generation payment security tools. Finally, DARPA has established a Cyber Grand Challenge that seeks to automate cyber defense, fielding the first generation of machines that can discover, prove and fix software flaws in real-time, without any assistance. If successful, the speed of autonomy could someday blunt the structural advantages of cyber attackers.

Conclusion

AI is one of many areas of innovation in the U.S. economy right now. At least to date, AI has not had a large impact on the aggregate performance of the macroeconomy or the labor market. But it will likely become more important in the years to come, bringing substantial opportunities—and our first impulse should be to embrace it fully. The biggest worry I have about AI is that we will not have enough of it, and that we need to do more to make sure we can continue to make groundbreaking discoveries that will raise productivity growth, improving the lives of Americans and people throughout the world. However, it is also undeniable that like technological innovations in the past, AI will bring challenges in areas like inequality and employment. As I have tried to make clear throughout my remarks today, I do not believe that exogenous technological developments solely determine the future of growth, inequality, or employment. Public policy—including public policies to help workers displaced by technology find new and better jobs and a safety net that is responsive to need and ensures opportunity —has a role to play in ensuring that we are able to fully reap the benefits of AI while also minimizing its potentially disruptive effects on the economy and society. And in the process, such policies could also contribute to increased productivity growth—including advances in AI itself.

Notes to Figures

Figure 1
Source: Conference Board, Total Economy Database; CEA calculations.

Figure 2
Source: International Federation for Robotics, World Robotics 2015.

Figure 3
Source: Bureau of Labor Statistics, Occupational Employment Statistics; Frey and Osborne (2013); CEA calculations.

Figure 4
Source: World Wealth and Income Database.

Figure 5
Source: Organisation for Economic Co-operation and Development.

References

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. 2014. "Polanyi's Paradox and the Shape of Employment Growth." NBER Working Paper No. 20485.

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 & Company

Byrne, David, John Fernald, and Marshall Reinsdorf. 2016. "Does the United States have a Productivity Slowdown or a Measurement Problem." FEDS Working Paper No. 2016-17.

Council of Economic Advisers (CEA). 2014. "Chapter 6: The War on Poverty 50 Years Later: A Progress Report." *Economic Report of the President.*

_____. 2016a. *Economic Report of the President.*

_____. 2016b. "The Long-Term Decline in Prime-Age Male Labor Force Participation." Report.

Executive Office of the President (EOP). 2016. "Big Data: A Report on Algorithmic Systems, Opportunity, and Civil Rights." Report.

Frey, Carl and Michael Osborne. 2013. "The Future of Employment: How Susceptible are Jobs to Computerization." Oxford University.

Federal Trade Commission (FTC). 2014. "Data Brokers: A Call for Transparency and Accountability." May.

Furman, Jason. 2016a. "The Truth About American Unemployment." *Foreign Affairs* 95 (4): 127-138.

_____. 2016b. "The United States and Europe: Short-Run Divergence and Long-Run Challenges." Remarks at Bruegel.

Furman, Jason and Krista Ruffini. 2015. "Six Examples of the Long-Term Benefits of Anti-Poverty Programs." Council of Economic Advisers blog.

Goldin, Claudia and Lawrence Katz. 2008. *The Race Between Education and Technology.* Cambridge, MA: Belknap of Harvard UP.

Gordon, Robert. 2016. *The Rise and Fall of American Growth: The U.S. Standard of Living Since the Civil War*. Princeton, NJ: Princeton University Press.

Graetz, Georg and Guy Michaels. 2015. "Robots at Work." Centre for Economic Policy Research (CEPR) Discussion Paper No. DP10477.

Murray, Charles. 2006. *In Our Hands: A Plan to Replace the Welfare State*. Washington: AEI Press.

Office of Management and Budget. 2016. "Advancing Economic Opportunity and Mobility: Meeting Our Challenges: The President's Fiscal Year 2017 Budget." Fact sheet.

Peri, Giovanni, Kevin Shih, and Chad Sparber. 2014. "Foreign Stem Workers and Native Wages and Employment in U.S. Cities." NBER Working Paper No. 20093.

Rhodes, Elizabeth, Matt Krisiloff, and Sam Altman. "Moving Forward on Basic Income." Y Combinator blog.

Schmitt, John, Heidi Schierholz, and Lawrence Mishel. 2013. "Don't Blame the Robots. Assessing the Job Polarization Explanation of Growing Wage Inequality." Economic Policy Institute.

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: PublicAffairs.

Syverson, Chad. 2013. "Will History Repeat Itself? Comments on 'Is the Information Technology Revolution Over?'" *International Productivity Monitor* 25 (2): 37-40.

_____. 2016. "Challenges to Mismeasurement Explanations for the U.S. Productivity Slowdown." NBER Working Paper No. 21974.

Western, Bruce and Jake Rosenfeld. 2011. "Unions, Norms, and the Rise in U.S. Wage Inequality." *American Sociological Review* 76 (4): 513-537.

White House. 2016a. "FACT SHEET: Launching the Data-Driven Justice Initiative: Disrupting the Cycle of Incarceration." Press release.

_____. 2016b. "FACT SHEET: President Obama Announces Winner of New Smart Manufacturing Innovation Institute and New Manufacturing Hub Competitions." Press release.