The Councilor: A Journal of the Social Studies

Volume 81 | Number 2 Article 3

October 2020

The Economics of Artificial Intelligence: A Primer for Social Studies Educators

Scott Wolla
Federal Reserve Bank of St Louis, scott.a.wolla@stls.frb.org

Cover Page Footnote

The views expressed in this paper are my own and do not necessarily reflect the views of the Federal Reserve Bank of St. Louis or the Federal Reserve System. Acknowledgements: The author thanks The Calvin K. Kazanjian Economics Foundation for initiating this project. The author also thanks Carolyn Backus, George Fortier, Lydia Johnson, Shaun Lysen, Michael MacDowell, John Siegfried, and Mary Suiter for their comments and suggestions.

Introduction

Although the modern expression "robot apocalypse" refers to a fear of technological advance, fear of automation is nothing new. For example, in 1589, Queen Elizabeth of England refused to grant the inventor of a mechanical knitting machine a patent, fearing it would put knitters out of work (Ip, 2017). In the early 19th century, English textile artisans called Luddites attempted to prevent the mechanization of the textile industry, fearing (correctly) that machines would replace labor in the industry. Even economists such as John Maynard Keynes worried about widespread technological unemployment "due to our discovery of means of economising the use of labour outrunning the pace at which we can find new uses for labour" (Keynes, 1933, p. 3). The newest wave of technological anxiety also includes fear of artificial intelligence (AI). A 2017 Pew Research survey found high levels of anxiety about automation and AI, with 72 percent of respondents expressing worry about a future where robots and computers can do many human jobs (Smith & Anderson, 2017). The issue—and the anxiety that surrounds it—is a common topic in the national news cycle. For example, a recent headline from *Business Insider* suggests that machines may replace half of human jobs (Thompson, 2016).

Of course, sensational headlines and exaggerated concerns can cloud an issue. But history offers plenty examples of disruptive technological change. Horses and mules once provided a significant part of the work within the economy. "Horse labor" seemed resistant to technological change. Even between 1840 and 1900, a time when the telegraph replaced the Pony Express and railroads replaced the stagecoach and Conestoga wagon, the number of horses and mules grew sixfold to more than 21 million. These animals played an essential role in farm work and an important role in carrying people and goods across the country in both rural and urban settings. Most Americans would have found it difficult to imagine an economy not dependent on horse labor. However, once the right technology came along, horse labor became obsolete. By 1960, the internal combustion engine provided most of the "horse power" in the U.S. economy and the population of horses had dwindled to just three million—a decline of nearly 88 percent (Brynjolfsson & McAfee, 2015). In just over half a century, horse labor had become uneconomical. Some fear that human labor faces a similar fate—that AI will replace it.

Economics, which is identified in the National Council for the Social Studies College, Career, and Civic Life (C3) Framework for Social Studies State Standards (NCSS, 2013) as one of the four core academic disciplines in social studies, provides a useful framework to analyze the effects of AI on society. First, AI provides social studies teachers with a current-issues "hook" that can be used to link valuable economic content to a relevant, contemporary issue that students care about. Second, AI provides these teachers with an issue that connects to key social justice issues such as economic opportunity and income inequality. Third, AI is the subject of many articles and lesson specifically designed for social studies teachers to use in the classroom. And finally, AI gives social studies teachers a chance to discuss key objectives such as optimal economic policy, college and career readiness, and investment in higher education. In short, AI is an important topic to address in the social studies classroom.

This paper provides teachers with the conceptual background required to address automation and AI within three of the four C3 social studies disciplines directly and to also include the fourth without much effort. It first addresses the current discussion of AI, which is rooted in the historical arguments that have been made since the beginning of the Industrial Revolution. It also provides context for understanding historical people's "Luddite" perspectives and describes how many of the same fears exist today relative to AI. This historical perspective is reflected in the C3 Framework, which requires students to "analyze how historical contexts shaped and continue to shape people's perspectives" (D2.His.5.9-12). Second, the implication that automation and AI may exacerbate inequality pushes societal discussion about it squarely into the civics realm. This concern aligns with the public policy proposals voiced by Elon Musk, Bill Gates, and Mark Zuckerberg, which are described and analyzed in this paper. These proposals connect well with the C3 Framework, which suggests that students "evaluate public policies in terms of intended and unintended outcomes, and related consequences" (D2.CIV.13.9-12). And, of course, many of the core concepts in the discussion about automation and AI—such as labor markets, income, productivity, economic growth, and standards of living—are found in economics and are defined and explained in this paper. While several of the economics C3 Framework standards apply, the most applicable one in this case requires students to "explain why advancements in technology and investments in capital goods and human capital increase economic growth and standards of living" (D2.Eco.13.9-12). And, while one step away from the content of this paper, the topic is conducive to addressing related issues in geography. For example, the C3 Framework requires students to "evaluate how political and economic decisions throughout time have influenced cultural and environmental characteristics of various places and regions" (D2.Geo.5.9-12).

Key Terms and Concepts

Scarcity lies at the heart of economic thought. It is the central economic problem and the basis of the economic way of thinking. Scarcity is the condition that exists because there are not enough resources to produce everyone's wants. Because resources are scarce, people must make choices. When they make choices, they face tradeoffs and incur opportunity costs.

Economists often describe the production process as a mix of labor and capital (both physical and human) resources (McConnell, Brue & Flynn, 2012). *Labor* is the quantity and quality of human effort directed toward producing goods and services. It is the work people do for income or, more specifically, wages. *Capital resources* (or *physical capital*) are the tools and equipment that have been produced and are used to produce other goods and services. They are used repeatedly in the production process. Capital resources are easily confused with technology. As economists define it, *technology* includes the knowledge, processes, and techniques used to produce goods and services. In other words, technology is all the intangible features embodied in the physical capital. Think of an iPhone: for a business, it is physical capital; but the difference between the original iPhone and the iPhone X is a difference in technology (Wolla, 2018). As the example suggests, technology changes. This idea is captured in the term *technological advance* (or *technological progress*), which is an advance in overall knowledge in a specific area. This insight—that output cannot simply be explained by increases in labor supply, human capital, and accumulation of physical capital alone—resulted in a Nobel Prize in Economic Sciences for Robert Solow, who summarized the idea well: "In the long term we know the only way to get sustained faster growth is to have sustained faster technological progress" (USB Nobel Perspectives, n.d., para. 6). Nobel Laureate economist Paul Romer (n.d.) explained technological progress as growth that springs from better recipes (processes), not just from more cooking (labor). *Human capital* is the knowledge and skills that people obtain through education, experience, and training. Increases in human capital can improve the productivity of labor resources.

Although the terms robot, automation, and AI are often used interchangeably, there are distinctions. A *robot* is any device or algorithm that does what humans once did, from thermostats to dishwashers to airline search sites. Unlike other physical capital or forms of technology, robots can be programmed to perform many tasks and do not need a human operator (Acemoglu & Restrepo, 2017). Robots are very good at doing repetitive tasks that can be programmed ahead of time. These kinds of routine and repetitive tasks are susceptible to *automation*, which is the automatically controlled operation of an apparatus, process, or system by mechanical or electronic devices. Once a process is automated, the production can occur without human assistance. Economists Jason Furman and Robert Seamans (2018) describe *AI* as "a loose term used to describe a range of advanced technologies that exhibit human-like intelligence, including machine learning, autonomous robotics and vehicles, computer vision, language processing, virtual agents and neural networks" (p, 2). The literature often describes *robots* and *AI* as variations of the same movement toward increasing *automation*, as will be seen in the remainder of this paper.

Underlying the discussion of labor, AI, and automation is the economic concept of *productivity*, which is the ratio of output per worker per unit of time. An increase in productivity enables a firm or an economy to produce the same amount of output with fewer inputs (or more output with the same level of inputs). For firms, increased productivity results in lower production costs and higher profits; when these effects appear for the overall economy, this translates to a higher standard of living.

Literature Review

Despite the hype and anxiety often expressed in the popular press, economists seem much less worried than the general population that humans will be replaced by computers and robots. In fact, many economists are likely to see AI as the latest chapter in a very long story of technological advance (Autor, 2015). When new technology is developed, it typically disrupts certain labor markets; but the economy adjusts in ways that result in more employment (Acemoglu & Restrepo, 2018).

In contrast, the popular perception and views expressed in the media often incite anxiety. David Autor (2015), a labor economist at the Massachusetts Institute of Technology and a leading researcher in the area of automation, says that journalists and even expert commentators tend to overstate the degree to which automation will substitute for labor. In doing so, they often ignore how labor often complements automation and how automation increases productivity, earnings, and the demand for labor. Economists Daron Acemoglu and Pascual Restrepo (2018) describe a false dichotomy that has taken hold in both the popular press and academic circles: "On the one side are the alarmist arguments that the oncoming advances in AI and robotics will spell the end of work by humans, while many economists on the other side claim that because technological breakthroughs in the past have eventually increased the demand for labor and wages, there is no reason to be concerned that this time will be any different" (p. 1). They suggest that the truth lies somewhere in the middle: humans will not be pushed to the sidelines, nor will AI have a trivial effect on the labor force.

Acemoglu and Restrepo (2018) built an economic model to explain their understanding of the relationship between capital and labor. This model provides a useful context for understanding the past and provides a framework for thinking about the future of technological advance. In this model, jobs are broken into tasks, which serve as the central unit of production, and these tasks can be completed either by labor or by capital. Individual tasks can be analyzed as better suited to labor or capital such that, in any given "job," some tasks might be automated while others are completed with human labor. In this way, productivity is enhanced; but human labor is still an essential part of the production process. As technology advances, it changes the mix of tasks completed by human labor and those completed by robots. The model identifies the following four processes at work:

The first process is the Displacement Effect. The model assumes that both labor and capital have comparative advantages for different tasks, which means that each (labor and capital) have varying opportunity costs. Labor is sometimes better suited for some tasks than capital and vice versa. In other words, the relative productivity of labor varies across tasks. When capital is less expensive (relative to labor) at the margin, the firm will substitute capital for labor for those tasks. The transition occurs as automation increases the number of tasks that can be cost-effectively completed using capital. On its own, the displacement effect results in a decrease in the demand for labor and a decrease in the equilibrium wage rate. Anxiety about automation has historically emphasized the displacement effect. However, the model suggests that three countervailing effects are at work.

The second process is the productivity effect. Investment in capital is intended to increase productivity. Increases in productivity decrease the cost of production and increase the demand for labor for non-automated tasks that complement the tasks completed by machines (Acemoglu & Restrepo, 2016). Historically, this can be seen with the advent of the ATM machine. Initially, some feared this technology would completely substitute the labor of bank tellers, making them obsolete. Rather, the number of bank tellers employed actually increased. ATM technology decreased the operating costs of bank branches to the point that banks found it advantageous to open more branches. With additional branches, they hired additional tellers. Thus, tellers per branch might have been reduced, but more branches overall increased employment (Bessen, 2015).

Although the number of bank tellers increased, tellers still needed to adjust their skills to meet the demands of the changing labor market. Rather than specializing in the tasks that ATMs had automated, tellers now specialized in relationship banking—forging “good will” with customers and selling them other bank services, which are difficult tasks to automate. Autor (2016) captures a similar idea in his economic framework. In his model, humans have a comparative advantage in tasks involving skills not easily replaced, including creativity, building relationships, and problem solving. For example, applying his framework to the ATM shows that ATM technology replaced some routine tasks of teller jobs, allowing tellers to take on more cognitively demanding and relational aspects of the banking business.

Higher productivity also has the potential to affect the demand for a broad selection of goods and services. Since higher productivity often translates into lower prices for goods and services, the productivity effect also leads to higher real incomes (i.e., income adjusted for inflation) for consumers. Higher real income results in an increase in the demand for all goods and services (including those not automated). As the demand for goods and services increases, the demand for the labor which produces them also rises. For example, the mechanization of agriculture reduced food prices, leaving households more income to spend on non-agricultural goods and services. The increased demand created employment opportunities for those displaced by mechanization (Herrendorf, Rogerson & Valentinyi, 2013). Overall, the productivity effect has been powerful in countering the displacement effect. The risk of automation and AI, then, is when either is just productive enough to bring the positive productivity gains described above (Acemoglu & Restrepo, 2018).

The third process is the capital accumulation effect. The displacement effect occurs when automation or AI substitutes for human labor. As technological progress continues, older capital is increasingly replaced by even more-productive capital. This process of "capital deepening" increases the productivity of tasks already automated. As such, additional labor is not displaced, because labor had already been replaced by capital. Instead, the deepening of automation further increases productivity, with the same effects as from the productivity effect: the demand for labor increases and real incomes for households increase as well, which result in greater demand for goods and services.

Acemoglu and Restrepo (2018) argue that up to this point (the sum of the displacement, productivity, and capital accumulation effects), automation results in a more capital intensive production process, which increases productivity more than wages. This increased productivity means that, on its own, the capital accumulation effect will cause the share of labor in national income to decrease over time. Historically, however, another force has pushed the production mix in the other direction—making production more labor intensive.

The final process is new tasks. Throughout history, the emergence of new jobs, industries, and tasks have complemented periods of automation. The mechanization of agriculture in the early 20th century resulted in a large increase in employment in factory work for new farm-related industries (Kuznets, 1966), for example, for a growing farm equipment industry (Olmstead & Rhode, 2001) and cotton milling (Rasmussen, 1982). Without the new demand for workers from these new industries, it would have been impossible to employ the millions of workers displaced from the agriculture sector. More recently, Acemoglu & Restrepo (2018) found that about half of the employment growth from 1980 to 2010 resulted from new tasks and job titles. Without the demand for new jobs in factory work (production, engineering, accounting, supervision, and management) in the latter half of the 19th and 20th centuries, it would have been nearly impossible to employ the millions of people exiting the agricultural sector and other labor-intensive jobs as automation replaced their jobs. Again, new tasks has the opposite effect of automation because it generates additional labor demand and increases the share of labor in national income.

Even though new types of jobs arise from automation's destruction of traditional jobs, the transition is not necessarily a smooth one. The risk is that, because of constraints, adjustments in the labor market will occur slowly and weaken any productivity gains that occur in the economy at large. One primary adjustment is the need for workers to acquire new skills, which requires the education system to keep up with the changing demands for skills (Brynjolfsson & McAfee, 2011). As more tasks are automated, the complementary tasks that human labor provides will change as well. As more automation occurs, any mismatch between skills and technologies is bound to complicate the adjustment process (Acemoglu & Restrepo, 2018). Unfortunately, there is little concrete information about what types of skills will complement the new technologies. Broadly, Autor (2015) forecasts that skills such as creativity, literacy, numeracy, adaptability, problem solving, and common sense will become increasingly important in coming decades. Just as the education system provided the workforce with the necessary skills to transition from an agricultural economy to an industrial economy, so again the education system will need to be an important part of a successful transition to the future economy.

Prediction: The Essence of AI

The year 1995 was pivotal: Microsoft released Windows 95, the U.S. government removed the final restrictions on commercial use of the internet, and Netscape celebrated its initial public offering (IPO). With these events, the internet moved from a technology primarily for academia and government to one for the economic mainstream. And frequent references to the New Economy (in 1995) suggested that a new economic framework was in play. Economists, however, realized that this was not a new economy per se, but rather one experiencing a technological change to which all the old economic models still applied. Using a standard economics framework, Agrawal, Gans, and Goldfarb (2018b) suggest that the internet was a general purpose technology that reduced the cost of distribution, communication, and search. And, in doing so, the internet changed nearly every industry over the next few decades.

In many ways, recent advances in AI technology remind people of 1995—that is, AI seems to be at a turning point. The recent advances in AI are a function of better computing power, better algorithms, big data, and advances in machine learning. Agrawal, Gans, and Goldfarb (2018a) argue that technologies labeled as AI can be thought of as prediction technologies. And, because prediction is a key input to decision making, AI has the potential to affect every decision. Advances in AI-powered prediction have decreased the cost of prediction. And, as the law of demand would suggest, as the cost of prediction falls, people will demand a greater quantity of it (i.e., they will find new ways to incorporate AI prediction).

Agrawal, Gans, and Goldfarb (2018b) define prediction as "the process of filling in missing information. Prediction takes the information you have, often called 'data' and uses it to generate information you don't have" (page 24). They provide simple, everyday examples: Netflix predicts which television shows and movies you might enjoy based on what data it collects about your's and other's behavior. Amazon similarly predicts which books, movies, music, and other products you might purchase. Amazon's Alexa, Apple's Siri, and Google Assistant predict what information you need when you ask them a question. Gmail predicts which words you intend to use in your message. Pandora predicts which new songs you might enjoy. Credit card companies predict which transactions are fraudulent. And Google Translate predicts what you are trying to communicate and expresses it in another language.

Advancements in AI prediction expand the range of what can be automated. Agrawal, Gans, and Goldfarb (2018a) use the self-driving car as an example of the difference that AI can make in programming. They explain that autonomous vehicles have been possible for some time but have largely run on tracks or pre-programmed routes in controlled environments. Putting an autonomous vehicles on a city street seemed improbable until recently because of the near endless "if-then" statements required to program a car to operate in an open (and much more complex) environment. However, recent improvements in AI prediction have turned the if-then problem into a prediction problem, allowing self-driving cars to emerge much earlier than people had projected just a few years ago. Rather than having to program if-then responses for every possible scenario, engineers now focus on a single prediction problem: What would a good human driver do? This prediction framework is then combined with machine learning. Machine learning is a category within the field of AI that intends to give machines the power to learn (Agrawal, Gans, & Goldfarb, 2019). In the case of driving, the autonomous car is equipped with cameras, radar, and lasers that act as the car's eyes and ears. In the training phase, as a human drives the car, the AI processes data about how the human acts and reacts in a variety of situations. The more the AI observes the human, the better it becomes at predicting the actions a human would take. As its predictions becomes better, performance improves. And, at some point, the AI becomes capable of driving the car without the human driver (Agrawal, Gans, & Goldfarb, 2018a).

Google Translate is a good example of how AI can change a process. The program, first offered in 2006 and built on "statistical machine translation," initially translated words one at a time, often leading to clumsy phrases and mistakes (Lewis-Kraus, 2016). Then in November 2016, the quality of its translations transformed nearly overnight when the platform switched to "neural machine translation"—a machine learning AI strategy. With this strategy, algorithms are used to discover patterns in the training data they are exposed to (De Jesus, 2017).

Google translate uses this strategy in the context of human communication. As the program is used, it identifies language patterns. It then takes what it learns and predicts the intended translation from one language to another. The program receives reinforcement as it collects more data. The more data it collects, the more it learns, and the better it becomes (Agrawal, Gans, & Goldfarb, 2018b).

Prediction is valuable because it is an input into decision making, and decision making is everywhere. But prediction is not decision making in itself. Other decision-making inputs include data, actions, and judgment. Judgment is the process of determining what the reward is to a particular action in a particular environment, or knowing which predictions to make and what to do with them. At this point, judgement is still a skill that humans contribute to the process. In short, prediction and judgement are complements. From an economic perspective, as prediction becomes cheaper, human decision making (judgement) becomes more important. For example, consider radiology. AI prediction technology is already allowing machines to detect tumors better than humans can because of its ability to recognize patterns. Will radiologists be replaced by machines? Similar to the way the ATM changed the jobs of bank tellers, AI will likely change the jobs of radiologists (and many other professions). In the future, radiologists will spend much less time looking at images. Even today, radiologists spend much of their time working with primary care doctors, explaining the results of tests, and helping them decide the best course of action for patients. Over time, this consulting role will grow more important as the role of examining images is increasingly performed by an AI prediction machine. Further, as the cost of AI imaging falls, the amount of imaging used in the diagnosis process will likely increase. The increase in volume can be expected to offset some of the decline in image analysis (Agrawal, Gans, & Goldfarb, 2018b).

In summary, some key insights can be drawn from the economics of prediction: First, as prediction becomes less costly, we use it more. Second, when the cost of AI prediction is low, it decreases the value of its substitute (human prediction). Third, when the cost of AI prediction is low, it increases the value of it complements (such as data and human judgement).

How Will AI and Robotics Affect Workers?

It is clear that the use of robots has increased dramatically. In the United States there were 0.49 robots per thousand workers in 1995. By 2017, there were 1.79 robots per thousand workers (Bharadwaj & Dvorkin, 2019)—a trend likely to continue. McKinsey Global Institute studied tasks within more than 800 occupations to determine the percentage of a job that could be automated using current technology. This approach is similar to the framework that Acemoglu and Restrepo (2018) developed because it recognizes that every job is made up of many tasks with differing potential for automation. The study found that less than 5 percent of occupations are candidates for full automation. However, roughly 60 percent of occupations could have 30 percent or more of their tasks automated. Because these expected changes will occur at the task level (not the job level), they will result in significant job redefinition. For example, diagnosis of many health issues could be effectively automated so that diagnosis and triage could be combined in emergency rooms; this would allow doctors to focus on the most acute or uncommon cases. Additionally, automation would allow mortgage lenders to spend less time processing paperwork, creating more time to review the exceptions and meet with clients. Both of these examples free the expert to focus on higher-value work (McKinsey and Company, 2017). The bottom line is that more occupations will change than be automated away; but there will be growing pains as displaced workers retrain and existing workers learn to use new technology. Inevitably, continued automation will change the workplace and the types of labor employers demand.

In the past, labor-intensive, low-wage jobs were in the most danger of being replaced by automation. Economists suggest that AI will impact future workers differently: the labor market will expand at the high and low ends of the job-skills spectrum, and the greatest substitution will occur at the middle of the job-skills spectrum, a phenomenon called employment polarization (Bharadwaj & Dvorkin, 2019).

On one end of the spectrum are high-income, cognitive jobs that require creative problem solving and tasks complemented by technology. Job opportunities will continue to occur in traditionally high-income sectors such as healthcare, finance, and law (with some automation occurring in each). But Autor and Salomons (2019) also identify a growing set of occupations they label “frontier jobs” that involve the production, installation, and maintenance of new technologies such as robot integration, search engine optimization, and wind turbine maintenance.

The other end of the spectrum, which is also growing, are jobs in low-income manual and service occupations that rely on dexterity, interpersonal communication, and physical proximity to the job; these tasks are difficult to automate (Autor & Dorn, 2013). For these occupations, which include home health aides, landscapers, and maintenance workers, only a very small percentage of activities could be automated. The McKinsey Global Institute expects that the demographics of an aging population combined with the individual care they will likely need will result in a large increase in jobs as registered nurses, nursing assistants, personal care aids, and home health aides from 2016-2030 (McKinsey and Company, 2017). The institute also forecasts increased demand for construction workers as infrastructure is modernized and housing patterns change and for teachers, administrators, and teaching assistants as the demand for education increases. Autor and Salomons (2019) identify a category of growing occupations they label "wealth workers" that provide labor-intensive, in-person services to affluent consumers in high-wage, urban labor markets. These occupations include those for yoga instruction, sommelier services, pet care, personal training, and counseling and are neither technologically advanced nor highly paid.

Middle-income jobs are at the greatest risk from technological advance (Bharadwaj & Dvorkin, 2019). While previous waves of automation were confined largely to industry and agriculture, AI threatens middle-income jobs because they often follow well-understood rules and procedures increasingly easily codified in software and executed by computers (Autor, 2016). Computer software has already automated several middle-income tasks in retail, wholesale, and business services. AI-powered technologies can "retrieve information, coordinate logistics, handle inventories, prepare taxes, provide financial services, translate complex documents, write business reports, prepare legal briefs, and diagnose diseases" (p. 4, Acemoglu & Restrepo, 2018). AI-powered technologies are also set to become much better at these tasks during the next several years (Brynjolfsson & McAfee, 2011). McKinsey (2014) forecasts office support positions (such as payroll clerks and data entry) to be among the hardest hit. The shrinking of the middle class from employment polarization could result in a more stratified society. Thus, as the number and nature of the tasks that AI can perform increases (and threatens tasks previously done by well-educated professionals), so will anxiety about the long-term impact of AI on the economy and society in general.

Predicting future employment is difficult because it depends on accurately forecasting the rate of technological progress. As it turns out, previous predictions have not been very accurate. For example, in "Why People Still Matter," Frank Levy and Richard Murnane (2004) argued that replicating human perception would be difficult to automate; thus, driving in traffic would be insusceptible to automation. Six years later, Google announced that it had modified several Toyota Priuses to be fully autonomous (Brynjolfsson & McAfee, 2011). In 2003, David Autor, Frank Levy, and Richard Murnane (2003) identified tasks they believed were and were not susceptible to automation, with legal writing and truck driving among those insusceptible. Current analysis suggests that advances in AI will soon automate both legal writing and truck driving (Frey & Osborne, 2017).

Policy Responses to Automation

The two main groups providing potential policy responses to automation and AI are those who fear the "robot apocalypse" and those who believe automation and AI will generate more jobs than they destroy in the long run. Some who fear the "robot apocalypse" often argue that scarcity will cease to be a problem; in other words, they believe that although the demand for labor will be greatly reduced as robots do all the work, goods and services will be plentiful. In such a case, the central economic problem will be the distribution of the goods and services produced among the people (Autor, 2015). In essence, society will need to address the allocation problem of who will get to consume the new abundance of goods and services.

In our current economic system, jobs are plentiful and differences in income largely determine who can buy goods and services. This system provides both a strong incentive to work and an efficient system (albeit unequal) for allocating goods and services. In a world where machines have replaced human labor, income will be largely directed to the owners of the companies who utilize automation and AI. This suggests that much more income will flow into the accounts of fewer, wealthier individuals. Some fear that today's workers will have less of an opportunity to earn income in the future as more and more human jobs are replaced by machine labor. In this framework, inequality has the potential to rise to a greater (perhaps staggering) level and require resources to be allocated differently.

Mark Zuckerberg (Facebook) and Elon Musk (Tesla) suggest that the benefits of automation and AI be used to fund continuous education and universal basic income (UBI) (Sodha, 2017; Weller, 2017). UBI is unconditional income paid by the government to all citizens, regardless of whether they are working or unemployed, wealthy or impoverished. Proponents suggest this reallocation would provide basic and sustainable consumption for households and could replace the various welfare and assistance programs run by the state and federal governments.

Taxation would also become a central economic problem in a post-scarcity world with few human workers. In 2017, individual income taxes generated about 48 percent of tax revenue (Office of Management and Budget). If human labor becomes largely automated and many people no longer earn income through labor, tax revenue would need to be generated in some other way to sustain the functions of government. Bill Gates suggests that the government should tax the work done by robots to compensate the workers they replace. Gates says that the tax could be used to finance jobs taking care of the elderly and working with children in schools, services for which human labor is well-suited. Gates also suggests the robot tax would be useful to slow the speed of automation, which Gates sees as an added benefit (Delany, 2017).

Economist Larry Summers (2017) disagrees with Bill Gates. He does not believe AI will eliminate scarcity. Further, he says that because robots increase productivity, they increase our standard of living and create wealth for society. As such, taxing robots would reduce these benefits, which would be counterproductive. Summers also suggests that taxing robots will likely drive production off-shore, further hurting American workers. Rather than taxing robots, Summers suggests a larger role for government to counter the problems with structural joblessness that will likely result. He suggests reforms in education and training systems, wage subsidies for people with severe employment problems, and major investments in infrastructure to further boost productivity and provide employment opportunities.

Like Summers, most economists suggests that scarcity will continue to be the basic economic problem and the new tasks that emerge from new industries will generate enough jobs to ensure full employment in the long run. They suggest that policy should focus on an education system that provides workers with the skills they need to engage in a growing and changing economy.

Education: The Importance of Increasing Human Capital in the Face of Automation

Technological advance changes the mix of tasks in the production process. Because it changes the production process, it also changes the skills demanded by employers. Difficulties adapting to these changing demands will cause some workers to be displaced or forced to retrain to keep their positions (Autor, 2015). Education has historically played a major role in mitigating the adverse effects of labor market shifts in the wake of technological advance. For example, in the United States, the large supply of labor displaced by the innovations in agriculture in the early 20th century coincided with the "high school movement," which increased enrollment in U.S. high schools from 18 percent to 71 percent. This dramatic increase in the human capital of the American workforce helped smooth the transition from agriculture to industry (Goldin & Katz, 2008).

Future job growth lies in the tasks that will be developed as new industries and tasks emerge. Economists suggest that the jobs of the future will bundle human skills and judgment with technological skills. As such, Andrew McAfee, co-director of the MIT Initiative on the Digital Economy, suggests that students pursue a double major: one in liberal arts to develop problem-solving, creativity, and critical-thinking skills and another in a STEM (science, technology, engineering, and mathematics) area to develop quantitative and technological skills (Regalado, 2012). While this option is not possible for all students, it is clear that workers must increasingly acquire the skills necessary to ensure that technology is a complement rather than a substitute for their human capital. In the future, education will more likely not end with a high school or post-secondary education; employability will likely necessitate constantly upgrading skills and education as technology changes.

For educators, many have feared that education technology might someday substitute for classroom teachers. Acemoglu, Laibson, and List (2014) use the economic theory of comparative advantage to argue that the Internet will not substitute for teachers, but rather lead to specialization in teaching tasks. They argue that the Internet is more likely to replace the lecture and information transmission aspect of education and increase the demand for teachers who specialize in face-to-face student interaction to lead discussions and active learning. They suggest that online lectures and content will give all students access to "superstar" instructors, which will reduce education inequalities for those in underserved communities.

Of course, the challenge is to provide an education system that supplies the skills and training *now* for jobs and tasks that will exist in the future—but are currently unknown. If the education system is not able to provide workers with the necessary skills (problem-solving and critical-thinking skills as well as quantitative and technical training), adjustment to the new economic conditions will be impeded. Any mismatch between skills and technologies will slow the adjustment of employment and wages (Acemoglu & Restrepo, 2018). Slow adjustment will impact the welfare of displaced workers, but it will also impact the potential productivity enhancements to the economy more broadly. In short, if certain skills are complementary to new technologies but those skills are lacking in the labor force, the productivity gains will be smaller than what they could have been. The danger is that productivity-enhancing automation replaces labor but does not increase productivity enough to increase the demand for the non-automated tasks that complement the automation (Acemoglu & Restrepo, 2018).

Social Studies Education and Automation and AI

It is essential that people see beyond the fear of the "robot apocalypse" and make rational decisions about their human capital. In order for students, parents, teachers, and policymakers to make good economic decisions about the future, they must have a solid foundation of economic knowledge and the ability to use the economic way of thinking—a rational, decision-making model built on a foundation of scarcity, tradeoffs, and opportunity cost. Upon that foundation, teachers can build a framework more specific to automation and AI, such as knowledge about labor markets, productivity, and the roles of human and physical capital. While these tend to be seen as economics-specific concepts, they can be introduced into any of the social studies disciplines. In fact, these concepts are core to the C3 Social Studies Standards, the National Voluntary Content Standards in Economics, and many state standards. For example, automation and AI issues tie directly to the C3 Social Studies Framework by addressing topics central to understanding the national economy and requiring high school students to "explain why advancements in technology and investments in capital goods and human capital increase economic growth and standards of living" (D2.Eco.13.9-12). The issues also align with the same dimension at lower grade levels (D2.Eco.13.K-2, D2.Eco.13.3-5, and D2.Eco.13.6-8) as well as the C3 Framework in civics, history, and geography. In addition, automation and AI issues and the core concepts described above also align with several of the Voluntary National Content Standards in Economics, including #1 (Scarcity), #2 (Decision Making), #4 (Incentives), #6 (Specialization), #13 (Income), and #15 (Economic Growth).

A variety of resources are available now: The National Council for the Social Studies offers ready-made classroom resources that align with the economic content. Day (2018) provides a lesson that applies the Inquiry Design Model (IDM) and C3 Framework to the topic by asking the inquiry question "Should We Tax Robots?" He includes three supporting questions that align well with the content in this article: 1) What is the effect of technology in the workplace?, 2) What is the effect of technology at home?, and 3) What is the effect of technology on employment in the U.S. economy? The lesson uses a variety of sources for students to investigate the topic, including a picture of a factory in the 1960s and comparison picture of a factory in the 2010s, a transcript from an interview with Bill Gates (who comments on the issue of taxing robots), U.S manufacturing data, U.S. unemployment data, and interview data gathered by students. To provide students an overview, *Social Education* published "The Economics of Artificial Intelligence and Robotics" (Wolla, Schug, & Wood, 2019).

The concepts identified in this present paper can be built into any social studies course. Proper scaffolding might involve starting with foundational concepts such as scarcity and opportunity cost and then moving to economic systems, labor markets, productivity, and investment in human capital. Free teaching resources on these topics are widely available from reliable sources such as the Council for Economic Education, the U.S. Bureau of Labor Statistics, and Federal Reserve Banks of Dallas, Richmond, and St. Louis (Table 1).

Table 1: Teaching Resources

Key Concept: Scarcity
Opportunity Cost (online learning module) This course introduces opportunity cost, a fundamental concept in economics. An understanding of opportunity cost contributes to informed decision making at the personal and societal levels. **Source: Federal Reserve Bank of St. Louis** https://www.econlowdown.org/resource-gallery/opportunity_cost
There is No Such Thing as a Free Lunch (video) The study of economics is built on the foundation of three very important concepts: scarcity, choice, and opportunity cost. This episode of the Economic Lowdown Video Series explains why there is no such thing as a free lunch. **Source: Federal Reserve Bank of St. Louis** https://www.stlouisfed.org/education/economic-lowdown-video-series/no-such-thing-free-lunch

Key Concept: Economic Systems
Circular Flow (video) This episode of the Economic Lowdown Video Series explains the circular flow model. Viewers will learn how households and businesses interact in the market for resources and in the market for goods and services and see how money keeps the whole process moving. **Source: Federal Reserve Bank of St. Louis** https://www.stlouisfed.org/education/economic-lowdown-video-series/episode-6-circular-flow
Factors of Production (podcast) Factors of production are the resources people use to produce goods and services; they are the building blocks of the economy. Economists divide the factors of production into four categories: land, labor, capital, and entrepreneurship. **Source: Federal Reserve Bank of St. Louis** https://www.stlouisfed.org/education/economic-lowdown-podcast-series/episode-2-factors-of-production
Key Concept: Labor Markets
Who Decides Wage Rates? (lesson) In this lesson, students play the role of either buyers or sellers of labor to examine the interconnectedness of individuals and companies in labor markets. Students learn that the demand and supply for labor determine market wage rates and that wages depend, in part, on individual productivity. **Source: *High School Economics* (lesson 13), Council for Economic Education** http://hseconomics.councilforeconed.org/
The Labor Market (video) This episode of the Economic Lowdown Video Series explains the basics of the labor market. Viewers will learn how the laws of supply and demand determine the wages and quantities of labor employed in various labor markets. **Source: Federal Reserve Bank of St. Louis** https://www.stlouisfed.org/education/economic-lowdown-video-series/episode-4-the-labor-market
Key Concept: Productivity
Will Robots Take Our Jobs? (article) Robots are in the headlines, and many of us wonder if they'll also take over our jobs. Is the "Robot Apocalypse" upon us, or is this part of a larger trend that's been occurring for much of human history? This issue of *Page One Economics*® discusses that and more. **Source: Federal Reserve Bank of St. Louis** https://www.stlouisfed.org/education/page-one-economics-classroom-edition/will-robots-take-our-jobs
What is Productivity? (video) Explore the meaning of productivity with the Bureau of Labor Statistics and learn how productivity growth can lead to improvements in our lives and the well-being of our nation. **Source: U.S. Bureau of Labor Statistics** https://www.bls.gov/video/?video=mRxICdUYaCs

Key Concept: Investment in Human Capital
Invest in What's Next: Life After High School (online module) This online course helps high school students evaluate their choices for their first major financial decision—what path to pursue after high school. The three online lessons include Exploring My Options, Budgeting for My Future, and Evaluating My Plan. **Source: Federal Reserve Banks of Richmond and San Francisco** https://www.investinwhatsnext.org/
Navigate: Exploring College and Careers After High School (workbook) This resource is an introduction to the investigation of careers and college and is designed for 7th through 9th grade students. Many students and their families may not be familiar with the pathways to education after high school, and Navigate provides information to begin preparing for success. **Source: Federal Reserve Bank of Dallas** https://www.dallasfed.org/educate/navigate.aspx

Conclusion

Many people suffer from anxiety about technology and its effects on the workforce; they imagine a world where human labor has been displaced by robots and the *distribution* of goods and services (not scarcity) is the central problem. While the general population has concerns about how automation and AI will affect the economy, economists are generally optimistic about their effects (Autor, 2015). Productivity-enhancing technology has changed the economy in dramatic ways over the past two centuries, but it has not made human labor obsolete or eliminated the problem of scarcity. In 1966, another period of automation anxiety, economist, computer scientist, and Nobel Laureate Herbert Simon wrote, "Insofar as there are economic problems at all, the world's problems in this generation and the next are problems of scarcity, not of intolerable abundance. The bogeyman of automation consumes worrying capacity that should be saved for real problems" (Simon, 1966). This statement is still true today.

Many economists see the current wave of new technology and automation as a trend that has been occurring for most of human history and one that will continue in the future (Autor, 2015). The challenge in their view it to equip workers with the skills needed in the future to complement emerging technology—to avoid being substituted by it (Acemoglu & Restrepo, 2018). This challenge has implications not only for the choices of students and parents, but also for the decisions of schools and governments as they choose how much to invest, what parts of the curriculum to invest in, and what academic standards will be required and tested.

Automation and AI provide social studies educators with a topic that is current and relevant, has historical context, and provides opportunities to connect to ideas embedded throughout the social studies curriculum. Social studies educators can contribute by ensuring students have the understanding they need to make sound decisions about their economic futures. Overall, while technological progress creates uncertainty, education has the potential to mitigate many of the adverse effects of labor market disruptions and alleviate the fear surrounding the growth of automation and AI.

References

Acemoglu, D., Laibson, D., & List, J. A. (2014). Equalizing superstars: The internet and the democratization of education. *American Economic Review*, 104(5), 523-27.

Acemoglu, D., & Restrepo, P. (2018). Artificial intelligence, automation and work. *National Bureau of Economic Research*, NBER Working Paper, No. 24196.

Acemoglu, D., & Restrepo. (2017). Robots and jobs: Evidence from US labor markets. *National Bureau of Economic Research*, NBER Working Paper, No. 23285.

Acemoglu, D., & Restrepo, P. (2016). The race between machine and man: Implications of technology for growth, factor shares and employment. *National Bureau of Economic Research*, NBER Working Paper, No. 22252.

Agrawal, A., Gans, J., & Goldfarb, A. (2018a). Prediction, judgment and complexity. *National Bureau of Economic Research*, NBER Working Paper, No. w24243.

Agrawal, A., Gans, J., & Goldfarb, A. (2018b). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Press.

Agrawal, A., Gans, J., & Goldfarb, A. (2019). Economic policy for artificial intelligence. *Innovation Policy and the Economy*, 19(1), 139-159.

Autor, D. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3-30.

Autor, D., & Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review*, 103(5), 1553-97.

Autor, D. H. (2016). Will automation take away all our jobs? [video file]. Retrieved from https://www.ted.com/talks/david_autor_why_are_there_still_so_many_jobs?language=e.

Autor, D., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, 118(4), 1279-1333.

Autor, D., & Salomons, A. (2019). New frontiers: The evolving content and geography of new work in the 20th Century. *NBER Conference*, Toronto. Retrieved from https://conference.nber.org/conf_papers/f129906.pdf.

Bessen, J. (2015). Toil and technology. *Finance and Development*, 52(1).

Bharadwaj, A., & Dvorkin, M. A. (2019). The rise of automation: How robots may impact the U.S. labor market. *The Regional Economist*, Federal Reserve Bank of St. Louis, Vol. 27, No. 2. Retreived from https://www.stlouisfed.org/publications/regional-economist/second-quarter-2019/rise-automation-robots.

Brown, P. C., Roediger III, H. L., & McDaniel, M. A. (2014). *Make it stick: The science of successful learning.* Harvard University Press, Cambridge, MA.

Brynjolfsson, E., & McAfee, A., (2011). *Race against the machine: How the digital revolution is accelerating innovation, driving, productivity, and irreversibly transforming employment and the economy*. Digital Frontier Press, Lexington, MA.

Brynjolfsson, E., & McAfee, A. (2015). Will humans go the way of horses? Labor in the second machine age. *Foreign Affairs*, Federal Reserve Bank of Philadelphia. Retrieved from https://www.philadelphiafed.org/-/media/research-and-data/events/2015/fed-policy-forum/papers/brynjolfsson-humans_horses.pdf?la=en.

Council for Economic Education. (2010). *Voluntary National Content Standards in Economics*, 2nd ed. Council for Economic Education, New York, NY.

Retrieved from http://councilforeconed.org/resource/voluntary-national-content-standards-in-economics/.

Day, S. H. (2018). Should we tax robots? In Swan K., John L. K., & S.G. G. (Eds.), *Teaching the College, Career, and Civic Life (C3) Framework, Part Two* (pp.177-186). Silver Spring, Maryland: National Council for the Social Studies.

De Jesus, C. (2017, January). Artificial intelligence: What it is and how it really works. *Futurism*. Retrieved from https://futurism.com/1-evergreen-making-sense-of-terms-deep-learning-machine-learning-and-ai

Delany, K. (2017, February 17), "The robot that takes your job should pay taxes, says Bill Gates", *Quartz.* Retrieved from https://qz.com/911968/bill-gates-the-robot-that-takes-your-job-should-pay-taxes/.

Frey, C. B., & Osborne, M. A. (2017), "The future of employment: How susceptible are jobs to computerization?" *Technological Forecasting and Social Change*, 114, 254-280.

Furman, J., & Seamans, R. (2018). AI and the economy. *National Bureau of Economic Research,* NBER Working Paper, No. 24689.

Goldin, C., & Katz, L. F. (2008). Mass secondary schooling and the state: The role of state compulsion in the high school movement. In D.L. Costa & N.R. Lamoreaux (Eds.), *Understanding long-run economic growth: Geography, institutions, and the knowledge economy* (pp. 275-310). University of Chicago Press.

Herrendorf, B., Rogerson, R., & Valentinyi, A. (2013). Two perspectives on preferences and structural transformation. *American Economic Review*, 103(7): 2752–2789.

Ip, G. (2017). Workers: Fear not the robot apocalypse. *Wall Street Journal*, September 5. Retrieved from https://www.wsj.com/articles/workers-fear-not-the-robot-apocalypse-1504631505.

Keynes, J.M. (1933). Economic possibilities for our grandchildren. *Essays in Persuasion*, 358-373.

Kuznets, S. (1966). *Modern economic growth*. New Haven: Yale University Press.

Olmstead, A., & Rhode, P. (2001). Reshaping the landscape: The impact and diffusion of the tractor in American agriculture, 1910-1960. *The Journal of Economic History,* 61(3): 663–98.

Levy, F., & Murnane, R. J. (2004). *The new division of labor: How computers are creating the next job market*. Princeton University Press.

Lewis-Kraus, G. (December 14, 2016). The great A.I. awaking. *The New York Times Magazine*. Retrieved from https://www.nytimes.com/2016/12/14/magazine/the-great-ai-awakening.html.

McConnell, C. R., Brue, S. L., & Flynn, S. M. (2012). *Economics: Principles, problems, and policies, 19th edition*. McGraw-Hill Irwin.

McKinsey and Company. (2017). *Jobs lost, jobs gained: Workforce transitions in a time of automation*. Retrieved from https://www.mckinsey.com/mgi/overview/2017-in-review/automation-and-the-future-of-work/jobs-lost-jobs-gained-workforce-transitions-in-a-time-of-automation.

National Council for the Social Studies. (2013). *The college, career, and civic life (C3) framework for social studies state standards: Guidance for enhancing the rigor of K-12 civics, economics, geography, and history*. Silver Spring, MD: NCSS.

Office of Management and Budget (OMB), *Table 2.2—Percentage composition of receipts by source: 1934–2023*. Retrieved from https://www.whitehouse.gov/omb/historical-tables/.

Rasmussen, W. (1982). The mechanization of agriculture. *Scientific American*, 247(3): 76–89.

Regalado, A. (2012). When machines do your job. *MIT Technology Review*, July 11. Retrieved from https://www.technologyreview.com/s/428429/when-machines-do-your-job/.

Romer, P. (n.d.). Economic growth. *The Library of Economics and Liberty.* Retrieved from http://www.econlib.org/library/Enc/EconomicGrowth.html.

Simon, H. A. (1966). Automation: A letter in response to 'Where Do We Go from Here?' *New York Review of Books*, March 17, 1966.

Smith, A., & Anderson, M. (2017). *Automation in everyday life*. Pew Research Center. Retrieved from http://www.pewinternet.org/2017/10/04/automation-in-everyday-life/.

Sodha, S. (2017). Mark Zuckerberg's got some cheek, advocating a universal basic income. *The Guardian.* Retrieved from https://www.theguardian.com/commentisfree/2017/jul/10/mark-zuckerberg-universal-basic-income-facebook-tax,

Summers, L. (2017). Robots are wealth creators and taxing them is illogical. *Financial Times, March 5.* Retrieved from https://www.ft.com/content/42ab292a-000d-11e7-8d8e-a5e3738f9ae4.

Thompson, C. (2016). Machines may replace half of human jobs. *Business Insider*, February 16. Retrieved from http://www.businessinsider.com/machines-may-replace-half-of-human-jobs-2016-2.

USB Nobel Perspectives (n.d.). *Why do some economies grow faster than others?* Retrieved from https://www.ubs.com/microsites/nobel-perspectives/en/laureates/robert-solow.html.

Weller, C. (2017). Elon Musk doubles down on universal basic income: 'It's going to be necessary.' *Business Insider,* February 13. Retrieved from http://www.businessinsider.com/elon-musk-universal-basic-income-2017-2?r=UK&IR=T.

Wolla, S. (2018). Will robots take our jobs? *Page One Economics*®, Federal Reserve Bank of St. Louis. Retrieved from https://research.stlouisfed.org/publications/page1-econ/2018/01/02/will-robots-take-our-jobs.

Wolla, S., Schug, M., & Wood, W. (2019). The Economics of artificial intelligence and robotics. *Social Education*, National Council for the Social Studies, 83(2), pp 83-87.