Chapter 7

Adapting to Technological Change with Artificial Intelligence while Mitigating Cyber Threats

Although technological change has always had significant effects on economic activity, artificial intelligence (AI) and high-speed automation are among its most important recent manifestations. The expansion of computing power and availability of big data have fueled remarkable advances in computer science, enabling technology to perform tasks that traditionally required humans and significant amounts of time. However, along with these advances' prospects for encouraging continued productivity growth, they also threaten to significantly disrupt the labor market, particularly among people whose work involves routine and manual tasks. Astute policymaking will play an integral role in leveraging technology as an asset for the country, while mitigating potential disruptions.

The first section of this chapter briefly defines AI and corresponding advances in computer science. AI's most distinctive feature is that it can be used to manage a wide range of highly complex tasks with little required supervision, relative to conventional technology. This general applicability broadens the types of tasks where AI could plausibly be a substitute for human labor, underscoring both the economic promise of AI and its potential risks.

The second section places AI within the broader historical context of technological change and highlights the CEA's predictions for its short-, medium-, and long-run effects on productivity and wages. Although we may experience a span of years where AI substitutes for human-based labor for many tasks, AI,

like much technological change, will ultimately benefit labor through greater productivity and real wage growth.

The third section explores AI's heterogeneous effects and automation across industries and the skill distribution. Using autonomous vehicles as a case study, we show that one of the key factors for understanding the impact of technological change on employment is the price elasticity of demand. AI is expected to have a positive net effect on industrial employment, though there could be subsector-specific price declines based on changing consumer demand.

The fourth section pivots to the possible risks of technological advances. Building on findings in the 2018 *Economic Report of the President* on the cost of cybersecurity breaches, we analyze how measurement problems related to these breaches make it difficult to estimate their costs. We present new data from 2018 on the pervasiveness of cybersecurity vulnerabilities and the paucity of firms' responses to them across *Fortune* 500 companies.

The fifth and final section highlights the role of policy and the considerable strides that have been taken by the Trump Administration during the past two years. The Administration will continue to embrace technological change, while maximizing its promise and minimizing its risk.

Recent years have seen enormous advances in computer science, leading to skyrocketing hardware and software capabilities. The refinement of computers continues to advance at a rapid rate. The computational power that took up enormous refrigerated rooms a few decades ago has been miniaturized to a fraction of its former size. Moreover, computer scientists and engineers have made remarkable discoveries in artificial intelligence (AI) and automation. These advances have complemented years of rapid growth in computer processing power, along with the explosive growth in the availability of digitized data. According to two prominent scholars, "the key building blocks are already in place for digital technologies to be as important and transformational to society and the economy as the steam engine" (Brynjolfsson and McAfee 2014, 9).

In last year's *Report*, we highlighted one aspect of the rapid diffusion of computer technology: the increasing exposure of the economy to malicious cyber activity—for example, cybercrime. We found that cybercrime had

expanded so much that in 2016 alone that it caused up to $109 billion in harm to the economy. Yet computers have, of course, created many more benefits than costs, and their rapid evolution promises to fundamentally transform the economy in the decade ahead. In 2016, President Obama's Council of Economic Advisers published a sweeping report outlining the likely economic impact and policy challenges of accelerating technological change. One metric of how rapidly the sector is advancing is that already, in 2018 and 2019, enough change has occurred so an update of the previous reports is essential for meeting the challenges of the next decade and beyond. We look ahead in wonder at the possibilities of advanced thinking machines, but also worry that automation will proceed at such a rapid pace that many workers in today's economy will suddenly find themselves superfluous or disconnected from competitive job opportunities. We also consider the additional cybersecurity risks posed by the increased reliance on information technology.

In this chapter, we dig deeper than we did a year ago into the promise and risks of the ongoing computer science revolution. We begin by reviewing the latest developments in AI and automation, discussing their likely economic effects. The central theme of the first section of this chapter is that a narrow, static focus on possible job losses paints a misleading picture of AI's likely effects on the Nation's economic well-being. With technological advances, specific types of legacy positions are usually eliminated, though new jobs and evolving work roles are created—increasing real wages, national income, and prosperity over time. Automation can complement labor, adding to its value; and even when it substitutes for labor in certain areas, it can lead to higher employment in other types of work and raise overall economic welfare. This will likely be what happens as AI transforms more and more aspects of the economy, though new challenges will arise about cybersecurity. In the years to come, AI appears poised to automate tasks that had long been assumed to be out of reach. Thus, we also analyze the important role of reskilling, apprenticeship initiatives, and future hiring processes to help mitigate the potentially disruptive employment effects of technological change and automation throughout the skill distribution.

One key question for economists today is whether—in addition to improving traditional productive processes—AI will alter processes whereby creative new ideas are generated and implemented. In other words, is AI simply the next phase in automation, or is it a real break from the past with unique implications? We explore both possibilities, but conclude that AI is likely to have major effects on the value of different skill sets and the rate at which they appreciate and depreciate. In particular, in the long run, aggregate wages will be higher because of these new advances.

We then turn to an update of our previous research on the economic vulnerabilities associated with the diffusion of technology and mobile computing capabilities into virtually every corner of our lives. Technology is leading to

new and constantly evolving complex security challenges because individuals, firms, and governments are already reliant on interconnected and interdependent technology. Whereas past conflicts unfolded on land, sea, and air, future conflicts and criminal activity will increasingly take place in cyberspace. Drawing on new data, we document that cyber vulnerabilities are quite prevalent—even in *Fortune* 500 companies with significant resources at their disposal. Although these new data do not allow us to update our 2018 estimate of the economic costs of malicious cyber activity, the latest data suggest that our previous estimate might have been too low, given the underreporting of cybercrime. We conclude by discussing the initiatives that are being implemented by the Trump Administration and the policy challenges that lawmakers will likely face in the years ahead.

What Is Artificial Intelligence?

Although there is no universal definition of artificial intelligence (AI),[1] the Future of Artificial Intelligence Act of 2017 (H.R. 4625), for example, defines AI as "any artificial system that performs tasks under varying and unpredictable circumstances, without significant human oversight, or that can learn from their experience and improve their performance. . . . They may solve tasks requiring human-like perception, cognition, planning, learning, communication, or physical action."[2] These intelligent systems generally use machine learning to form predictions and adaptively make adjustments based on new information in their environment (Russell and Norvig 2010). Because AI has such a wide array of applications across sectors and disciplines, it is viewed as a general purpose technology and important source of economic growth (Agrawal, Gans, and Goldfarb 2018). Automation technologies usually focus on automating a specific process, or multiple commonly understood processes, to reduce labor intensity, which differs greatly from highly complex, human-like decision logic, which has already been observed in the emerging embodiments of AI.

Although the general concepts and algorithms within AI are decades old, AI has emerged as an especially powerful and widely applied tool for

[1] A recent study by Deloitte (2017) contains survey results that point out ambiguity in how many top executives and everyday citizens define AI.

[2] Similarly, in the National Defense Authorization Act for Fiscal Year 2019, "the term 'artificial intelligence' includes the following: (1) Any artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight, or that can learn from experience and improve performance when exposed to data sets. (2) An artificial system developed in computer software, physical hardware, or other context that solves tasks requiring human-like perception, cognition, planning, learning, communication, or physical action. (3) An artificial system designed to think or act like a human, including cognitive architectures and neural networks. (4) A set of techniques, including machine learning, that is designed to approximate a cognitive task. (5) An artificial system designed to act rationally, including an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communicating, decision making, and acting."

Figure 7-1. Error Rate of Image Classification by Artificial Intelligence and Humans, 2010–17

Error rate (percent)

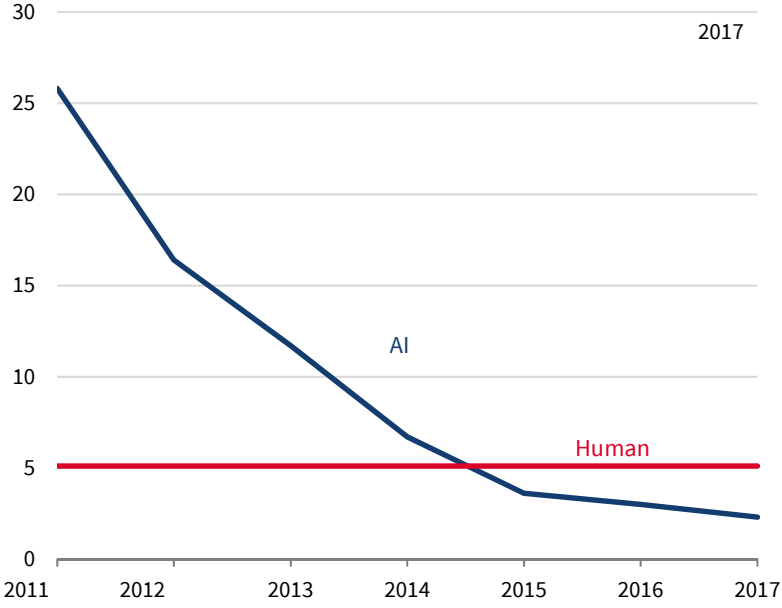

Sources: Russakovsky et al. (2015); CEA calculations.

performing not only existing tasks much more efficiently but also new tasks that were traditionally viewed as infeasible. To give just one example, researchers have created AI algorithms capable of classifying images even more reliably than humans can do under certain conditions, and at a much faster rate and scale than ever before (figure 7-1)—although these algorithms can still be tricked by savvy programmers (CSAIL 2017). More examples abound in other areas, ranging from natural language processing to theorem proving (Artificial Intelligence Index 2017). Other types of computer science and AI advances include solutions to automate high-skill human cognitive tasks, such as automated reasoning and intelligent decision support systems (Arai et al. 2014; Davenport and England 2015; Kerber, Lange, and Rowat 2016; Mulligan, Davenport, and England 2018).

The convergence of two factors have made these remarkable advances possible. First, accumulated decades of sustained growth in technology have led to an explosion in computing power. As Gordon Moore (1965) first observed, computing power historically doubles every 18 months. These advances have led to an increase in transistor density, which, combined with the declining cost of manufacturing integrated circuits, have led to a staggering increase in

computing power (Brynjolfsson and McAfee 2014).[3] Moreover, lower manufacturing costs for hardware have been complemented by annual price declines in cloud computing of 17 percent between 2009 and 2016 (Byrne, Corrado, and Sichel 2018).

Second, the colossal increase in data availability has complemented the surge in computing power, allowing researchers to develop and test AI algorithms on much larger data sets.[4] The emergence of big data has been driven by "digitization," which means the ability to take different types of information and media, ranging from text to video, and convert them into streams of ones and zeros—"the native language of computers and their kin" (Brynjolfsson and McAfee 2014, 37). Researchers have also found creative ways to convert different types of digital media into comprehensive sets of numeric quantities, which often involve "feature engineering," or optimizing the permutations of data inputs, to produce reliable predictions (Arel, Rose, and Karnowski 2010).

Machine Learning

Machine learning (ML) is integral to the design and implementation of AI (Russel and Norvig 2010). Unlike computers, which tend to execute a set of prespecified rules, AI is defined by the ability to learn and adapt to its environment.[5] There are three main types of ML algorithms—supervised, unsupervised, and reinforcement learning—which we summarize in the next paragraphs (Hastie, Tibshirani, and Friedman 2009).

First, supervised learning algorithms take a set of descriptive variables that are matched with a corresponding label ("outcome variable") and "learns" the relationship between the two. For example, to predict college attainment, a researcher could use data on whether the individual has a college degree, together with a set of individual characteristics, such as parental education and gender, to estimate classification models. Supervised learning algorithms take a subset of the sample and search for the parameters that best fit the data based on a prespecified objective function.

Second, unsupervised learning algorithms, in contrast to supervised ones, take a set of feature variables as inputs and detect patterns in the data. Though these algorithms have not been as prolifically applied as supervised

[3] An integral part of the efficiency gains among producers of computer equipment is the rapid decline in effective prices of semiconductors due to advances in chip technology (Triplett 1996). These empirical patterns have also continued during the past decade. For example, Byrne, Oliner, and Sichel (2017) find that semiconductor prices fell by 42 percent, relative to the meager 6 percent decline in the producer price index between 2004 and 2009.

[4] Computer scientists often refer to the process of developing and testing AI algorithms as "training." The process refers to estimating model parameters on a subsample, subsequently using the estimated parameters to predict out-of-sample. The quality of the out-of-sample prediction is used to, sometimes iteratively, tune model parameters.

[5] Russell and Norvig (2010, 43) remark that algorithms in deterministic settings are not a form of AI because they are executing a set of preprogrammed tasks.

learning algorithms, they are often used to simplify otherwise computationally demanding problems by reducing the number of variables that need to be kept track of, sometimes referred to as "dimensionality reduction" (Bonhomme, Lamadon, and Manresa 2017).

Third, reinforcement learning algorithms have been among the most influential class of algorithms in the emerging set of AI and big data applications. Unlike supervised and unsupervised algorithms, reinforcement learning algorithms do not require complete representation of input/output pairs, but rather only require an objective function. This function specifies how the intelligent system responds to its environment under arbitrary degrees of stochasticity (i.e., the extent to which it involves a random variable). Consider the game of chess, which contains millions of potential moves. Though individuals face cognitive limitations that preclude internal simulation of thousands, and potentially millions, of scenarios simultaneously, "deep learning" reinforcement learning algorithms have largely overcome these limitations. For example, Google's new AI algorithm, AlphaZero, defeated the world's best chess engine, Stockfish. Unlike Deep Blue—the IBM supercomputer that defeated Garry Kasparov, the world's leading chess champion in 1997—AlphaZero trained itself to play like a human, but at an unprecedented scale and aptitude (Gibbs 2017).

One way a reader can picture this evolution of computing power is by considering the computer modeling of sports outcomes. It is now common for commentators at sporting events to announce midgame the probability, given the current score, that the team that is currently ahead in the score will indeed win the game. At one point during the 2017 American football championship game Super Bowl LI, the New England Patriots had a mere 0.3 percent chance of victory (ESPN Analytics 2018). This probability was calculated based on data from previous games and an analysis of the percentage of times that a team went on to win after trailing by a certain margin deep into the third quarter. Algorithms used by various networks and media platforms allow for these odds to be constructed from historical performance data of past teams that have been in similar situations.

Moreover, as with other games, like chess, estimating probabilities of winning can grow in complexity because of real-time interactions between the players, as well as the astronomical number of possible outcomes that can be reached, even without repetitive actions between the start and end of a game. In a game with finite outcomes, given an enormously powerful computer and a set of initial conditions describing the configuration of pieces on the board, a program could explore all possible moves and responses from that state and "solve" the game. The optimal computer would then, for a given player, recommend a move from that initial state associated with the highest probability of victory for that player.

However, because there are infinite possible future states associated with almost every state of the world in a chess match, software must discover the types of moves that tend to lead to victory because exploring all future paths and developing a discrete solution is impossible for a problem with infinite outcomes. A computer equipped with AI, however, allows for a combination of human rationality with computing probabilities of victory. This provides improved predictions that can lead the AI algorithm to "play" the game, rather than attempting to solve it.

Applications of AI Technology

Today, facial recognition is possible because data (e.g., images) can be not only digitized but also collected and analyzed at scale. Suppose our AI machine, in addition to assessing the remaining possible outcomes, could also discern the identities of the players themselves and use this information to further revise its predictions based on knowledge about the two players. For instance, the probabilities of victory associated with an advantageous position would need to be updated if player 1 was an amateur and player 2 was a professional. However, if player 1's position was so advantageous that the odds of victory were 99.7 percent, even someone as talented as the professional could lose if forced to start from a severely disadvantaged position. In addition to assessing situations from a static perspective, an AI algorithm that can discern the identity of the player through facial recognition can choose strategies that are tailored to the player's weaknesses.

Another example of how AI can complement society and human tasks is through its effects on the delivery and production of educational services. One of the primary types of AI educational applications are personalized learning algorithms that allow instructors to tailor information to the unique ways that individuals learn. For example, Georgia State University sends customized text messages to students during the college enrollment process, which Page and Gehlbach (2017) find is associated with a 3.3-percentage-point increase in the probability that individuals will enroll on time.

Similarly, Arizona State University uses adaptive and hybrid learning platforms that enable teachers to offer more targeted learning experiences (Bailey et al. 2018). These platforms provide instructors with real-time intelligence to assess how well their students understand each concept, allowing instructors to pivot, when needed, to improve the learning experience. In sum, economists find significant returns on student outcomes from these "edtech" programs (Escueta et al. 2017). Given that at least 54 percent of all employees will require significant reskilling and/or upskilling by 2022 (World Economic Forum 2018), educational institutions will need to become increasingly adaptive, finding ways to integrate technology to simultaneously reduce costs, improve quality, and raise agility.

AI systems have mastered tasks that have traditionally been performed by humans. One way of measuring the breadth of these AI-based applications is to examine the clusters of emerging research content. Using the universe of Scopus and Elsevier articles, Elsevier (2018, 34) identified seven clusters of AI capabilities, including "machine learning and probabilistic reasoning, neural networks, computer vision, natural language processing and knowledge representation, search and optimization, fuzzy systems, and planning and decisionmaking." Moreover, using the subset of papers that have been uploaded to the research platform arXiv, Elsevier (2018) finds that articles about core AI categories that are posted on arXiv have increased by 37.4 percent in the past five years.

These sustained research efforts will continue to expand AI's capabilities. Indeed, Brynjolfsson and McAfee (2014, 52) remark that "we're going to see artificial intelligence do more and more, and as this happens costs will go down, outcomes will improve, and our lives will get better." Already, AI is being applied in four main areas of the marketplace, according to Lee and Triolo (2017): (1) the Internet (e.g., online marketplaces); (2) business (e.g., data-driven decisionmaking); (3) perception (e.g., facial and voice recognition); and (4) autonomous systems (e.g., vehicles and drones). Take, for instance, the domain of perception AI. One discovery helps individuals who have historically been visually impaired to use a device with digital sensors that can survey the physical environment and create sound waves through the bones of the head. The technology clips onto eyeglasses, and after being oriented toward text within the user's vision and signaled to read the source by the wearer, the device reads and verbalizes the text (Brynjolfsson and McAfee 2014). Similarly, Brynjolfsson and McElheran (2016) also illustrate how manufacturing establishments using data to influence their decisionmaking exhibit greater productivity than their counterparts. Companies in the digital economy will increasingly compete based on their ability to use data efficiently and strategically.

Technological Progress and the Demand for Labor

This section explores the interaction between technological progress and the demand for labor. First, it gives a brief history of technological change and work. Then it describes the effects of technological progress on investment and wages. Finally, it considers how specialization and comparative advantage affect trade between people and machines.

A Brief History of Technological Change and Work

Do technological advances reduce employment? That is not a new question—concern about job losses caused by automation dates back at least two centuries. During the early 19th century, English artisans (Luddites) in the rapidly changing textile industry famously attempted to destroy the mills and

automated machine looms that they believed threatened their livelihoods. Despite the opposition of the Luddites to automation, the next two centuries witnessed a transition to mechanization of much of the physical labor performed by workers (Galor and Weil 2000). The agriculture sector provides a notable example. Tractors replaced horsepower and manual labor in 19th-century plowing work, and labor-intensive manual tasks were mechanized (Rasmussen 1982). Similar examples abound among many types of skilled artisanal work after the introduction of machine tools, as well as the transformation of manufacturing after advances such as steam power and electricity.

Automation's effects on labor are no longer confined to manufacturing and agriculture (Brynjolfsson and McAfee 2014; Autor 2015; Polson and Scott 2018). Computers and constantly evolving software have eliminated the need for many of the administrative and clerical tasks that had long been performed by white-collar workers in commercial business. Indeed, before the word "computer" referred to a microprocessor on a desk, it was a job title for a person who laboriously performed simple arithmetic or more complex mathematical calculations. Today, an accountant or financial specialist can do in seconds what would have once taken hours or days of painstaking computation by a team of educated people. An online tax preparation system can do much of what a professional certified public accountant might have done, while being faster and more accurate. White-collar work environments are likely to undergo further disruptive changes as AI technologies continue expanding into logistics and inventory management, financial services, complex language translation, the writing of business reports, and even legal services. Even medical diagnoses are likely to involve AI technologies in the foreseeable future.

Economists and policymakers have long studied the question of job displacement caused by technological advancement. In just one example, in 1964 Congress authorized the National Commission on Technology, Automation, and Economic Progress to study the effects of technological advancement, particularly in relation to unemployment. The commission's 1966 report included the finding that "technology eliminates jobs, not work" (Bowen 1966, 9). In a more contemporary discussion, David Autor (2015, 5) noted that "journalists and even expert commentators tend to overstate the extent of machine substitution for human labor and ignore the strong complementarities between automation and labor that increase productivity, raise earnings, and augment demand for labor." Though the introduction of new technologies can create job displacement, examining technological change from a historical perspective shows that these transformations do not lead to permanently lower employment, but rather an increase in demand for new tasks (Mokyr, Vickers, and Ziebarth 2015).

Effects of Technological Progress on Investment and Wages

Capital investments, such as in machines and software, embody AI, which Brynjolfsson, Rock, and Syverson (2017) call a general purpose technology. New investments that embody AI are expected to be more like ("closer substitutes for") labor than traditional capital investments were. Here, we begin by relating capital to labor and productivity and explain why labor is expected to receive most of the net benefits from AI in the long run. In particular, we argue that, though AI is expected to increase real wages on average, the economy has three phases of adjustment where the wage effects are different. In the anticipation phase, real wages are somewhat elevated as businesses begin to switch to activities that are intensive in cognitive tasks, but still do not have machines to adequately perform those tasks. Then, AI arrives and can fill many of the positions, temporarily depressing real wages during the implementation phase as workers compete with the new machines. In the long run, business formation catches up with the new technology and real wages are higher.

Growth in labor productivity can come from changes in three distinct factors: a rise in the quality of labor, which can occur with greater education, training, or skill attainment; a rise in capital, which occurs when firms invest in productive inputs, such as machines, factories, or computers; or a rise in what economists call total factor productivity (TFP), which pertains to other determinants of productivity, ranging from regulatory frictions to unmeasured quality improvements (Solow 1957).

TFP growth often increases real wages and the return to capital in the short run because it makes the factors more productive.[6] A greater return to capital also stimulates additional investment leading to business creation and growth. As a result of the additional capital, real wages rise and, because new capital competes with old capital, the return to capital declines. Indeed, a century or more of economic growth has increased real wages by more than a factor of five (Fisk 2001; Zwart, van Leeuwen, and van Leeuwen-Li 2014), while the return to capital has been almost constant over time (Caselli and Feyrer 2007; Mulligan and Threinen 2011). Nearly all the long-run benefits of TFP go to labor by reducing the effective prices of goods and services or by raising total compensation (Caselli and Feyrer 2007; CEA 2018c).

Although real wages trend up and the return to capital does not, as discussed above, labor's share of gross domestic product (GDP) can be constant, rising, or falling, depending on the type of technological change and the degree to which the new investment substitutes for labor in the production process. In other words, some types of TFP growth may reduce labor's share of GDP in the long run even while the entire benefit from TFP growth goes to workers in the form of higher real wages. For example, Karabarbounis and Neiman (2014)

[6] Our discussion of wages in the text that follows views it as representing all compensation from work, including fringe benefits.

show that the decline in the relative price of investment goods (e.g., due to the expansion of information technology and computers) helps to account for the decline in the labor share.

Although the TFP growth occurring during most of the 20th century did not reduce labor's share of national income (Kaldor 1961), AI might reduce it in the long run to the degree that it is more substitutable for labor than 20th-century capital investments were. The transition to a labor-substitutable AI is illustrated in figure 7-2 from the perspective of the capital market. Because a downward-sloping capital demand curve shows the relationship between the amount of capital and its marginal contribution to output, the area under the curve up to the equilibrium amount of capital is equal to the total amount of output. This output is divided between capital and labor, with capital's income equal to the rectangular area, which has dimensions equal to the amount of capital and the rental rate per unit of capital. In the figure, the triangular area above the rectangle is the output not paid to capital, which is labor income.

The arrival of AI makes new capital investments more productive, which is why the capital demand curve is shifted up by the discovery. Initially, AI investments earn returns greater than the normal capital return, as at the point *b* in figure 7-2, which stimulates more investment. The additional investment begins to drive down the return to capital, but more slowly than investment did in earlier eras, because the new investment does not compete as directly with existing capital, which is why the new demand curve is flatter than the old one. In the long run, the return to capital falls back to normal, the economy is at point *c* in figure 7-2, and labor income has increased by the amount of the shaded area *L*.[7] Labor's share is lower in the long run than it was before AI arrived, as shown in the diagram by the fact that the rectangular increment to capital income is disproportionate to *L*. Ironically, the addition to capital income is a symptom of more investment and real wage growth due to the assumption that AI investments are more substitutable for labor than older types of capital.

In the short run, after the arrival of AI, new investment that is a good substitute for labor reduces real wages to the extent that human workers compete with AI for jobs and the additional business formation is not yet complete. This phase resembles the commonly expressed concern that workers would be harmed by AI. In terms of figure 7-2 the capital rental rate *r* at point *b* is temporarily elevated, at the expense of labor income. However, it is important to also consider the phase *before* AI arrives. Here, real wages are elevated by the anticipation of AI because businesses are formed with the expectation that they will eventually have both human and machine labor, but in the meantime will need to perform their operations entirely with human labor.

[7] In the limit in which AI is a perfect substitute for human workers, the area *L* is zero. The subsection of this chapter titled "Trade between People and Machines" explains why the perfect-substitution case is ruled out by market forces.

Figure 7-2. The Effect of AI on the Amount of Capital and the Distribution of Factor Incomes

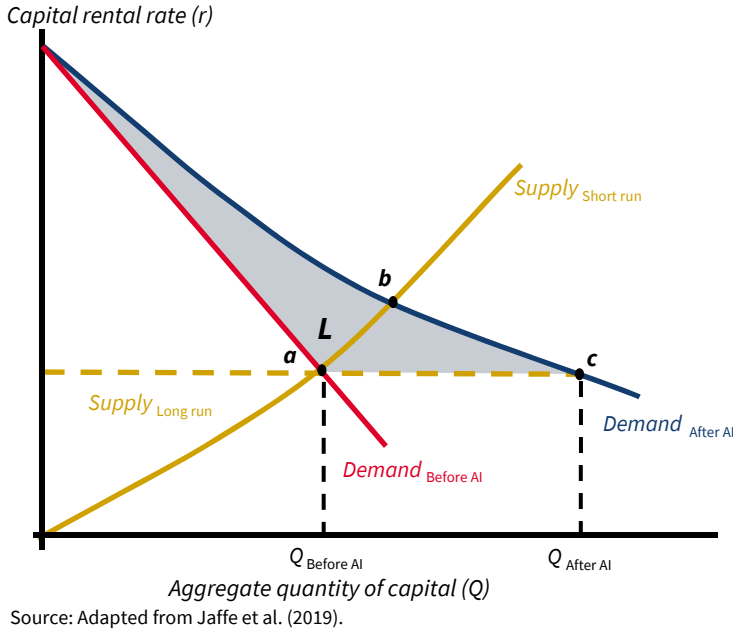

Source: Adapted from Jaffe et al. (2019).

This stylized discussion highlights the situation that though AI can depress real wages for a period if it is a good substitute for labor, ultimately AI will raise real wages above what they were before AI because of the investment and increased productivity that it stimulates. These conclusions are consistent with not only theoretical models of economics featuring AI in general equilibrium (Aghion, Jones, and Jones 2017) but also with evidence on how the introduction of robots raised labor productivity across 17 countries between 1993 and 2007 (Graetz and Michaels 2018). Moreover, taking the information technology (IT) revolution as an analogue, Autor, Katz, and Krueger (1998) show that the introduction of computers led to strong and persistent growth for skilled workers, which accounts for the increased demand (and subsequent expansion of supply) for workers who have gone to college.

In summary, even though AI is expected to temporarily decrease real wages, in the long run it will increase real wages, on average, because of the investment it stimulates. The next section highlights the role of comparative advantage behind the reallocation of tasks across and within sectors of the labor market (Acemoglu and Autor 2011), explaining how firms will apply AI in ways that are complementary to labor and therefore have a more positive effect on real wages, and a less negative effect on labor's share of GDP than shown in figure 7-2.

Trade between People and Machines

When and how much is AI likely to substitute for human tasks? The principle of comparative advantage tells us that human workers can benefit from being in the same market with machines, even if these machines excel at many traditionally human tasks. The benefit comes from workers' specialization in the tasks humans can do better than machines, or at least the tasks where humans are at the smallest disadvantage (Autor 2015). Specialization allows the machines to be used on their best tasks without wasting resources on tasks that people can do at a lower opportunity cost. To put it another way, even if it were technologically possible to let machines do all tasks, and do them better than humans do, an owner of the machines would sacrifice profits by deploying them without regard for specialization.

Consider the operation of a store that requires cashier tasks, communication with suppliers, the delivery of products, and arranging displays. The AI machines perform the arrangement tasks 10 times better (in terms of speed and accuracy) than humans, and perform the other tasks 20 times better. Given comparative advantage, and assuming that the machines are cheap enough to justify using them for any task, profit-maximizing deployment will have workers performing the arrangement tasks, thereby freeing up machines to do the other tasks where they are especially productive. The theory of comparative advantage means that humans inevitably have a comparative role to play, even if they do not have an absolute advantage in every task.

Moreover, the choice of which machines to deploy is not merely determined by what is technically possible with engineering and computer science.[8] Robocop, *Star Wars*' C-3PO, and other near-human machines are great entertainment, but in many situations they would be poor investments precisely because of their close similarities to humans.[9] Because machines and AI are ultimately another form of capital, designing machines to complement, rather than substitute for, humans will be more profitable. In other words, the potential for specialization implies that producers will look for ways to magnify differences with people. For example, Abel and others (2017) explain how providing algorithms with expert (human) advice—part of a broader class of "Human-in-the-Loop Reinforcement Learning"—can improve various aspects of learning and prediction.

[8] Consider the analogous case of agricultural tobacco production. Though some countries, like Brazil, display very labor-intensive tobacco production (Varga and Bonato 2007), U.S. production of tobacco is highly mechanized (Sykes 2008). For a similar illustration from cotton production, see FAO (2015). In this sense, the mere presence of capital does not guarantee its use; the opportunity cost of labor in an economy will drive the division of labor and degree of specialization. Lagakos and Waugh (2013) formalize these insights within a general equilibrium Roy model with agriculture and nonagriculture sectors.

[9] Research in human–machine interaction finds situations in which people can more easily and intuitively work with robotic partners when the robots look and behave in ways similar to humans. In these cases, people can project human expectations of how robots should act, and thus do not need to be trained (or study user manuals) in order to figure out how to work with the robot.

The purposeful acquisition of comparative advantage has long been observed in human labor markets (Becker and Murphy 1992). Consider an electrician and a carpenter who work together to build a high-quality house. Their comparative advantage is obvious at the time that they are building the house, but neither of them was born with his or her specialized skills. They both chose to specialize knowing directly—or perhaps indirectly, through market prices—that they would be a more valued member of a construction team if they could excel at carpentry, or excel at electrical work, rather than having mediocre skills at both types of tasks. Robotics research already suggests that productivity is enhanced when machines specialize (Nitschke, Schut, and Eiben 2012). Also see, for example, box 7-1, which describes the Defense Advanced Research Projects Agency's (DARPA's) initiatives regarding "partnering with machines." In light of these examples of complementarity between AI and humans, the entertainment industry's anthropomorphic portrayal of robotics and artificial intelligence is somewhat misleading about how much these types of investments will substitute for human workers.

The concern, of course, is that the price associated with human tasks will decline to a point where humans are driven out of the workforce and are not incentivized to work. For example, some manufacturers might find that production is cheaper with complete automation, rather than by retaining a mix of some human employees and AI. However, specialization and trade also occur at the market level. A robot-intensive business may engage in one phase of production, selling its output to a person-intensive business at a later phase of production. In this sense, even if certain tasks traditionally performed by humans are instead now done by machines, humans will nonetheless hold a comparative advantage for other tasks and thus will continue to play a role in production processes.

Although there are some concerns about complete automation of human activities (Frey and Osborne 2017), the emerging empirical evidence suggests that the main effects of AI and automation are on the composition of tasks within a job, rather than on occupations in general. For example, Brynjolfsson, Rock, and Mitchell (2018) introduce an index of suitability for machine learning (SML), and they find that, though most occupations have at least some tasks that are SML, few (if any) have tasks that are all SML. Similarly, Nedelkoska and Quintini (2018) use data on skills across occupations and 32 countries, and they find that, though 14 percent of jobs are likely to be automated by over 70 percent, 26 percent of jobs face a change of automation of 30 percent or less. The key observation is that, as automation progresses, workers will increasingly be drawn to the jobs and tasks that are more difficult to automate. Astute policymaking will nonetheless play a role in promoting workforce development, particularly for less educated workers—through, for example, the Pledge to America's Workers, which we discuss later in the chapter.

Box 7-1. DARPA: Strategic Investments in Artificial Intelligence and Cybersecurity

The Defense Advanced Research Projects Agency (DARPA) is focused on a future where AI is a complement to humans in the production of goods, services, and ideas—that is, where humans can safely "partner with machines" more as colleagues, rather than as tools (DARPA 2018a). To facilitate this vision, DARPA is actively funding the development and application of a so-called third wave of AI technologies that would result in intelligent machines capable of reasoning in context. In particular, DARPA announced a $2 billion, multiyear investment in new and existing programs in September 2018. These investment areas include "security clearance vetting or accrediting software systems for operational deployment; improving the robustness and reliability of AI systems; enhancing the security and resiliency of machine learning and AI technologies; reducing power, data, and performance inefficiencies; and pioneering the next generation of AI algorithms and applications, such as 'explainability' and commonsense reasoning" (DARPA 2018a).

DARPA has already piloted a number of successful programs, including the Cyber Grand Challenge in 2016—a competition that showcased the state of the art in Cyber Reasoning Systems (DARPA 2018b). Competing systems played an "attack-defend" style of "Capture the Flag," where contestants were tasked with developing AI algorithms to "autonomously identify and patch vulnerabilities in their own software while simultaneously attacking the other teams' weaknesses" (Hoadley and Lucas 2018).

Although conventional cybersecurity programs may take up to several months to find and patch problems, the competing and largely rules-based algorithms found the bugs in seconds. According to DARPA (2016), "the need for automated, scalable, machine-speed vulnerability detection and patching is large and growing fast as more and more systems . . . get connected to and become dependent upon the Internet." The major innovation in the Cyber Grand Challenge was the demonstration that AI can play both an offensive and defensive role. DARPA continues to build out these human-machine cyber detection capabilities for pinpointing and addressing vulnerabilities through its Computers and Humans Exploring Software Security program, known as CHESS. The activities funded by CHESS involve helping computers and humans work collaboratively through tasks, such as finding zero-day vulnerabilities at scale and speed.

The Uneven Effects of Technological Change

This section delineates the uneven effects of technological change. It first considers these changes' differential effects by occupation and skill. Then it explores the scale and factor-substitution effects of an industry's technological progress and how they moderate the effect on labor. Finally, the section asks when we will see the effects of AI on the economy.

Differential Effects by Occupation and Skill

Many types of technological change affect workers and industries in heterogeneous ways. For example, the widespread adoption of computers and information technology during the past several decades has enormously increased productivity for certain types of workers, but has brought comparatively little or no productivity enhancement for others (Acemoglu et al. 2014). Because earnings are determined by workers' productivity, such changes in technology are expected to have varying effects on workers with different sets of skills, such as workers with or without a college or graduate education (Katz and Murphy 1992).

Economists have concluded that "skill-biased technical change" can account for most of the observed rise in earnings disparities between some higher-skilled workers (whose productivity was greatly enhanced by technology, like computers) and some lower-skilled workers (who were less affected), which was amplified during the IT revolution (Autor, Katz, and Krueger 1998; Autor, Levy, and Murnane 2003). This disparity is in part explained by the complementarity between capital and certain types of skills (Krusell et al. 2000). In the context of AI and automation, the complementary relationship means that there is processing power that mainly benefits workers who use computer technology. In this sense, the more rapid increase in earnings among college-educated workers, despite the corresponding rise in the supply of these workers, represents a skills premium for individuals who can leverage technology to augment their productivity (Juhn, Murphy, and Pierce 1993).

The Scale and Factor–Substitution Effects of an Industry's Technological Progress

Technological progress allows an industry to produce the same output with fewer inputs (e.g., workers). At first glance, we might therefore expect workers to leave the industry and find work elsewhere. One could point to the example of changes in agriculture in the 20th century, when the agricultural employment share dropped from 41 to 2 percent between 1900 and 2000 (Autor 2015), at the same time that agricultural TFP rapidly increased (Herrendorf, Rogerson, and Valentinyi 2014). See box 7-2 for an example of technological change in the agricultural sector that has fueled productivity.

Box 7-2. Technological Change in Agriculture and Rural America

Agriculture has been one of the sectors experiencing rapid technological change, including the computer science revolution. For example, output per hour in the agricultural sector grew annually by 4.3 percent between 1948 and 2011, whereas it grew annually by 2.4 percent in manufacturing (Wang et al. 2015).

For example, precision agriculture—which refers to a broad class of AI applications allowing for precise control over agricultural inputs based on detailed, site-specific data—has allowed farmers to improve the productivity of soil by better understanding the characteristics that are most conducive to growth within a specific geographic area; see figure 7-i for evidence on its incidence across peanut and soybean farming. Moreover, these systems contain sensors that allow farmers to monitor crop yields and self-guided tractors and variable rate planters that vary their seeding and fertilizer rates based on fertility and past yield data. In brief, these technologies have allowed corn and soybean farmers, among others, to produce more at lower costs (Schimmelpfennig 2016).

Figure 7-i. Precision Agriculture Use in Peanuts and Soybeans

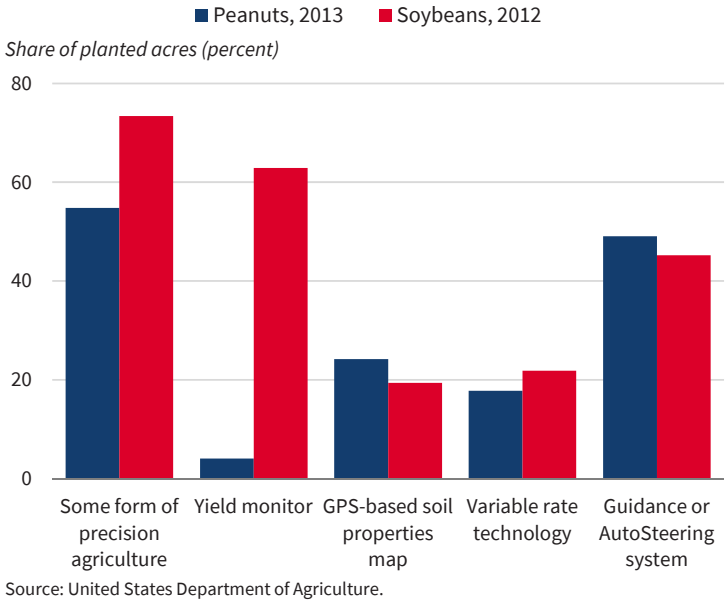

Source: United States Department of Agriculture.
Note: GPS = Global Positioning System

AI is also used in animal agriculture. For example, over 35,000 robotic milking systems are in operation globally on dairy farms. According to Salfer and others (2017), farms using robotic milking systems are much more productive, selling 43 percent more milk per hired worker and 9 percent more milk per cow. Moreover, rather than displacing humans, the introduction of automation in dairy farms has allowed labor and management to reallocate their time toward maintaining animal health, analyzing records, and managing reproduction and nutrition on the farms. For example, John Deere runs a two-year associate degree program to help its employees not only stay current on the latest farming machine tools but also acquire new skills in data science (Burkner et al. 2017).

However, rural Americans have not always seen the gains of technological progress (Forman, Goldfarb, and Greenstein 2012). Motivated by these disparities, President Trump signed Executive Order 13821 in January 2018 (White House 2018c), expanding and streamlining access to broadband in rural America. Given the importance of high-speed Internet access for data science capabilities, connectivity in rural America is essential for its economic competitiveness. Moreover, the Trump Administration is committed to investing in and promoting workforce development through, for example, the Pledge to America's Workers, which we discuss in below in this chapter's main text.

However, as an industry's productivity advances, it is producing each unit of output at a lower cost and thereby selling at lower prices. Consumers of this output respond by purchasing more, which is a force toward more industry employment known as the "scale effect" on labor demand. The productivity revolution in agriculture did result in more production and higher sales of food. However, because consumers' demand for agricultural output is price inelastic—consumers spend less of their budget on agriculture when it becomes cheaper—the "factor-substitution effect" dominated the scale effect on the demand for agricultural labor.[10]

If demand for a good is price elastic—meaning that consumers spend more of their budget on the good when prices fall—then cost-reducing technology might raise that sector's shares of employment and GDP. Consider the recent history of taxi dispatchers, who take calls from individuals desiring a ride and direct a driver to the pickup point. About a decade ago, companies discovered how to use a smartphone to perform the tasks of the dispatcher, and these companies famously distributed such an app to millions of smartphone users. The result was a dramatic increase in the number of people working in the transportation industry, broadly understood to include drivers for Uber,

[10] The decomposition of labor demand into scale and factor-substitution effects is usually attributed to Alfred Marshall (1890) and John Hicks (1932).

Lyft, and other ride-sharing platforms. By observing what happened to overall employment in the industry (which provides rides for passengers, and which now includes ride sharing in addition to traditional taxis), we can see that it had price-elastic demand. The cost reductions associated with the new technology increased the number of rides even more than it increased the number of humans giving rides.

Although there is some difficulty in measuring participants in the sharing economy in ways that are directly comparable with traditional taxi employment, there is emerging evidence of its expansion. For example, JPMorgan Chase (2018) found that the share of families generating earnings on transportation platforms over the course of a year increased to 2.4 percent of the labor force in March 2018 after the inception of ride sharing in about 2010 (figure 7-3).[11] A large part of the increase came from the introduction of 460,000 driver-partners in just three years under the Uber platform alone (Hall and Krueger 2018). Increasing empirical evidence suggests that these ride-sharing applications not only have provided significant flexibility for drivers (Chen et al., forthcoming; Koustas 2019) but also have generated social welfare benefits for those who are not platform participants (Cohen et al. 2016; Makridis and Paik 2018).

These ride-sharing applications are an early, pre–autonomous vehicle (AV) manifestation of transportation as a service. Whereas transportation has traditionally been about assets (i.e., vehicle ownership), it may increasingly move toward services as more AVs enter the market. For example, even though PricewaterhouseCoopers (PwC) estimates that the transportation sector may require 138 million fewer cars in Europe and the U.S. by 2030 (PwC 2018a), it also estimates that the market for shared, on-demand vehicles may grow to $1.4 trillion by 2030, in comparison with $87 billion in 2017 (PwC 2018b). Though predicting the growth in the AV market is outside the scope of this Report, the emerging patterns in ride sharing and AVs are illustrative examples of the impact of technological change.

When Will We See the Effects of AI on the Economy?

Some economists have noted a puzzling productivity paradox with the historical and ongoing patterns described above. Although most researchers agree that the recent advances in AI and automation promise production possibilities that are even greater than the initial emergence of the digital

[11] The National Academies (2017) also cite estimates pointing toward growth from 10 to 16 percent in alternative work arrangements between 2005 and 2015. According to Katz and Krueger (2018), who did a survey in November 2015, 0.5 percent of workers report working through an online intermediary. Though there is debate about the measurement of alternative work arrangements, a recent assessment by Katz and Krueger (2019) concludes that, despite the only modest increase in these arrangements obtained from the 2005 and 2017 Contingent Work Surveys in the Current Population Survey, this survey's data are likely underestimates.

Figure 7-3. Share of Respondents Reporting Income from Ride-Sharing Platforms in the Past Year, 2013–18

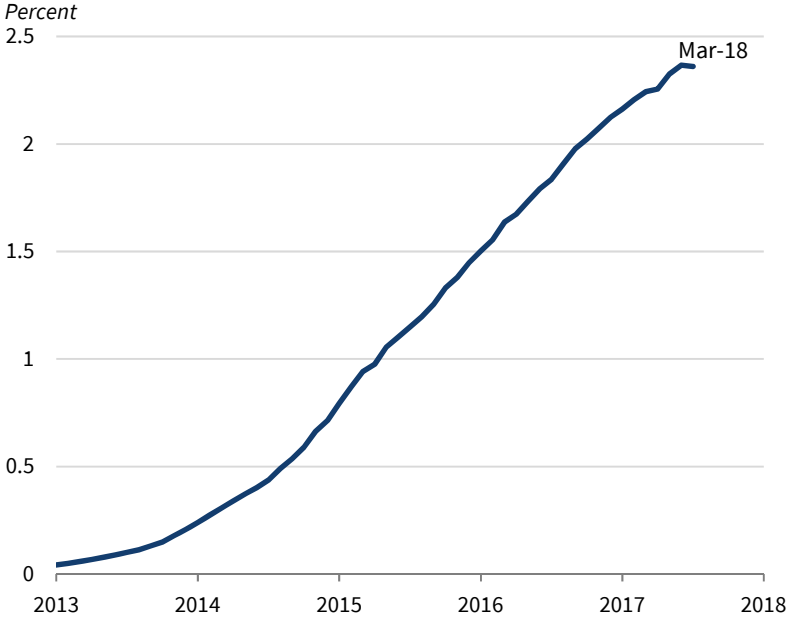

Source: JPMorgan Chase (2018).

economy (Brynjolfsson and McAfee 2014), the growth of labor productivity, at least in the way it has traditionally been measured, has been surprisingly sluggish.[12] For example, in contrast to the 2.8 percent annual growth in aggregate labor productivity seen in the United States between 1995 and 2004, its annual growth between 2005 and 2015 was only 1.3 percent (Syverson 2017). This pattern is consistent with growth across other economies; Syverson (2017) found the annual growth rate in labor productivity was 2.3 percent between 1995 and 2004 in 29 sampled countries, but fell to 1.1 percent between 2005 and 2015.

If technological change and the adoption of AI have been especially rapid during the past decade, what can account for the slower growth of labor productivity? One possibility is that the productivity effects of technology may have been oversold (Gordon 2000) and the period of rapid growth of the Information Age was a temporary aberration in a long-run trend toward slower technology-related productivity growth (Gordon 2018). However, Oliner and Sichel (2000) show, using a multisector neoclassical growth model with both IT and non-IT capital, that the increase in IT and corresponding efficiency gains account for two-thirds of the increase in labor productivity for the nonfarm

[12]As the Nobel laureate Robert Solow famously said, "You can see the computer age everywhere but in the productivity statistics."

business sector over the 1990s.[13] Moreover, Byrne, Oliner, and Sichel (2013) apply the same framework and fit more recent data between 2004 and 2012, suggesting that there is no inconsistency with theory. Jorgensen and Stiroh (2000) also obtain slightly lower contributions to growth from computer hardware because they use a broader definition of output. Yet another related explanation is that the expansion of credit in the early 2000s led to a misallocation of investment into less productive sectors, creating a drag on growth (Borio et al. 2016). However, productivity has recently ticked up (e.g., see chapter 10 of this *Report*). Therefore, secular stagnation and the misallocation of investment do not appear to be viable explanations.

Another possibility is that our official estimates of growth and productivity fail to capture many of the recent gains from technological advancement. Many of today's new technologies involve little or no direct cost to consumers, but give them great utility. These developments include, for example, Internet social networks, information search capabilities, and downloadable media. A quick Internet search today can yield information that, a few generations ago, would have required a team of individuals searching a university library—such benefits are not captured in our measurement of GDP. Though these benefits are clearly important factors behind consumer welfare (Brynjolfsson, Eggers, and Gannamaneni 2018), mismeasurement between 2005 and 2015 would need to be unrealistically high to account for the sluggish GDP growth, relative to the overall trend (Syverson 2017).

Perhaps the strongest argument for why productivity statistics in recent history have not shown the expected benefits from the new technologies is that, for practical reasons, there have so far simply been lags between productivity and the widespread implementation of AI and ML. The theoretical genesis of this argument is an insight from Paul David (1990). Much as the dynamo and the computer were fundamental components of a broader technological infrastructure, AI is a similar general purpose technology. Although these discoveries often have immediate effects on productivity, their full impact is not realized until all the complementary investments are made, thereby creating a lag with investment. Brynjolfsson, Rock, and Syverson (2017) apply this logic to AI, reconciling the productivity paradox. Under their preferred interpretation of the data, we are simply awaiting the results of a necessary trial-and-error process and the productivity benefits will eventually be realized.

[13] An integral part of the efficiency gains among producers of computer equipment is the rapid decline in effective prices of semiconductors due to advances in chip technology (Triplett 1996). Byrne, Oliner, and Sichel (2017) find that semiconductor prices measured with a hedonic index fell at an estimated annual rate of 42 percent between 2009 and 2013, much faster than the 6 percent decline experienced by the microprocessor producer price index series that provides a broader measure that subsumes semiconductors.

Cybersecurity Risks of Increased Reliance on Computer Technology

Although technological advances and the emergence of AI have the potential to raise productivity and economic growth, the widespread reliance on technology also exposes the economy to new threats of malicious cyber activity. Cyber threat actors may be nation-states, cyber terrorists, organized criminal groups, "hacktivists" (individuals or collectives that aim to advance their social agenda through cyber interference), or simply disgruntled individuals. These threats transcend the typical boundaries of conflict, which have been analyzed through the lens of land, sea, and air. However, the emergence of the "Internet of Things" implies that anything connected to the Internet is vulnerable to malicious cyber intrusions, introducing threat vectors throughout the Internet ecosystem (Hoffman 2009).

Malicious cyber activity imposes costs on the U.S. economy through the theft of intellectual property and personally identifiable information, denial-of-service attacks, data and equipment destruction, and ransomware attacks. The CEA estimated this cost to be as high as $109 billion in 2016 (CEA 2018b). Most innovations, however, lead to little-understood risks, whether for new drugs or computer technologies. This section describes our current assessment of the scope of cyber vulnerabilities, how they vary by industry, and the factors that may exacerbate failures to adopt cybersecurity best practices.

Assessing the Scope of the Cyber Threat

The 2018 *Economic Report of the President* (CEA 2018b) estimated the 2016 costs of malicious cyber activity by adding up the costs experienced by the private sector, the public sector, and private individuals. It estimated the costs to the private sector using event-study methodology, whereby it quantified the loss of firm value as a result of an adverse cyber event. It estimated the costs to the corporate sector using event-study methodology, whereby it quantified the loss of firm value as a result of an adverse cyber event. The estimate further took into account the spillover effect of these costs to economically linked firms. On the basis of a sample of cyber incidents occurring between January 2000 and January 2017, the *Report* estimated that the total economic cost for 2016 ranged between $57 and $109 billion.

Although these event studies provide an important starting point for evaluating the costs of cybersecurity incidents, they presuppose that the timing of the event was reliably recorded and that investors knew the distribution of new risks induced by the event. However, to give just one example, when the largest recorded data breach, according to the Privacy Rights Clearinghouse, occurred in late 2013, it was not reported until September 2016 (Lee 2016). Delays between the time when an incident takes place and the time it is reported are a function of not only a firm's ability to identify the incident but

also of varying State laws that mandate disclosure (Bisogni 2016).[14] The affected firm's own estimate of the damage caused by the 2013 breach has been updated and increased on several occasions, illustrating how difficult it can be to accurately calculate the cost. Moreover, data on the number of records or systems that have been breached often contain significant measurement error and sampling variability.

In addition to reporting discrepancies across States, there are also discrepancies across sectors. Makridis and Dean (2018) study sector discrepancies using data from the Privacy Rights Clearinghouse and the Department of Health and Human Services to investigate the relationship between recorded breaches and firm outcomes. Though they find some evidence of a negative association between productivity and record breaches in the Health and Human Services data, where healthcare companies face greater disclosure requirements, they do not find such evidence in the data from the Privacy Rights Clearinghouse covering all sectors. Publicly traded companies, based on requirements from the Securities and Exchange Commission (SEC 2011), must provide timely and ongoing information in the periodic reports of material cybersecurity risks and incidents that trigger disclosure obligations. Beyond the Federal securities laws, other reporting standards in specific sectors, like the Health Insurance Portability and Accountability Act, may result in disclosures of other data breaches that are not material.

Since 2009, the National Cybersecurity and Communications Integration Center (NCCIC) of the Department of Homeland Security (DHS) has served as the Nation's flagship cyber defense, incident response, and operational integration center. The NCCIC serves as the national hub for cyber and communication information, technical expertise, and operational integration, operating a 24/7 watch floor tasked with providing situational awareness, analysis, and incident response capabilities to the Federal government; private sector stakeholders; and State, Local, Tribal, and Territorial Partners. Through this process, DHS has been collecting robust data on the types of incidents that are having an impact on the Nation. Furthermore, the Federal Bureau of Investigation (FBI) also maintains CyWatch, a 24/7 command center for cyber intrusion prevention and response operations based on consensual monitoring and third parties that report to the FBI. CyWatch monitors must notify companies whose network security has been breached (34 U.S.C. § 20141 creates an obligation for Federal law enforcement agencies to notify victims of a crime). After notification, CyWatch shares information with its partner law enforcement agencies—including the Department of Defense, DHS, and National Security

[14] Using data from the Privacy Rights Clearinghouse, Bisogni, Asghari, and Van Eeten (2017) estimate that adoption of the "inform credit agency" and the "notification publication by informed attorneys general" State provisions would increase the number of publicly reported cybersecurity breaches by at least 46 percent.

Figure 7-4. Cybersecurity Breaches That Were Made Public, 2005–18

Number of breaches

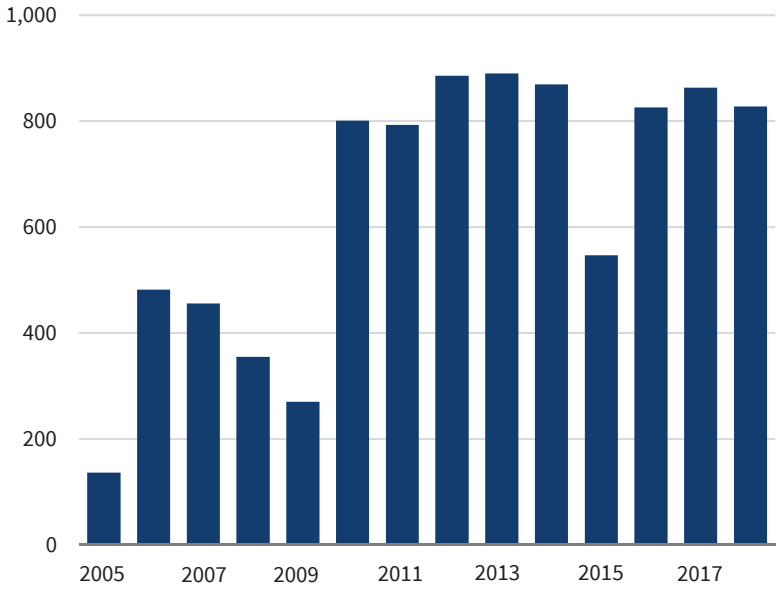

Sources: Privacy Rights Clearinghouse; CEA calculations.

Agency—to improve preparedness and attribution behind attacks and guide appropriate responses.[15]

Despite the serious limitations associated with data from the Privacy Rights Clearinghouse, they nonetheless provide a time series proxy for the increased frequency of data breaches since 2005; see figure 7-4. Although there is an upward trend in cyber breaches between 2005 and 2018, these data largely understate the number of data breaches (Bisogni, Asghari, and Van Eeten 2017; ITRC 2019). The Internet Crime Complaint Center, a partnership between the FBI and National White Collar Crime Center, gives victims of cybercrime an accessible reporting mechanism for alerting the authorities about suspected criminal or civil violations. Although not directly comparable, the 2017 "Internet Crime Report" announced a total of 301,580 complaints of cyber breaches in 2017. Even though these complaints represent a broader range of potential Internet crimes, the number far exceeds the 863 publicly reported incidents.

Recommending possible solutions for these cyber vulnerabilities requires an accurate understanding of their sources. We suggest that there are at least

[15] Though exact attribution in cyberspace is possible, it requires not only technical expertise but also leadership and information sharing and coordinating across the layers of an organization (Rid and Buchanan 2015).

two underlying drivers behind the above-mentioned empirical regularities. First, organizations could lack informational awareness. Much like the quantitative management science literature on the adoption of best practices in business (Bloom et al. 2013), many organizations might simply not be aware of basic cyber hygiene practices. Second, the executives of organizations could suffer from incomplete incentives to promote cybersecurity practices. If, for example, financial metrics are easier to measure, relative to cybersecurity, then managers might allocate too little effort to cybersecurity due to a "multitasking problem" (Holmstrom and Milgrom 1991). Particularly because cybersecurity breaches generate network externalities, the private sector could underinvest in cybersecurity (Gordon et al. 2015).

Our preceding evidence on the lack of many basic cybersecurity practices among the most profitable companies in the U.S. economy suggests that a lack of information awareness and a lack of resources are unlikely to be the primary culprits behind existing vulnerabilities. Moreover, the "Cybersecurity Framework" of the National Institute of Standards and Technology's (NIST 2014), which details best practices, is publicly available and has been disseminated through many channels. These facts suggest that the alternative culprit could be incomplete incentives arising from agency problems within organizations that lead managers to overlook cyber hygiene.

Information sharing and dissemination of best practices must remain a priority, particularly for small businesses that are more likely to lack the resources or infrastructure to search out and implement best practices. In particular, information needs to be publicly available, transparent, and shared to disseminate best practices and call attention to dangerous practices. For example, Gal-Or and Ghose (2005) show that industry-based information sharing and analysis centers can lead to improvements in social welfare, but the degree of competition in the marketplace is an important moderating factor that determines whether a firm participates. In particular, unless firms in an industry understand the downside associated with their vulnerability to cyberattacks, they may not realize the gains that can come from collaboration through information sharing.

Many security operations companies also provide a source of market discipline by promoting transparency and information vis-à-vis cyber vulnerabilities (such organizations that raise firms' awareness of cybersecurity flaws are often referred to as "white hat hackers"). Conversely, a survey by Malwarebytes (2018) suggests that roughly 1 in 10 U.S. security professionals admit to considering participating in "black hat hacker" activity, which involves exploiting discovered cybersecurity vulnerabilities for financial gain. Roughly 50 percent of security professionals say they have known or know someone involved in black hat hacking activities.

Potential Vulnerabilities by Industry

The prevalence of cyber threats suggests that firms are relatively unprepared to protect themselves. Indeed, according to Hiscox (2018a), in 2017 nearly three-quarters of organizations based in the United Kingdom, the United States, Germany, Spain, and the Netherlands failed basic cyber readiness tests. Even though the United States ranks higher than most countries in cyber readiness (Makridis and Smeets 2018), its preparedness is still poor enough to concern policymakers studying the impact of cyber insecurity on the U.S. economy.

To better understand these cybersecurity risks at a more granular level, Rapid7, an Internet security firm whose business model involves collecting publicly observable data on cybersecurity practices of any firm with an Internet presence, shared its 2018 data for *Fortune* 500 companies with the CEA. Using public data and a proprietary methodology, Rapid7 matches uniquely identified Internet protocol addresses of Internet-connected devices to a specific firm. Though the security scan is voluntary, only 4 percent of *Fortune* 500 firms opt out. These data show that the majority of *Fortune* 500 companies are vulnerable to cyberattacks, and thus fail to take even the most basic security measures. And though there are many metrics for gauging vulnerabilities, we focus here on an important and transparent metric: whether email has been configured for protection against spam.

Motivated by the frequency of phishing email attacks, which are the most common method used by malicious cyber actors to penetrate network security, configuring a secure email network is one of the first lines of defense. One metric for email security is whether the organization has adopted the Domain-Based Message Authentication, Reporting & Conformance (DMARC) protocol. Although it is not a panacea for all types of phishing attacks, DMARC allows senders and receivers to authenticate whether a message is legitimately from a sender. Adopting DMARC for email makes it easier for organizations to not only identify spam and phishing messages, but also to keep them out of employees' inboxes, thereby reducing the probability that an employee accidentally clicks on a link. Moreover, properly configured DMARC records are able to actively quarantine or reject emails that are a threat to safety by allowing the message's sender to signal to the recipient that the message is protected by a Sender Policy Framework and/or as DomainKeys Identified Mail. We note, however, that DMARC is only one metric out of many and that having it does not guarantee cyber safety.

Figure 7-5 reports the percentage of all *Fortune* 500 firms without a DMARC email configuration, together with value added, across industrial sectors. This figure illustrates significant exposure across industries, ranging from 40 percent of firms in business services to 93 percent of those in chemicals that are not implementing DMARC protocol. Moreover, although we do not interpret the relationship between value added and a lack of DMARC as causal, the data

suggest that a 10-percentage-point increase in share of firms without DMARC in a sector is associated with $345 billion less in value added in that sector (in 2017 dollars). This suggests that greater adoption of DMARC could avoid breaches and phishing scams.

Given that the combined market value of the *Fortune* 500 firms is over $21 trillion, these results suggest that much of this value may be exposed to cyber thefts of intellectual property, various destructive and ransomware attacks, and the destruction of reputational capital. Moreover, as outlined in the 2018 *Economic Report of the President*, an attack on entities—especially large, publicly traded *Fortune* 500 firms that are part of the Nation's critical infrastructure—could have effects throughout the U.S. economy, affecting other firms in the supply chain and individual customers. Given the limited preparedness among *Fortune* 500 companies—manifested by not only the failure to adopt DMARC, but also a range of other cyber vulnerabilities detailed by Rapid7 (2018)—an additional concern is that smaller firms may have even less robust cybersecurity measures in place (Hiscox 2018b).

The Federal government continues to modernize its cyber practices. OMB and DHS worked together to transform the Trusted Internet Connection (TIC) policies and processes so that Federal departments and agencies can take advantage of common and advanced cloud computing capabilities to meet their requirements. AI is not specifically identified in the policy updates, but departments and agencies are now able to use outside expertise in the cloud, which can include using AI and other methods, while continuously meeting appropriate cybersecurity and privacy controls. In alignment with the action steps identified in the *Report to the President on Federal IT Modernization* (American Technology Council 2017), those cooperating in the interagency effort continue to identify if there are any real or perceived policy limitations, by working through cases of real-world use that support their current and future needs. This continuous approach is instrumental for realizing the value of AI and other methods that best meet national needs.

The Federal government is more prepared than the private sector to protect against phishing attacks, which are a primary method for hackers to gain access to enterprises, due to the 2017 Binding Operational Directive 18-01, which introduced requirements for agencies to enhance email and web security. Using data from the 2018 Federal "Cyber Exposure Scorecard," figure 7-6 plots the number of government agencies with various email configurations. In the figure, "fully rejects" means that an organization has properly configured its email, whereas "no rejections" means that it is vulnerable to an attack. Government agencies' use of the DMARC email configuration is 47.9 percent, which is better than the average of 26 percent in the private sector. Moreover, of the 1,018 Federal second-level "dot-gov" domains, 86 percent have a valid DMARC record with a policy of "reject." Though adoption of DMARC is only one of many indicators of cyber hygiene, and was linked to the implementation

Figure 7-5. Industries That Are Most Lacking the DMARC Protocol Among *Fortune* 500 Companies by Value Added, 2017

Value added by industry (billions, 2017)

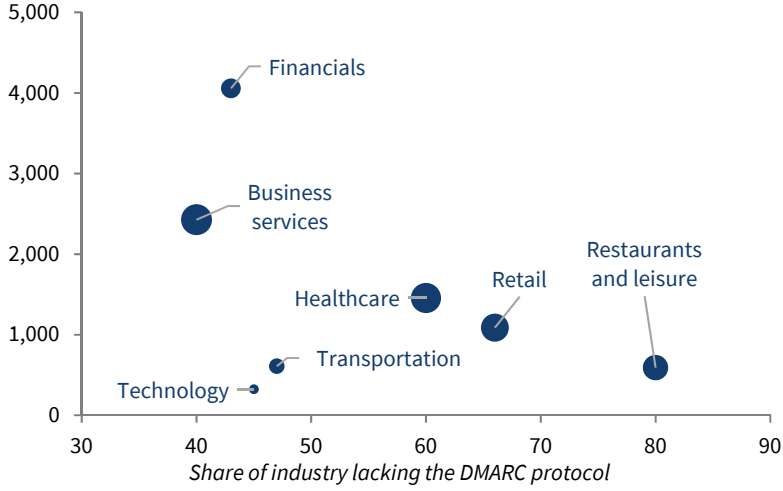

Sources: Rapid7; Bureau of Labor Statistics; Bureau of Economic Analysis; CEA calculations.

Note: DMARC = Domain-Based Message Authentication, Reporting & Conformance, which is an email validation system designed to detect and prevent the use of forged sender addresses for phishing and email-based malware. Points are scaled by industry employment in 2017, and only the top 10 sectors (ranked by employment) are plotted.

Figure 7-6. DMARC Protocol Use Across Government Agencies, 2018

Number of agencies

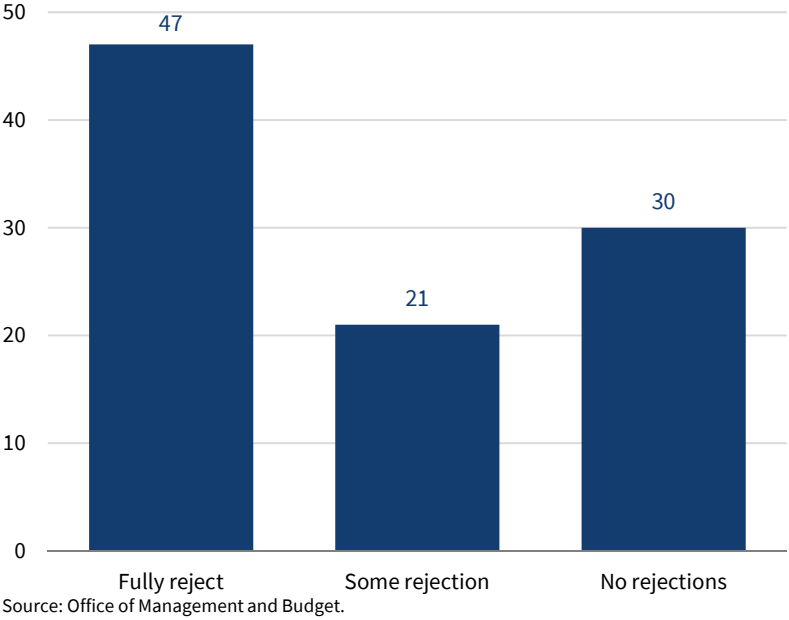

Source: Office of Management and Budget.

of Binding Operational Directive 18-01 across Federal agencies, these results nonetheless suggest that Federal cyber best practices could set an example for the private sector.[16]

The Role of Policy

This section discusses the longer-run policy implications of both AI advancement and cybersecurity issues, and details the Trump Administration's current policies in these areas. The discussion highlights the Administration's priorities for AI readiness and implementation, reskilling, and cybersecurity initiatives to contend with the changing nature of work and emerging technological threats.

Policy Considerations as AI Advances: Preparing for a Reskilling Challenge

As discussed in earlier in this chapter, economists agree that technological change resulting from AI will affect the structure of the demand for labor in the years to come (Brynjolfsson and McAfee 2014; Agrawal, Gans, and Goldfarb 2018). One potential challenge that policymakers could face as AI advances is an increase in the number of workers who need new skills to find work in a changed labor market. Reskilling efforts, both for workers whose jobs have been displaced by technology and for those who need new skills to operate new technologies, could become more urgent as the demand for labor enters a new phase of its decades-long evolution. For example, the World Economic Forum (2018) found in a sample of firms that at least 54 percent of all employees will require significant reskilling and/or upskilling by 2022.

In 2016, the Obama Administration's Council of Economic Advisers examined the economics of AI, including its possible effects on jobs in the future, predicting that "2.2 to 3.1 million existing part- and full-time U.S. jobs may be threatened or substantially altered," by AI. In addition, it predicted roughly 364,000 self-employed "drivers" (ride-sharing workers) would be at risk from a shift toward use of autonomous vehicles as of May 2015 estimates (CEA 2016, 15). However, they also concluded that other workers could see a rise in productivity and increasing demand for certain skills. They identified four areas that could see a rise in labor demand: (1) engaging with AI to complete tasks, (2) developing new AI tools, (3) supervising and maintaining AI tools to ensure they are achieving the desired aims, and (4) responding to paradigm shifts where entirely new approaches are needed (CEA 2016). Because the jobs most vulnerable to automation are concentrated among lower-paid, less-educated workers, reskilling programs could play an important role in helping avert further wage polarization and reallocating skills to where they are most

[16] Although it is also possible that the Federal government does not perform as well in other dimensions, the data from Rapid7 (2018) suggest that the sample of *Fortune* 500 companies also are exposed in other important dimensions of basic cybersecurity practices.

needed. The CEA (2016) made three primary recommendations: (1) investing and developing AI for its many benefits in both the public and private sectors, (2) educating and training workers so they are prepared for the jobs of the future, and (3) helping workers transition across jobs to ensure shared gains from technological change.

More recently, in discussing how automation may interact with the economy and workforce, the CEA (2018a) has referred to an observation made in a report by the National Academies (2017, 140), that continued advance of information technology implies "workers will require skills that increasingly emphasize creativity, adaptability, and interpersonal skills over routine information processing and manual tasks." This report also reiterates findings by the Organization for Economic Cooperation and Development (OECD 2018), among others, that workers who have not obtained a college degree are most at risk for displacement by automation. Similarly, motivated by the declining college and cognitive skills premium—as documented by Beaudry, Green, and Sand (2016); Valletta (2016); and Gallipoli and Makridis (2018)—individuals in occupations that involve greater IT-based tasks have continued experiencing rising wage premiums. All these pieces of empirical evidence point to the need for digital skills in the emerging labor market.

Policymakers may also address the concern that job losses from automation could disproportionately affect those who are least able to afford the tuition costs of reskilling programs up front, and those who are least likely to be able to sustain a forfeiture of labor income for the duration of the reskilling period. Gallipoli and Makridis (2018) find that individuals in jobs that tend to require more routine and manual skills are especially exposed to the growing demand for IT-based tasks. Another factor to consider in future policymaking is the unpredictable nature of disruption on the workforce. In determining federally funded programs to address displaced workers, the CEA (2018a, 21) cautions against programs targeting specific industries, instead suggesting that "keeping programs as flexible as possible reduces the need for continual re-optimization and increases the return on Federal dollars spent."

In addition to studying reskilling challenges, the Trump Administration has also established the President's National Council of the American Worker to develop and implement a strategy aimed at expanding educational attainment, training, and nontraditional degree programs that will prepare workers for the emergence of automation and AI (White House 2018a). Chapter 3 of this *Report* discusses the reskilling challenge in detail, including the job openings rates by industry.

The opportunity for reskilling is perhaps greatest in the field of cybersecurity, where there is a shortage of skilled workers (Burning Glass 2018). Figure 7-7, for example, uses 2018 data from CyberSeek (2018)—a partnership between Burning Glass Technologies, the Computing Technology Industry Association, and the National Initiative for Cybersecurity Education—to

Figure 7-7. Supply-and-Demand Ratio for Cybersecurity Jobs, 2018

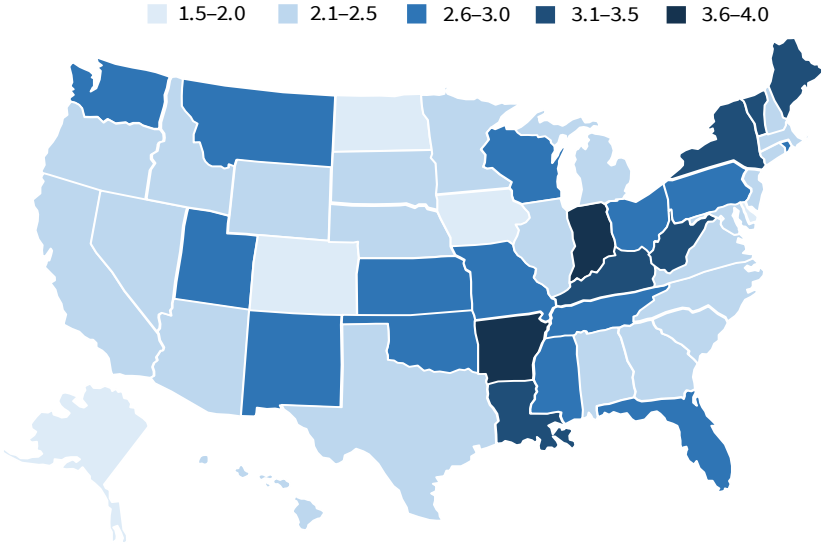

Source: CyberSeek (2018).

characterize the ratio of supply and demand for cybersecurity workers across locations (e.g., States). Although no State has a ratio less than 1, the vast cross-sectional heterogeneity highlights how different State labor markets face very different intensities of shortage (e.g., the District of Columbia has a value of 1.4, vs. Kentucky, which has a value of 3.2). To put these numbers in perspective, a value of 2 means that half of a State's existing cybersecurity workforce would need to change jobs every year to meet new postings, underscoring the amount of turnover that would be required to meet the skills gap.

The Administration's Policies to Promote Cybersecurity

It is essential that the Federal government and the private sector promote cyber best practices and cyber hygiene. For example, as discussed above, many Federal agencies have properly configured their email systems with DMARC. DHS's National Cybersecurity Assessments and Technical Services team determined that 71 of the 96 Federal agencies surveyed have cybersecurity programs that are either at risk or at high risk, for at least four reasons, according to OMB (2018a); in the next paragraph, we summarize these factors from the "Federal Cybersecurity Risk Determination Report and Action Plan" (White House 2018a).

Government agencies, along with the private sector, are not always aware of the situational context and/or the resources that exist to tackle the current threat environment. For example, 38 percent of the Federal cyber

incidents that were reported in 2018 did not specifically identify an attack vector. Organizations continue to adopt best practices, but there can be challenges with implementation. For example, only 49 percent of agencies have the ability to detect white-list software running on their systems.[17] Moreover, the lack of network visibility means that agencies may be unable to detect data exfiltration. For example, only 27 percent of agencies report that they have the ability to detect and investigate attempts to access large volumes of data. Finally, the lack of organizational and managerial policies surrounding the ownership of cybersecurity risk results in chief information officers or chief information security officers who lack the authority to make the relevant organization-wide decisions, but are nonetheless charged with the responsibility of maintaining network security. For example, only 16 percent of agencies achieved the government-wide target for encrypting inactive data.

These challenges are only going to grow, given the proliferation of data and increasing use of machine learning. Countries and malicious actors may turn toward counter-AI operations that attempt to alter and/or manipulate data (Weinbaum and Shanahan 2018). Individuals throughout the Federal civilian government, Department of Defense, intelligence community, and private sector will need to evolve to meet the expectations with identifying, protecting, detecting, responding, and recovering from threats in a timely manner. The Trump Administration—particularly through OMB, in partnership with the Department of Homeland Security, NIST, and the General Services Administration—is working to actively address these shortcomings. For example, the update to the TIC initiative is only one component of a broader effort by the Federal Chief Information Security Officer Council to obtain and test use-cases, particularly from the private sector (OMB 2018c). Moreover, as discussed in box 7-1, DARPA is developing new AI capabilities that help national security personnel more rapidly and reliably identify and address cybersecurity threats.

The Administration's Policies to Maintain American Leadership in Artificial Intelligence

The Trump Administration's AI agenda prioritizes advancing U.S. leadership in AI as well as helping the Nation's workforce adapt to the changes that are coming. As evidenced in the Administration's 2017 and 2018 budget priorities memoranda and highlighted at the White House AI summit in May 2018, the Administration continues to prioritize research-and-development funding for AI research and computing infrastructure, machine learning, and autonomous systems (OSTP 2018). To complement these active financial investments, the Administration also chartered the Select Committee on Artificial Intelligence under the National Science and Technology Council. This committee advises the White House on interagency research-and-development priorities, to foster

[17] An application white list refers to a set of applications that are authorized to be present according to a well-defined benchmark (Sedgewick, Souppaya, and Scarfone 2015).

collaboration between the private sector and academia, to identify opportunities to leverage Federal data and computational resources, and to improve the efficiency of government planning and coordination. The recent Executive Order on "Maintaining American Leadership in Artificial Intelligence" has formalized these commitments by calling for increased prioritization of investments, engaging in development of standards, and training and workforce development initiatives (White House 2019).

Second, the Administration has implemented policies that are conducive to more rapid economic growth and innovation by removing regulatory barriers, including those on the deployment of AI-powered technologies. In September 2017, the Department of Transportation released an update of the 2016 Federal Automated Vehicles Policy, providing nonregulatory guidance for AV developers, which was later further updated in October 2018 to provide a framework and multimodal approach to the safe integration of AVs into the surface transportation system. Similarly, the Administration is developing new rules in compliance with the Space Policy Directive–2 to streamline the licensing process for commercial space enterprises (White House 2018d). The Administration is also taking steps internationally to ensure that there is a level playing field for AI technologies. For example, at the World Trade Organization, and in trade agreements like the United States–Mexico–Canada Agreement, the Administration is protecting U.S. intellectual property and limiting the ability of foreign governments to require disclosure of proprietary computer source code and algorithms. These actions will better protect the competitiveness of our digital suppliers, and promoting access to government-generated public data, to enhance innovative use in commercial applications and services (USTR 2018).

Third, the Administration has begun integrating advances in AI and related technologies to improve the delivery of government services to the American people. The President's Management Agenda calls for the use of automation software to improve the efficiency of government services and maximize the applications of Federal data to help evaluate and modify Federal programs (OMB 2018b). In addition, in April 2017, the Department of Energy (DOE) and the Department of Veterans Affairs launched the Million Veteran Program Computational Health Analytics for Medical Precision to Improve Outcomes Now—known as CHAMPION—which uses high-performance computing infrastructure in the DOE National Laboratories to analyze large quantities of data and make recommendations that focus on suicide prevention and enhanced predictions and diagnoses of diseases (DOE 2017).

Recognizing that AI holds promise not only for greater economic opportunity but also for national security aims, the Trump Administration has directed considerable resources and leadership into targeted strategic investments, particularly at the nexus of AI and cybersecurity. One example, as discussed in box 7-1, is the Defense Advanced Research Projects Agency (DARPA 2018c),

which is actively investing in a "third wave" of AI technologies to make AI more transparent and accessible for deployment across both the public and private sectors. In particular, these initiatives focus on identifying ways for humans to use AI as tools for more effectively completing their tasks and maintaining network security.

To complement these broad-based research-and-development funding priorities, the Administration signed a memorandum directing, "Secretary of Education DeVos to place high quality STEM [science, technology, engineering, and mathematics] education, particularly Computer Science, at the forefront of the Department of Education's priorities" (White House 2017b). The Department of Education is working to devote over $200 million a year in grant funds toward these STEM and computer science activities, in addition to exploring other administrative actions that will advance computer science in K–12 and postsecondary institutions. Moreover, box 7-3 describes the emerging National Cyber Education Program, which is a prime example of an initiative focused on increasing the supply of STEM talent, specifically for the cybersecurity field.

The Administration's Implementation of the National Cyber Strategy

In addition to the National Security Strategy (White House 2017a), the Administration has also developed the comprehensive 2018 National Cyber Strategy, the first of its kind in over 15 years, to address the cybersecurity challenges of the coming decades (White House 2018b). This strategy's fourfold overarching goals mirror the pillars of the 2017 National Security Strategy; we paraphrase and synthesize these four objectives here, together with their priority areas.

The first objective is protecting the American people, the Homeland, and the American way of life. To do this, the Administration is securing Federal networks and information, securing critical infrastructure, and combating cybercrime and improving incident reporting. Three priorities associated with this objective involve improving risk management and incident reporting practices, modernizing Federal technology and security systems, and streamlining processes and roles and responsibilities.

The second objective is promoting American prosperity. To accomplish this, the Administration is fostering a vibrant and resilient digital economy, encouraging and protecting U.S. ingenuity, and developing a superior cybersecurity workforce. The priorities associated with this objective include promoting an agile and next-generation digital infrastructure, protecting intellectual property, and creating a pipeline and incentive structure that cultivate highly skilled cybersecurity and technology workers.

The third objective is to preserve peace through strength. To do this, the Administration is enhancing cyber stability through norms of responsible

Box 7-3. Educating the Cyber Workforce of Tomorrow

One of the most commonly cited workforce challenges within both the public and private sectors is the shortage of skilled workers. According to recent estimates from International Information System Security Certification Consortium—known as ISC^2—there is a shortage of 2.9 million cybersecurity employees globally (ISC^2 2018). Moreover, numerous survey results suggest that organizations are increasingly more likely to report a shortage of cybersecurity skills (Oltsik 2018; Burning Glass 2018).

Although there is debate about the its magnitude, there is a general recognition that more workers are needed to fill the increasing demand for cybersecurity skills, particularly as the paths by which hackers can gain access to computers and network servers expand in the growing digital economy. A national program that could help cultivate a new generation of cyber professionals prepared to meet the needs of the government, the defense community, and the private sector constitutes an Administration priority for both national security and the economy.

One example of a long-run and scalable solution is the National Cyber Education Program, which is a joint public–private initiative supported by the Trump Administration that seeks to inspire and educate children in elementary through high school about potential career paths and tools for careers in cybersecurity. This program is a multipart, public–private education initiative within the NIST Framework and with themes from the National Integrated Cyber Education Research Center at its core and strong support and leadership from a large educational services firm that serves 30 million K–12 students and 3 million teachers through its online education platform. This initiative includes these features:

1. Core curricular cyber content for grades K–12.
2. Virtual professional development for improving skills among STEM and cybersecurity educators to deliver content effectively and across disciplines.
3. Transformative learning tools and curricula for students to promote both technical content and real-world applications.
4. A career portal for connecting students with cybersecurity opportunities in government and the private sector, as well as regional conferences that provide access to counselors, educators, and industry professionals.
5. Tools for cybersecurity industry partners to engage their local communities, particularly schools, through volunteerism and mentorship.

The National Cyber Education Program has an estimated total budget of $20 to $25 million, which will be provided by a combination of committed private sponsors.

behavior and attributing and deterring unacceptable behavior in cyberspace. A priority related to this objective is countering malign cyber influence with information operations and better intelligence.

The fourth objective is to advance American influence. To accomplish this, the Administration is promoting an open, interoperable, reliable, and secure Internet and building international cyber capacity. Two priorities related to this objective include developing partnerships across the public and private sectors to promote innovation and cutting-edge technologies and promoting free and secure markets worldwide. As discussed in box 7-3, the National Cyber Education Program is an example of a public–private initiative that empowers teachers with the resources to improve learning outcomes and career pathways for students, particular for the emerging cyber workforce.

The Trump Administration is advancing these four objectives through a combination of short- and long-run efforts. In the long run, U.S. policymakers seek to prioritize an active and prepared pipeline of technology workers with mastery of information security practices. In the short run, the United States will continue strengthening network security, especially in critical infrastructure sectors. OMB issued a memorandum in May 2018 detailing the risk assessment process, which builds upon the Federal Information Security Modernization Act of 2014 Chief Information Officer metrics from 2017 and the Inspectors General metrics from 2016 (OMB 2018a). These metrics are based on the NIST Framework for Improving Critical Infrastructure Cybersecurity (NIST 2014), which provides best practices to which both public and private organizations can adhere, and aims to create predictability and encourage the adoption of best practices throughout government. Although no system in today's geopolitical environment is completely secure, these actions are setting the groundwork for a safe and secure digital infrastructure; see box 7-4 for a discussion of how Estonia became one of the world's leading countries in digital infrastructure.

Further Artificial Intelligence and Future of Work Policy Considerations

Motivated by the increasingly rapid pace of technological change and its implications for individuals, there are several lines of inquiry about the role of government.[18] First, some have suggested, as part of the social safety net, the

[18] We do not, however, discuss in depth the concerns about AI reaching a point of singularity, or general intelligence, whereby algorithms can create new ideas on their own without human assistance. Though the concept of singularity and the prospect of accelerated knowledge creation could lead to a large gain in productivity (Nordhaus 2015), an alternative scenario is one where algorithms would begin to dictate decisionmaking over human judgment. These discussions are beyond the scope of this chapter and the bulk of ongoing policy deliberations.

Box 7-4. Estonia: A Case Study of Modern Cybersecurity Practices

Although residents of Estonia rarely had access to electronic devices or the Internet a few decades ago, it has become an economic success story and digital leader in its region. Between 1995 and 2017, its real GDP grew by 141.5 percent (vs. 69.8 percent in the United States). According to the Estonian government, 99 percent of public services were available online as of 2017. Estonia does not use a centralized or master database, but rather X-Road—a software platform that allows links among its public and private e-service databases. According to the Estonian government, X-Road saves over 800 years of working time every year, reducing bureaucracy and raising efficiency (Vainsalu 2017).

Though Estonia "was, effectively, a disconnected society" in the early 2000s, moving toward a digital economy through the introduction of its X-road infrastructure has allowed the country to raise productivity and become more secure (Vassil 2015). Consider, for instance, queries involving vehicle registration data. Typically, this search would require three police officers working for about 20 minutes; but the X-Road software platform eases the retrieval of information, so a single officer can complete the search within a few seconds (Vassil 2015). All of Estonia's government services, ranging from collecting taxes to health records for personalized medical services, are made secure and readily accessible with the proper authentication credentials. These technological strides are arguably a major factor behind Estonia's emergence as one of the top countries for doing business, ranking as the most competitive tax system in the OECD, according to the Tax Foundation (2014), and as the seventh-most-free economy in the world, according to the Heritage Foundation (2018).

Interestingly, the number of queries through X-Road has grown exponentially, which is remarkable because similar digital services, such as data repositories and services, tend to grow linearly (Vassil 2015). An integral part of Estonia's success through X-Road has been its data security and privacy features. For example, citizens may use digital signatures, secured with a 2,048-bit encryption, to perform daily tasks such as banking and notarizing documents. Public safety has improved because the presence of digital identification cards has shortened response times to 10 seconds or less for 93 percent of emergency calls (Estonia 2018). In fact, as of 2018, the only legal transactions that one could not make online were marriage, divorces, and real estate. The core of these online activities is a 2000 digital signature law that created a framework for digital contracting.

Of course, the transition to a digital economy has come with increased targeting from other state and nonstate actors. Healthcare, energy, and the public sector face continuous cyberattacks, primarily from malware infections or outdated software. Perhaps Estonia's largest attack was in 2007; it involved distributed denial-of-service attacks that disabled computer networks, halting communication between the country's two largest banks

and causing reverberations for political parties. After the attack, Estonia established the NATO Cooperative Cyber Defense Center of Excellence in its capital, Tallin, in addition to founding the Cyber Defense League, which works to counter cyberattacks (Czosseck, Ottis, and Talihärm 2011). These increased security precautions and this institutional infrastructure have helped thwart attacks, including a large attempted attack on the country's digital identification cards, raising public confidence. The system is highly secure because access to databases via X-Road is gated via a secure identification card using two-factor authentication and end-to-end encryption (Estonia 2018).

Estonia has continued to prioritize improving its digital economy, in addition to developing a broader global network in partnership with other countries; see, for example, Estonia's "Digital Agenda 2020," which details plans to improve the well-being of its people and public administration through digitization (Estonia 2018).

provision of a universal basic income, which would help individuals potentially suffering from job displacement. Proponents argue, for example, that the scale of technological change is unlike anything developed countries have experienced in the past and, therefore, social safety nets must evolve to adapt to the new risks. However, a universal basic income would not only discourage work, especially in light of the existing social safety net (e.g., unemployment insurance and food stamps), but would also undermine the intrinsic value that work plays in creating meaning and purpose in peoples' lives (Opportunity America 2018).

Second, given the wide array of applications of AI for national security and warfare, there is an ongoing debate about whether AI should be regulated to prevent an "AI arms race" among countries (Taddeo 2018; Horowitz 2018). Particularly because AI is a general purpose technology (Agrawal, Gans, and Goldfarb 2018), the dual uses of AI developments mean that they will diffuse rapidly upon entering the private sector. One primary fear, for example, is that AI algorithms could make decisions about troop and/or drone deployments, which would put human lives at risk without the traditional human decision-making process. Much like the concerns about autonomous vehicles and passenger safety, some policymakers and researchers are calling for greater guidance on regulating AI when lives are at stake.

Third, although machine learning algorithms have been remarkably successful at predicting individual outcomes using increasingly accessible and granular data, many researchers and policymakers have voiced concern about the potential for these algorithms to propagate bias and discrimination (Kleinberg, Mullainathan, and Raghavan 2018). If the data on which algorithms are trained exhibit certain biases, then AI could propagate these biases on a wider and more subtle scale. Though these concerns are valid, the implications

for regulation are ambiguous. In particular, Kleinberg, Mullainathan, and Raghavan (2018) outline three conditions that are required for algorithmic fairness at the heart of these debates about algorithmic classification—showing that, except in special cases, no method can satisfy all three conditions simultaneously. In this sense, though concerns about algorithmic fairness ought to continue being voiced, policymakers should approach with caution when formulating policy to avoid simply reacting to the latest fad or worry.

Fourth, some are concerned that the emergence of big data and AI will pose a threat to competition because larger companies will be better equipped to train models on larger data (Seamans 2017; Bessen 2018). For example, companies with access to more data might be able to reduce business uncertainty by incorporating more information into their forecasts, thereby obtaining lower costs of capital (Begenau, Farboodi, and Veldkamp 2018). However, a countervailing force is the impact of AI on the cost of entry and creative destruction. For instance, the discovery and application of cloud computing allow firms to rent computer power and/or data storage. Aside from the 25 to 50 percent direct cost savings observed in government (West 2010), the indirect effects on entry costs and competition, particularly in concentrated markets, may be larger (Colciago and Etro 2013). Nonetheless, regulation and competition policy around big data and AI will remain an active ongoing debate.

Despite these general categories of concerns, caution is especially important when considering prospective regulation. For example, according to Stanford University's One Hundred Year Study of Artificial Intelligence, "The Study Panel's consensus is that attempts to regulate 'AI' in general would be misguided, because there is no clear definition of AI (it isn't any one thing), and the risks and considerations are very different in different domains" (Stanford University 2016). Moreover, because AI is an inherently global technology, regulation in one country could put companies that are competing in an international marketplace due to cross-country linkages at a significant disadvantage.

Conclusion

Recent advances in computer science and artificial intelligence technology are revolutionizing the U.S. economy. In many fields, tasks that traditionally required humans can now easily be performed by AI algorithms. Although these discoveries have the potential to "be as important and transformational to society and the economy as the steam engine," according to Brynjolfsson and McAfee (2014, 9), they are also creating known and unknown dependencies and challenges, such as accelerated polarization in the labor market and increased exposure to cybersecurity threats.

This chapter has defined and reviewed recent developments in AI and automation. Unlike traditional forms of information technology (e.g., computers) that require humans to provide instructions and programmatic

commands, intelligent systems are defined by their applicability to a wide range of tasks that need little supervision. For example, Google's new AI algorithm, AlphaZero, successfully trained itself how to play and subsequently defeat the world's best chess engine, Stockfish. Similarly, DARPA has also created tools capable of reliably and rapidly identifying cybersecurity vulnerabilities. Apart from these gaming and national security applications, AI is also frequently applied in the private sector—through, for example, data-driven decisionmaking business analytics and precision agriculture.

Drawing on historical examples, we have demonstrated the potential effects of AI technology on the U.S. labor market. Although advances in AI, and the introduction of technology more broadly, will inevitably change the composition of tasks and jobs by making some tasks typically performed by humans obsolete, we have shown in the text above that humans will continue to have an important economic function because of their comparative advantage over AI in other tasks, even if they do not hold an absolute advantage. This means that companies and entrepreneurs will find it more profitable to design technology capital that complements human capabilities. However, to alleviate the potentially adverse effects of AI on individuals and jobs that are more exposed to disruption, the Trump Administration has responded proactively by supporting and funding reskilling and apprenticeship initiatives in areas where humans retain a comparative advantage. For example, the Pledge to America's Workers, an initiative from the National Council for the American Worker, already has over 6.5 million pledges toward reskilling workers.

In addition, we have applied economic theory to analyze the wage patterns among industries that are adopting AI technology. In the initial anticipation phase, firms know that they will be more productive, but, because they currently lack the AI capital, raise real wages. However, in the arrival phase, which is typically the primary focus among the popular press, the introduction of AI substitutes for labor as workers compete with machines, thereby depressing real wages. But as business formation catches up with the new technology, real wages ultimately rise to levels above what they were before AI.

We have also explored ongoing cybersecurity vulnerabilities, along with future threats, as dependence on technology increases. The CEA (2018b) estimated the cost of attacks on these vulnerabilities to be $109 billion in 2016. Drawing on new data from Rapid7 across industries, we find that cybersecurity vulnerabilities are more pronounced than previously thought, even among well-established *Fortune* 500 firms. The prevalence of these vulnerabilities, coupled with the underreporting of public cybersecurity breaches, suggests that traditional measures of the cost of malicious cyberattacks may be greater than previously anticipated. We have discussed potential causes behind the failure to adopt cybersecurity best practices in the private sector, along with the policy implications, including tools already being used by the Federal government to prevent malicious cyberattacks and phishing attempts.

We conclude by highlighting the Trump Administration's current policy initiatives to tackle the risks posed by continued technological change in the labor market and new cybersecurity threats. The 2018 National Council for the American Worker, for example, has introduced initiatives to promote reskilling and apprenticeships to help workers transition into new and emerging jobs. For example, the Pledge to America's Workers already has over 6.5 million commitments to these aims by companies. In a similar vein, the 2018 National Cyber Strategy lays out a comprehensive framework for engaging and dealing with cybersecurity threats. For example, the "Federal Cybersecurity Risk Determination Report and Action Plan" (White House 2018a) establishes a detailed risk assessment process based on best practices from the NIST Framework to create predictability and the adoption of best practices throughout the Federal government. Moreover, by modernizing educational curricula and equipping teachers with new multimedia content and tools, the emerging National Cyber Education Partnership will help address the cybersecurity skills gap that currently threatens U.S. economic and national security.

The expansion of artificial intelligence and automation is already having profound effects on the U.S. economy and geopolitical landscape. Although we are only beginning to see their manifestations, and thus the full scale of potential threats and benefits cannot be entirely quantified, these changes pose both new challenges and opportunities. The Trump Administration is committed to policymaking that leverages technological change as an asset rather than a liability, to advancing economic gains for American workers, and to promoting best practices for our digital infrastructure so that America can remain the most prosperous and competitive country during the emerging technological transformation.