Susan Athey

The Economics of Technology Professor, Stanford Graduate School of Business

Written Testimony

House Committee on the Budget

Hearing on Machines, Artificial Intelligence, & the Workforce: Recovering & Readying Our Economy for the Future

September 10, 2020

Chairman Yarmuth, Ranking Member Womack and Members of the Committee,

      Thank you for inviting me to speak today. Artificial intelligence (AI) seems to inspire extreme views among policy analysts: some focus on a future where robots take all the jobs, while others argue that its effect will be no different than previous rounds of technological innovation. My own view is that AI has enormous positive potential that should not be ignored, and that governments and universities have a crucial role to play in ensuring that the potential is realized. At the same time, it will create some challenges, contributing to an era where workers transition more frequently and require more ongoing training and reskilling throughout their careers. If we invest carefully in making technology and innovation part of the solution, AI can achieve its potential and even contribute to addressing a variety of societal and fiscal challenges that loom before us.

      An important precursor to a policy discussion about AI is a grounding in the technology, including a clear framework for understanding what it can and cannot do. Once AI is demystified, it is easier to assess its potential for impact, both positive and negative, as well as to interpret the data we have so far about how it affects the economy.

Artificial Intelligence, Automation and Software

      Ideally, artificial intelligence refers to intelligent machines, where by "intelligent" we mean truly smart, for example capable of reasoning. In practice, most of what we've seen in the past fifteen years, and what I believe we will continue to see in the near future, can be better

thought of as "automation on steroids." Thinking in terms of automation can guide the discussion in a more realistic direction, and also helps remind us of the limitations of AI.

Traditional automation in turn refers to physical machines or software that follow pre-specified rules or routines for interacting with its environment without real-time human direction.

One kind of automation that has expanded dramatically over the last few decades consists of a set of rules conceived of and written by humans, and executed by software. An automated telephone system that requires you to choose from a series of menus is a simple example of this; these systems evolved from phones, to websites, and more recently to chatbots. However, more complex systems have also been programmed using human-specified rules; some of the earliest self-driving cars encoded in software a set of scenarios and correct responses. Many educational technology applications that advertise their ability adapt to the level of the student simply specify what the next learning exercise should be as a function ofthe student's answers to previous questions using a large decision tree created manually by human experts (Golub Capital Social Impact Lab, 2020).

Although most people wouldn't consider a phone menu either inspiring or enjoyable, over time, automation through websites and apps have created substantial customer value and saved consumers large amounts of time and money. Being able to check a bank balance quickly on the go can prevent an overdraft fee; and when airlines finally adopted technology to pull up your flight information based on your mobile phone number, it made it much faster to make last-minute changes to a travel itinerary. Although it can be frustrating to deal with pre-specified menus when you have a non-standard request, a well-designed automated system can allow consumers to get their needs met quickly, and on their own time, which may be especially important for people working long hours or for working parents. As more and more services digitize, we expect that a larger share of services directly provided by government or indirectly funded by government will be accessed through automated systems.

This automation, while not sexy, can have an impact on employment. As more people have mobile phones and learn to use bank mobile apps for more and more functions, there is less need for a bank branch or call center with human beings available to answer questions. Indeed, Totty (2020) documented that bank teller employment declined by 26% in the last decade, from 600,000 jobs in 2010 to about 442,000 in 2019, while the Bureau of Labor Statistics projects a further decline of 15 percent between 2019 and 2029 (BLS, 2019). Numerous analysts have predicted large declines in financial services employment in the near future, including not just front-line workers but also back-office workers whose jobs involve relatively routine information processing (Kelly, 2019).

This type of automation has increased substantially over the last decade, but much of the change is not directly due to advances in artificial intelligence. Instead, in my view the most impactful changes have resulted from an expansion of digitization of interactions and recording of digital data.

As firms interact with customers digitally, use software to manage supply chains, GPS to track locations, and have digital recordings of many processes, it becomes natural for them to use software to manage the data. There is an evolution through the steps of digitization, using software to automate and process data, to more sophisticated optimization of digital processes and machine learning from data. Although sometimes firms take several steps simultaneously, it is more common to evolve slowly through these steps.

Alongside digitization, some other important trends include the rise of cloud-based services (where according to analysts the global cloud computing market increased from under $1 billion to several hundred billion dollars annually over a single decade, and today provides services to the vast majority of businesses.) The U.S. leads the world in both supply of cloud computing services and in adoption (IDC, 2019). Software as a service replaced the need for company-owned computers and specialized software to do things like manage email, customer databases, and call centers.

Even technology firms often adopt "software as a service" solutions for many purposes rather than building their own applications. This can be efficient, as software development is done once by a specialized team of developers and then used across many organizations. It is also increasingly common for different software products to interact and share data through application programming interfaces (APIs), allowing for more of a "plug and play" environment and enabling firms to take advantage of the best software for each need. This in turn makes it easier for firms to enter in both the technology and non-technology sectors, as they can purchase the products they need off-the-shelf and scale their usage as they grow. This can reduce costs associated with functions ranging from email to human resources to accounting to marketing and sales; it can also position governments and firms to more easily take advantage of innovations in AI.

Artificial Intelligence and Machine Learning

The largest innovations in AI in the past decade have concerned automation using decision rules learned from data, rather than human-specified rules. This is called machine learning. For an educational application, machine learning might be used to learn what types of reading material to recommend to a student based on past reading. With machine learning, an algorithm takes data as input and performs tasks such as prediction and classification.

For example, an analyst building a machine learning algorithm might feed in "training data" consisting of digital files that represent photos together with the corresponding labels for the photos such as "cat" and "dog." Once trained, the algorithm will take as input an unlabelled photo and output its best guess of the label (that is, its guess of whether the photo is a cat or a dog). Or, an algorithm might take in data about patients in a hospital from their electronic medical records, and output a predicted probability of death for each patient or the predicted

number of days each patient will stay, predictions that can be important in a time of scarce hospital capacity, as in the COVID-19 pandemic.

Scientists and industry practitioners have built models like this for many decades, for applications like predicting the probability that a consumer will repay a loan. However, traditionally, the analyst had to do a lot of work to determine which variables were most important, and models were often overly simplified. Recent advances in machine learning allow the analyst to just feed in raw data, and the algorithm does a lot of the work to make sense of the data. In other words, the newer tools figure out what is important from the data, and then use the important factors to make predictions. This allows the same general purpose tools to be used across a wide range of applications. These tools work best when there is plenty of data, as well as a stable environment, so that the model's performance in the future is well approximated by how well it performs on a set of held out "test data" that is hidden from the analyst when building the model.

The fact that the algorithms take over a lot of the work of refining statistical models implies that algorithms can be very "general purpose," meaning that the same general approaches can be applied to a wide range of settings (see, e.g. Agarwal, Gans, and Goldfarb (2018) or Brynjolfsson, Rock, and Syverson (2020) for further discussion of this). For this reason, we have seen a rapid diffusion of machine learning across industries and applications.

At the same time, the fact that the algorithms are general purpose and emerge from data also create weaknesses. For example, the models can be difficult to interpret and understand without substantial additional effort and analysis; they are "black boxes" that translate input into output, making it challenging for even the engineers who build them to identify potential weaknesses. The goal of training is for the algorithms to perform well on average, but predictions may be poor for some realizations of input data, which can be a problem for individual decisions with certain characteristics, but further can create more widespread difficulties if the world changes in a way that makes the realizations of certain characteristics more likely. For example, unemployment or business failure may become more likely in sectors that previously did not experience it. In general, machine learning might make predictions based on relationships that are not stable in the long run. For example, factors that predict loan default may be different during the COVID-19 pandemic than at other times. In an image classification application, cats may be associated with indoor photographs, so that an algorithm bases its classification of a photo as showing a cat versus a dog using the background of the photo rather than characteristics of the animals, where characteristics of the animals themselves are more likely to remain constant across different countries, cultures, and environments for photo sharing.

It may be difficult for the engineers that build machine learning models to even understand the extent to which their models are fragile or unstable, let alone address the problems or prevent them. Well-documented examples where image classification failed for humans with darker skin highlight the pitfalls; once identified, many technology firms addressed

the particular problems that had been pointed out, but other problems likely remain to be identified and fixed. This remains an active and important area of research in both academia and industry.

For related reasons, off-the-shelf machine learning models are often not suitable for learning about cause and effect. A small research community focuses on addressing such challenges, and indeed much of my own research focuses on how to adapt machine learning to make it more useful for scientific discovery (Athey, 2017), disentangling cause and effect (Athey, 2018; Athey and Imbens, 2019), and to make prediction more stable and reliable (Chipman, 2018; Kuang et al, 2018; Kuang et al, 2020); but the most commonly applied algorithms in practice suffer from challenges that come with the territory of black-box, general purpose machine learning.

In spite of their limitations, there are many advantages of general purpose algorithms, and the general purpose nature of machine learning is complementary with the trends of digitization, cloud computing, and open source software. There has been a dramatic increase in the availability of open source software that provides tools to manage and analyze large datasets, as well as software to implement machine learning algorithms. Although it might at first seem surprising, large technology firms have put a large number of software packages and algorithms into the public domain, including algorithms that have been trained on their own large datasets. Even a few years ago, an engineer might need to spend a lot of time to train an algorithm that classifies images, but now it is possible to use off the shelf algorithms that companies like Facebook have trained on their large datasets of pictures. That frees engineers to focus on other parts of the problem, and can make it easier for academics to do research as well as for other AI firms to enter.

For example, as a university researcher I can store terabytes of data in public cloud infrastructure for research projects, and for just hundreds or thousands of dollars, I can analyze it in ways that would have only been possible inside large corporations a decade ago. Tools exist now that manage a lot of the work for distributing workloads and optimizing performance for statistical algorithms, and I can just focus on the analytic questions without worrying so much about the computational issues. My students can use publicly available image classifiers developed by Facebook to detect objects in images. Anyone from a student to a startup to a large company can use the technology. It is easy to experiment or prototype new models. What is often scarce is not the know-how or the infrastructure, but the data or the user base.

On the other hand, some of the most computationally intensive and cutting-edge applications of AI can be too expensive for startups or university researchers to carry out due to large computing needs, or they may require the types of data only available within large technology firms. For that reason, certain types of AI research have been challenging to carry out in academia, reducing the pace of publicly available innovation in those areas (see Lazer et al, 2020 for further discussion of policy issues around access to data and infrastructure for social

science research). Some have called for a greater investment in public infrastructure for AI to support research and innovation (Etchemendy and Li, 2020).

Overall, digitization has decreased the costs of implementing AI, which reinforces the advancement of practical applications of AI and facilitates the diffusion of AI across industries and applications. Cloud infrastructure has made it easier for companies to enter and compete providing software as a service. In turn, companies providing software as a service invest in research and development, over time adding features including those powered by machine learning, thus making the benefits of machine learning available more broadly. As long as markets are competitive, a lot of those benefits accrue to the businesses and consumers who use the products. Of course, and this is a topic for another day, market power can sometimes be a concern, particularly in situations where a single firm controls access to a large group of consumers and/or their data. However, in many cases, consumers have seen the quality and convenience of digitally-delivered services rise or the costs fall as a result of all of this innovation, as digital services often have low marginal cost of delivery. This potential is what captures the imagination of AI optimists. However, there are also many unintended consequences of all of this innovation that require careful attention, as I will discuss further below.

Artificial intelligence and the generation of novel creations

An even more recent set of advances in artificial intelligence concerns algorithms that are used to create original digital objects, such as stories, art, photographs or music. It is easy to see why an observer might get the impression that AI is truly intelligent when we see a computer create an original digital image that looks like a photograph of real human. Some of the photo filters my children enjoy on their smartphones are another example of this technology, where the apps produce an image of a person that is altered to look younger or older. This type of creation is often the result of what is known as a "generative adversarial network" or GAN (Goodfellow et al, 2014). This type of AI solves a very challenging problem; it is much easier for an algorithm to distinguish between images of cats and dogs than to create an original image of a cat. Telling images of cats and dogs apart might boil down to understanding the shape of cat ears, nose, and eyes separately, without needing to specify how they relate to one another within an image; but creating a believable photograph of a cat requires generating ears, nose and eyes that make sense together in terms of proportions, locations relative to one another in the image, colors, and shadows.

GANs work by combining two distinct, dueling algorithms, the generator and the discriminator, where the discriminator has access to a dataset of objects that are labelled as real. The generator algorithm creates fake objects, and while the discriminator builds a model that classifies objects as real and fake based on a dataset of including the generator's fake objects and those that it knows are real. The algorithms battle amongst themselves. The one that creates

fake objects keeps trying to make objects that trick the latest model built by the discriminator, and the discriminator keeps trying to find different ways to distinguish real from fake. The fact that the discriminator is automated (rather than a human judge) means that the generator can try a very large number of different options, and eventually it can find something that looks good in the sense that the discriminator can't distinguish it from something fake. But much like the example you've heard that a monkey would eventually write a classic novel if it typed long enough, the algorithm creating fake stuff is not smart, it just did a lot of trial and error against a computerized judge.

GANs already have a wide array of applications (Alqahtani et al, 2019; Brownlee, 2019), and we may expect more to be introduced, as they have only gained widespread use in the last few years. GANs can be useful when it is helpful to have lots of options to select from, but it is important that the objects are original or distinct. They can generate designs for images, fabrics, music, cartoons, email subject lines, etc. that are likely to be pleasing, but typically when applied, humans screen the generated options in a second step. GANs can also can make variants of existing objects with desired characteristics, e.g. a handbag with a specified shape decorated in a certain color or style, and GANs can also convert text to an image.

In scientific applications, GANs can be used to generate fake datasets for replication of studies, to test algorithms, or to protect privacy of, e.g., health data. Constraints can be added to the GAN, so that a GAN can be trained to generate objects that have desired scientific properties. They have been used for generating potential chemical or biological compounds with useful properties, for example candidate drug compounds that might deserve further research (Schmelzer, 2019).

GANs can be problematic as well. They can be used for "deep fakes," where convincing videos or images appear to show politicians or celebrities in a negative light, as well as to generate fake comments or social media posts, creating discord and anger among citizens. They can create fake identities that can be used for financial crimes or fraud as well. The fact that GANs can produce a large number of unique fakes make them particularly useful for applications where an actor repeatedly attempts to get past anti-fraud systems or algorithms designed to screen out robots from review systems and social media.

As impressive as GANs may be, there is a big gap between being able to, say, help a customer with a non-standard request and being able to spit out sentences that mimic human sentences. The technology doesn't think or reason, it just mimics patterns.

Artificial Intelligence and Autonomy

Autonomy, whereby an algorithm governs behavior of software or hardware in reaction to a changing environment, is another key concept in artificial intelligence. A robot that navigates obstacles to deliver an object is called autonomous; and recent advances have improved the ability of autonomous agents to learn decision rules. A type of algorithm known as

reinforcement learning is designed to solve a problem like winning a board game or climbing over a wall. The algorithms actively experiment, and learn by trial and error. They break a problem down into a set of "states of the world," e.g. the location of chess pieces on a game board, or what objects are perceived by a robot in a hallway. The algorithm then considers what states it can move to from the current state, and either moves to the best option, or experiments among alternatives to learn more about what the value is of being in different states.

Reinforcement learning generally requires a lot of training in a setting where mistakes are not too dangerous, for example when a computer can play billions of games of chess against other computers in order to learn what works. As impressive as it sounds, reinforcement learning still relies heavily on pattern recognition; it isn't developing theories of the world. The core ideas behind reinforcement learning have been around a long time, but its performance has advanced dramatically in large part because of improvements in machine learning that help the algorithm automatically learn ways to effectively simplify the state space. The algorithms learn simpler representations of the possible positions of chess pieces; just as chess students learn a point system to determine how good their position is in a game, the algorithms assign scores to different positions and suggest game moves that optimize the scores. The AI accomplishments are especially impressive when there are billions of billions of possible configurations of a game board, since human brains can't remember that much information, but fundamentally the algorithms are still just adding up wins and losses using a scoring system learned from many replications of the game. Algorithms created in this way might perform arbitrarily badly if something small changes about the environment, and they are only as good as the data they have seen in the past, whether through simulated game play, their own experiments in previous games, or data from past games played by humans. Thus, they still may not do well in never-experienced environments.

Artificial Intelligence Applications and Impact

We have seen that the most commonly used categories of AI either implement human-created decision rules, or make use of pattern recognition to derive predictions or decision rules. Despite the inherent limitations of these categories of AI, some awe-inspiring accomplishments have emerged; yet, it is important not to conclude from the accomplishments to date that some sort of general intelligence will follow quickly. As discussed above, pattern recognition has limitations, particularly performance in changing circumstances.

Machine learning and AI have had especially widespread adoption for applications when it is possible to update the model more quickly than the environment changes, and when prediction and classification tasks play an important role (Agrawal, Gans, and Goldfarb, 2018). Examples of applications include digitizing input from people, including handwritten forms and voice recognition. Image classification has many applications across industries ranging from finance (e.g. images for insurance claims) to medicine (diagnosis). For example, AI has been

used in medical imaging, detecting diseases or anomalies with specialist-level accuracy in some cases (Ruamviboonsuk et al., 2019). It is also used to predict the likelihood of a chemical reaction (Hao, 2020).

Machine learning and AI can be particularly useful in a setting where, in the absence of AI, human workers would need to make decisions in limited time. Although in principle, a human worker could make a better decision than an algorithm with sufficient investment of time, they might not have enough time to gather and absorb all of the relevant data in practice. An example is resume screening, where often a human screener might only briefly scan a resume before deciding if it should be prioritized for further consideration. The screener might not know all relevant information about the quality of the secondary or university school attended by the applicant, and they might not be familiar with the skills required for work at a particular firm in the worker's employment history. In contrast, an algorithm might be trained on thousands of resumes, and can quickly process a large set of characteristics of the individual. Of course, it is important to consider issues of bias that can arise when using algorithms in this way; careful attention to training data and active investment in expanding the pool of workers who advance to higher levels of the interview process are examples of approaches that may mitigate these issues.

A related application is in automation of worker screening for qualifications and availability. Alain Dehaze, CEO of Adecco, recently reported that his firm quickly recruited 16,000 workers in Europe using a purely digital process during the COVID-19 pandemic (Michaels, 2020).

AI is also used to prioritize resource allocation according to the risk posed by an individual or entity. It has been used in a variety of government applications, including the allocation of health inspectors to restaurants or home visits by child protective services workers (Schwartz et al, 2017). AI has also been used to help judges make bail decisions (Kleinberg et al, 2018). In principle, they can improve decision quality, accountability, and equity, as they replace rapid decisions by humans who may not have the time to consider all relevant information. However, it should also be obvious that these types of applications require great care to implement fairly and effectively (Glaberson, 2019), and more work is needed both by academics and regulators to ensure that implementations of these algorithms follow best practices and are evaluated for unintended consequences.

Applications of reinforcement learning include digital marketing, where reinforcement learning is commonly used to figure out which of many email subject lines or headlines work best to attract consumer clicks; and autonomous drones, robots, and delivery vehicles.

More broadly, according to McKinsey's 2019 AI impact survey across hundreds of firms, AI had been adopted in nearly every industry by 2018. Retail is growing most quickly, with 60 percent of respondents from retail reported that their companies have embedded at least one AI capability in one or more functions or business units, a 35-percentage-point increase from the 2018 survey (Cam et al., 2019). Funding of AI startups in the U.S. spans a wide range of industries. AI startups received $19.8B of investment in 2019, with top focus areas including

Data Tools (8.1% of all startups); Medical Technology (5.3%); Fashion and Retail Technology (4.7%); Text Analytics (4.7%), and Chatbots (3.9%) (Perrault et. al 2019, p. 92).

Impact on the Efficiency and Accessibility of Service Delivery

The impacts of automation and AI arise not just in terms of the cost of providing services, but also in the costs of receiving services. In particular, individuals seeking services like health and education often bear substantial costs in terms of transportation and time, which translate into lost income, outlays for child care, or lost sleep. For those with full-time jobs or caregiving responsibilities, it can be extremely frustrating and create economic hardship to waste time in waiting rooms, lines, or traveling in order to access services.

Providing services digitally allows people to access services when it is convenient, avoiding the need to take time off of work or obtain child care. Otherwise unused blocks of time can be used to accomplish tasks such as filling out forms or acquiring job-relevant skills, and a parent might engage in these activities while children are sleeping to avoid the need for child care. Citizens may be more satisfied with digital services in many cases, and lower income consumers in particular can benefit from improved access (Kuziemski, M., & Misuraca, G., 2020). For example, Medicaid patients may experience long travel times or waits to access health care, while providing access to telehealth has shown promising results in pilot studies (Koehn, 2016), and the widespread adoption of telehealth during COVID has opened up new possibilities to expand this access while simultaneously reducing costs. Access to high-quality medical care in rural areas is another important application.

In contrast, the lack of ability to reach consumers digitally can interfere with efficient provision of services. For example, poor IT infrastructure hampered state governments in their efforts to provide COVID relief to individuals quickly, and it limited the scope of early programs to deal with the pandemic (Bollag and Wilner, 2020). Even where digital provision was enabled, as in telemedicine, concerns remain about inequality of access (Weigel et al, 2020; Anderson and Kumar, 2017). This type of infrastructure and access must be addressed for governments to be able to take the next step and optimize service provision using artificial intelligence.

It is perhaps more straightforward to assess the potential for efficiency benefits in terms of reducing the cost of providing services through digitization and automation, which (after technology is developed) reduces the number of government workers needed to provide a service. This reduction in employment is obviously a challenge for the affected workers, but it increases the efficiency of government. In some cases, government funds can be reallocated to provide additional services in areas where it is harder for automation to substitute for human workers, as in the area of education or child care.

A number of studies have attempted to assess the state of digitization as well as the efficiency gains that are possible for governments (see, e.g. analyses of the degree of digitization of government services in the US (OECD, 2019); the extent of potential efficiency and cost

savings (Eggers et al, 2017); citizen preferences and satisfaction with government services (BCG, 2018); the use of AI in government services in developing countries (Lauron and Stamboel, 2018)).   Deloitte estimates that automation of federal government employee tasks could save between 96.7 million and 1.2 billion hours annually, with potential savings between $3.3 billion and $41.1 billion (Eggers et al, 2017). Categories of potential benefits include reducing the labor cost of providing government services and avoiding corruption.

There are some challenges to undertaking large-scale digitization programs. First, the infrastructure underlying digitization has the feature of requiring up front investments that pay off over years, which can be difficult for companies, let alone governments, to manage well, especially in an environment of changing technology. There are large risks for governments undertaking large IT projects, although there are also substantial risks of retaining antiquated systems, as described in testimony before this committee in July (Gerton, 2019). It is also important to remain attentive to inequality in IT skills, access and adoption that make for an uneven playing field when it comes to accessing digital services (Anderson and Kumar, 2017).

Impact on Productivity and Measurement Challenges

Like many previous technologies, it is hard to isolate the impact of AI in productivity numbers. Despite impressive improvements in AI, not to mention many other technologies, productivity growth has actually slowed down in the last fifteen years, from an average of over 2.4% per year between 1995-2005 to less than 1.3% per year since then (Brynjolfsson, Rock and Syverson, 2019). That is to say, the data appears to rule out very large productivity gains from AI or digitization in general. One explanation is that realizing the potential of new technologies "requires large intangible investments and a fundamental rethinking of the organization of production itself. Firms must create new business processes, develop managerial experience, train workers, patch software, and build other intangibles. This raises productivity measurement issues because intangible investments are not readily tallied on a balance sheet or in the national accounts." (Brynjolfsson, Rock and Syverson, 2020, p.2) In addition, all of the investments required to achieve the full potential of AI may be characterized by up front investments whose benefits take many years to realize, making it more difficult to isolate the impact.

Opportunities

Several types of opportunities stand out where AI can be an important part of the solution to societal challenges. A first area is finance. The financial services sector is on the front line of employment impacts from AI, but there are also a wide range of opportunities for technology to improve access to services and to reduce costs to the point that more low-income consumers can be served. However, challenges remain, in part because the regulatory framework for our financial system is designed around regulating processes followed by human workers, not regulating algorithms and automated systems. Many financial regulators lack experience and expertise in AI, and the existing regulatory structure is not designed to handle situations where

algorithms have a small but positive chance of making mistakes. Cost-benefit analysis is not the primary framework used by regulators to evaluate AI projects, and firms slow their investment in the face of regulatory uncertainty. Financial technology has the potential to be very progressive, reducing costs of serving low-income consumers and small businesses, providing opportunities for new startups, and introducing competition. Automation can ensure compliance and avoid a role for human biases, while well-designed algorithms can avoid introducing new forms of bias. A regulatory framework needs to address the benefits as well as the risks of new technology and innovation. Of course, risks are real, as AI can be used to automate the creation of exploitive or manipulative marketing as well as financial crimes (FINRA, 2020).

A second area where AI can be quite impactful is in the area of education and training. The COVID-19 pandemic dramatically accelerated adoption of a wide range of educational technology applications (Gilchrist, 2020), ranging from preschool applications like Khan Kids or IntellectoKids (Kotlov, 2020) to upskilling programs. Companies like Coursera expanded their free offerings and partnered with governments to allow unemployed workers to acquire new skills, such as basic technical skills identified by employers (Training Industry, 2020). The United States could do more to study the impact of such programs, and if effective, scale up access to a larger group of unemployed Americans.

Many educational technology companies are in a relatively early stage of growth. However, with the new, larger user bases they have attracted during the pandemic, companies have the opportunity to introduce more automation and artificial intelligence into the learning process, personalizing educational experiences, adapting to the student, and also potentially innovating in certification and interview practice (as artificial intelligence can be used evaluate students or create interactive scenarios for students to practice their skills). Given its outsized impact on society, it will be important to continue to nurture the industry providing education and training digitally, and to consider complementary investments that reduce frictions for workers in job transitions.

More generally, the changes brought by AI will contribute to increasing rates of worker transitions between jobs. As automation changes the organization of work, some jobs are likely to be eliminated, and newly created jobs may require different skills. If, as in the case of call centers, a large number of jobs in a single category is eliminated in a short period of time, workers may need to transition to a different type of work. The use of cloud computing and software as a service may lead to a variety of employers in the same industry automating processes at the same time, making the transitions more challenging for affected workers. A variety of frictions might interfere with their ability to transition.

Simon (2020) reviews research on labor market frictions, and finds that many causes of frictions are related to individuals needing to adapt by (1) learning new skills or (2) moving to new geographies with better opportunities. There is potential for AI and digitization more generally to provide solutions to both of these frictions, (1) by providing convenient, accessible, enticing and personalized education and training solutions, (2) reducing the importance of

geography by providing remote work opportunities, and (3) improving worker access to information and support in order to guide workers to better decisions and reduce the risk associated with transitions.

A variety of evidence (Chen et al, 2019) also points to the importance of flexibility of schedules for many workers, particularly workers who need to balance school or caregiving responsibilities with work. Remote work and technology-related work may lend themselves to flexibility, and the new approaches to work developed during the COVID-19 pandemic may enable more such flexibility. Other policies that complement worker flexibility may also become important, for example policies that facilitate continuous access to health care while workers combine retraining and part time work.

Artificial Intelligence for Worker Safety and Monitoring

Another category of application of AI concerns monitoring the activities and experiences of workers. Although the primary motivation for the introduction of AI might be worker safety or regulatory compliance, there are a number of potential unintended consequences.

First, consider some of the ways in which AI can be used to monitor workers. Worker communication is increasingly digital, as it takes place through email, company-sponsored chat platforms, or over digital conferencing systems. A variety of companies provide services designed to translate conference calls into written notes, and it is straightforward to use machine learning to learn to create risk scores for written communication. The relevant type of risk to be predicted may vary by industry, but, for example, AI is already being productively deployed to monitor communication within financial services companies to ensure compliance with regulations, such as prohibitions on insider trading or other illegal activities (IBM and Chartis Research, 2018; FINRA, 2020). Indeed, the term "regtech" has been coined for technology that helps firms comply with regulations, something that ultimately will be important to consider when thinking about cost-benefit analysis from regulations. Some regulations may become easier and cheaper to comply with and enforce in the digital era, especially when compliance processes can be automated and included in software as a service provided to firms in a given industry. On the other hand, compliance failures may occur at a larger scale when many firms use the same software.

It is also relatively straightforward for an employer to create a training dataset for a machine learning algorithm based on video or audio recordings of employee interactions with customers. A sample of video or audio can be watched and manually scored by human judges to create a labelled training dataset. After training, an algorithm can be applied to all video, creating a "predicted customer satisfaction score" for every minute of every recorded interaction. This allows large-scale, low-cost scoring of workers. Another application would be safety violations. Video can be labelled, or assembly line accidents or errors can be used as "labels" in training datasets for algorithms designed to identify undesirable actions by workers. Once trained, the algorithms can be applied to all recorded video, and those video segments that

received poor scores could be manually reviewed to assess whether the incident was, in fact, an example of undesirable employee behavior. This dramatically increases the efficiency of worker monitoring by focusing human time on the portions of video most likely to be associated with problems.

A related example concerns safety for fleets of drivers. A variety of companies sell software that is installed on the mobile phones of drivers, where drivers are required to use the software during working hours. The software uses the telemetry from the phone to identify whether the driver was speeding, or whether the other safety violations took place. Drivers with poor safety can be assigned to safety courses, or if necessary, terminated.

Similar algorithms could in principle be used in governments to ensure that government workers provide proper service, follow rules, and don't take bribes, similar to the use of body cameras for police. This type of application is in an earlier stage, but there have been some promising case studies in recent years (Aarvik, 2019).

However, as a society we will need to consider the broader implications and substantial risks that arise alongside the large scale monitoring of workers. A variety of problems might emerge, ranging from bias, as might occur if algorithms perform poorly on women, shorter people, or people with darker skin; to the potential that the existence of all of this tracking can create fodder for blackmail or extortion. It is hard to imagine that nothing embarrassing ever happens if you are constantly monitored as a worker or as a citizen. Privacy and security of data is already important to individuals in their leisure time, but it takes on different considerations in the workplace.

A final issue worth considering when AI is used to augment humans in the workforce is the extent to which reliance on AI reduces human attentiveness. Just as humans might "fall asleep at the wheel" when using an autopilot feature in a car, workers may fail to gain experience in certain types of tasks when the tasks are automated, or they may be insufficiently incentivized to pay attention and gather information when AI is introduced to assist them with tasks. Thus, it is important for firms and governments to consider the way in which AI changes both the information and incentives for workers when it is introduced into organizations and decision processes (Athey, Bryan, and Gans, 2020).

Additional Policy Considerations

A common mistake in thinking about AI is to focus on the harm to a small group of people, without considering the benefits that accrue to many. Many AI applications can be progressive, because automation and digital provision of services reduce marginal cost. On the other hand, technical change is taking place in a context where the costs of essentials such as housing and health care are rising faster than wages, and some of the consumer products and services made available through technology are consumed by wealthier individuals. It is thus important to consider overall effects as well as the context in which the innovation occurs.

AI raises a variety of additional opportunities and challenges for our economy. One is a misalignment of incentives by firms in the type of research and development that is prioritized. Firms have incentives to invest in cost reduction, but don't consider broader societal consequences; thus, they may be more likely to invest in labor-replacing technology. Research that helps explore labor augmenting technology may receive insufficient investment by the private sector. The Stanford Institute for Human-Centered Artificial Intelligence, where I am associate director, is attempting to prioritize research on labor-augmenting technology, but much more investment in this type of research and development would be beneficial and more aligned with societal objectives.

Another set of considerations concerns the geographic distribution of research, entrepreneurship, and investment. Historically, the United States has been a leader in AI innovation, but there is no guarantee that will continue without support for universities as well as high-skilled immigration. Given that some parts of the AI industry are prone to concentration due to scale economies, it can be important that the U.S. retains its leadership position in innovation if the U.S. hopes to be the home of the world's leading companies in the future. Furthermore, there is no guarantee that jobs created by AI will have broad geographic distribution within the U.S. For example, Bloom et al. (2019) suggests that between 1990 and 2007, large multinational firms offshored their production while creating new service-sector jobs. However, the lost jobs were in the U.S. heartland, while the new service jobs were created in high-education areas along the coasts, and were therefore taken up by very different people. Despite this shift in the location of jobs, geographic mobility has not been a widespread way of adjusting to the shock (Autor, Dorn, and Hanson, 2013; Autor et al., 2014).

Gruber and Johnson (2019) argue that universities, especially those with large medical centers, can anchor cities that are large enough to grow and create jobs, and that research and development investment (including investment funded by the government) can be effective at spurring the innovation that drives this growth. With high cost of living and congestion in the existing tech centers, together with tech firms embracing remote work, there may be opportunities for timely investments in these locations that contribute to their success, expanding the set of geographies where firms open offices with AI jobs. Technology companies consider the size of the engineering workforce as well as the opportunities to hire from universities when locating new satellite offices, and it has been common for veterans of larger technology companies to start new ventures after gaining experience, which can seed new communities of entrepreneurs outside the traditional technology hubs.

A further set of challenges facing the U.S. is related to demographics. Overall, our working population is aging. Several countries, including the United States, may face fiscal challenges due to demographics unless they increase either the birth rate or immigration. Varian (2018) did analysis to suggest that using the most aggressive estimates for the impact of automation, we are still more likely to face a worker shortage than surplus in the coming decades. He points out that it may be hard to predict the future impact of AI, but we have a lot

of certainty about how many 40 year olds there will be in 20 years, and our population is very likely to be much older. According to the Congressional Budget Office, the federal government spent about one-third of its budget on seniors in 2005 (Congressional Budget Office, 2019, pp. 12-14). By last year, the share grew to 40 percent, or $1.5 trillion. The share is forecast to rise to half of all non-interest spending, or $3 trillion, by 2029. This amounts to spending 10 percent of the nation's Gross Domestic Product on older adults.

Countries with aging workforces tend to invest more in automation. To mitigate demographic challenges, we may want to consider investments in research and development that supports the efficient provision of services to older populations, perhaps assisted by technology to make services safer and more affordable. In addition, AI can be used to augment older workers, making it possible for older workers to have second careers without compromising their health. For example, robots can help workers with challenging physical tasks, while AI can assist with tasks that traditionally required memory or attention to detail, as it can warn of anomalies and monitor performance in real time. This removes some of the more challenging components of work, allowing older workers to focus on activities that emphasize human interaction and provide fulfillment and stimulation.

Conclusions

As AI is adopted through government and the economy, it will be important for governments to keep a close eye on the myriad challenges raised by AI. Some of them have been raised in my testimony today, including bias, privacy, security, investment risk, reliability and fragility of machine learning models, and the need for thoughtful regulation that includes cost-benefit analysis. Other challenges, such as concerns about market power of large firms who have unique access to data, are important, but beyond the scope of this testimony.

In the coming decades, as AI plays a larger and larger role in our economy and in the provision of government services, transitions will be the new normal for workers. It will be important for governments to address the challenges faced by workers, and also to make the necessary investments to ensure that AI fulfills its potential as part of the solution. Since digitally provided services have low marginal cost of delivery, AI is well positioned to contribute to providing scalable and effective education, training, and access to government-provided and government-funded services.

References

Agrawal, A., Gans, J. and Goldfarb, A., 2018. *Prediction machines: the simple economics of artificial intelligence*. Harvard Business Press.

Alqahtani, H., Kavakli-Thorne, M. and Kumar, G., 2019. Applications of generative adversarial networks (gans): An updated review. *Archives of Computational Methods in Engineering*, pp.1-28.

Anderson, M. and Kumar, M., 2017. Digital divide persists even as lower-income Americans make gains in tech adoption. *Pew research center*, *22*.

Aarvik, P., 2019. Artificial Intelligence a promising anti-corruption tool in development settings. *U4 Report* 2019:1. https://www.u4.no/publications/artificial-intelligence-a-promising-anti-corruption-tool-in-development-settings.pdf

Athey, S., 2017. Beyond prediction: Using big data for policy problems. *Science*, *355*(6324), pp.483-485.

Athey, S., 2018. The impact of machine learning on economics. In *The economics of artificial intelligence: An agenda* (pp. 507-547). University of Chicago Press.

Athey, S. and Imbens, G.W., 2019. Machine learning methods that economists should know about. *Annual Review of Economics*, *11*, pp.685-725.

Athey, S.C., Bryan, K.A. and Gans, J.S., 2020, May. The allocation of decision authority to human and artificial intelligence. In *AEA Papers and Proceedings* (Vol. 110, pp. 80-84).

BCG, 2018. Realizing the Power of Digital Government, *Digital Government Citizen Survey* 2018. https://www.bcg.com/en-us/industries/public-sector/digital-transformation-technology

Bloom, N., Kurmann A., Handley, K. and Luck, P., 2019. The Impact of Chinese Trade on U.S. Employment: The Good, The Bad, and The Debatable, *Stanford mimeo*.

Bureau of Labor Statistics, U.S. Department of Labor, *Occupational Outlook Handbook*, Tellers, on the Internet at https://www.bls.gov/ooh/office-and-administrative-support/tellers.htm

Brownlee, J., 2019. 18 Impressive Applications of Generative Adversarial Networks (GANs) *Machine Learning Mastery*, June 14, 2019.
https://machinelearningmastery.com/impressive-applications-of-generative-adversarial-networks/

Brynjolfsson, E., Rock, D., & Syverson, C. (2019). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.) *The economics of artificial intelligence: An agenda* (pp.23– 57). National Bureau of Economic Research Conference Report. University of Chicago Press.

Brynjolfsson, E., Rock, D. and Syverson, C., 2020. The productivity J-curve: How intangibles complement general purpose technologies, *AEJ: Macroeconomics* (forthcoming).

Cam, A., Chui, M. and Hall, B., 2019. Global AI Survey: AI proves its worth, but few scale impact. *McKinsey Analytics.*

Chen, M.K., Rossi, P.E., Chevalier, J.A. and Oehlsen, E., 2019. The value of flexible work: Evidence from uber drivers. *Journal of Political Economy*, *127*(6), pp.2735-2794.

Chipman, Ian, 2018. Susan Athey: Why Business Leaders Shouldn't Have Blind Faith in AI. *Insights by Stanford Business*.
https://www.gsb.stanford.edu/insights/susan-athey-why-business-leaders-shouldnt-have-blind-faith-ai

Congressional Budget Office., 2019. The budget and economic outlook: 2019 to 2029. Washington, DC: Congress of the United States, Congressional Budget Office.

Eggers, W.D., Schatsky, D. and Viechnicki, P., 2017. AI-augmented government. Using cognitive technologies to redesign public sector work. *Deloitte Center for Government Insights*.
https://www2.deloitte.com/us/en/insights/focus/cognitive-technologies/artificial-intelligence-government.html

Etchemendy, J. and Li, F., 2020. National Research Cloud: Ensuring the Continuation of American Innovation.
https://hai.stanford.edu/blog/national-research-cloud-ensuring-continuation-american-innovation

FINRA, 2020. Artificial Intelligence (AI) in the Securities Industry, *FINRA RegTech White Paper*.
https://www.finra.org/rules-guidance/key-topics/fintech/report/artificial-intelligence-in-the-securities-industry

Gerton, T.W., 2019. Testimony of Teresa W. Gerton President and Chief Executive Officer National Academy of Public Administration Before the Budget Committee U. S. House of Representatives July 15, 2020.
https://budget.house.gov/sites/democrats.budget.house.gov/files/documents/Gerton_Testimony.pdf

Gilchrist, K. 2020. These millennials are reinventing the multibillion-dollar education industry during coronavirus, *CNBC*, June 8, 2020.
https://www.cnbc.com/2020/06/08/edtech-how-schools-education-industry-is-changing-under-coronavirus.html

Golub Capital Social Impact Lab, 2020. What really is 'adaptive learning'? *Medium*, January 22, 2020. https://medium.com/@gsb_silab/what-really-is-adaptive-learning-68e64c1f78c3

Glaberson, S.K., 2019. Coding over the cracks: predictive analytics and child protection. *Fordham Urban Law Journal*, *46*, p.307.

Goodfellow, Ian, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. "Generative adversarial nets." *In Advances in neural information processing systems*, pp. 2672-2680. 2014.

Gruber, J. and S. Johnson, 2019. *Jump Starting America. How Breakthrough Science Can Revive Economic Growth and the American Dream,* Public Affairs.

Hao, K., 2020. IBM has built a new drug-making lab entirely in the cloud. MIT Technology Review, August 28, 2020.
https://www.technologyreview.com/2020/08/28/1007737/ibm-ai-robot-drug-making-lab-in-the-cloud/

Kelly, J. 2019. From Layoffs To Declining Revenues, Banks Are In Trouble: Here Are Some Hiring Bright Spots, *Forbes*, September 12, 2019.
https://www.forbes.com/sites/jackkelly/2019/09/12/wall-street-banks-are-in-trouble-employees-axed-and-revenues-down-heres-some-hiring-bright-spots/#6f7b5de6dd44

Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J. and Mullainathan, S., 2018. Human decisions and machine predictions. *The Quarterly Journal of Economics*, *133*(1), pp.237-293.

Kuziemski, M., & Misuraca, G., 2020 . AI governance in the public sector: Three tales from the frontiers of automated decision-making in democratic settings. *Telecommunications Policy,* 44(6), 101976. https://doi.org/10.1016/j.telpol.2020.101976

Koehn, A. 2016. Apps and Technology Help Low-Income People Access Healthcare. *Medill Report Chicago*, March 1, 2016. https://news.medill.northwestern.edu/chicago/apps-and-technology-help-low-income-people-access-healthcare/

Kotlov, M., 2020. Is the Rise of Preschoolers' App Usage a Pandemic Boom or a Paradigm Shift?, *Hackernoon*, July 31, 2020. https://hackernoon.com/is-the-rise-of-preschoolers-app-usage-a-pandemic-boom-or-a-paradigm-shift-lh8l3eqx

Kuang, K., Cui, P., Athey, S., Xiong, R. and Li, B., 2018, July. Stable prediction across unknown environments. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (pp. 1617-1626).

Kuang, K., Xiong, R., Cui, P., Athey, S. and Li, B., 2020, January. Stable Prediction with Model Misspecification and Agnostic Distribution Shift. In *AAAI* (pp. 4485-4492).

Lauron, C. and Stamboel, I. 2018. Digitalization of government - Suits The *C*-Suite, *Business World*, February 18, 2018. https://www.bworldonline.com/digitalization-of-government/

Lazer, D.M., Pentland, A., Watts, D.J., Aral, S., Athey, S., Contractor, N., Freelon, D., Gonzalez-Bailon, S., King, G., Margetts, H. and Nelson, A., 2020. Computational social science: Obstacles and opportunities. *Science*, *369*(6507), pp.1060-1062.

Michaels, D., 2020. Remote Work Forever? Not So Fast, Jobs Guru Says, *The Wall Street Journal*, June 10, 2020. https://www.wsj.com/articles/remote-work-forever-not-so-fast-jobs-guru-says-11591790405

IBM and Chartis Research, 2018. AI in RegTech: A Quiet Upheaval. https://www.ibm.com/downloads/cas/NAJXEKE6.

IDC 2019, Worldwide Public Cloud Services Spending Forecast to Reach $210 Billion This Year, According to IDC, https://www.idc.com/getdoc.jsp?containerId=prUS44891519.

Totty, M., 2020. Rapid Change Is Coming to the Service Sector, *UCLA Anderson Review*, July 08, 2020.
https://www.anderson.ucla.edu/faculty-and-research/anderson-review/service-industrialization.

Training Industry, 2020. CINDE Announces: Coursera And Costa Rica Launch A Joint Program To Strengthen Industry 4.0 Skills And Train – Free Of Charge – 50,000 People To Confront The COVID-19 Crisis. June 16, 2020.
https://trainingindustry.com/press-release/it-and-technical-training/cinde-announces-coursera-and-costa-rica-launch-a-joint-program-to-strengthen-industry-4-0-skills-and-train-free-of-charge-50000-people-to-confront-the-covid-19-crisis/

OECD, 2019. Going Digital: Shaping Policies, Improving Lives, *OECD Publishing,* Paris,
https://doi.org/10.1787/9789264312012-en.

Schmelzer, R. 2019. 3 GAN use cases that showcase their positive potential. *E-Handbook: Neural network applications in business run wide, fast and deep*, August, 20, 2019.
https://searchenterpriseai.techtarget.com/feature/3-GAN-use-cases-that-showcase-their-positive-potentia

Schwartz, I.M., York, P., Nowakowski-Sims, E. and Ramos-Hernandez, A., 2017. Predictive and prescriptive analytics, machine learning and child welfare risk assessment: The Broward County experience. *Children and Youth Services Review*, *81*, pp.309-320.

Simon, L. 2020. Sources of Labor Market Upheaval - and How the Workforce Can Adapt to Constant Change. *Medium*, August 17, 2010.
https://medium.com/@gsb_silab/sources-of-labor-market-upheaval-and-how-the-workforce-can-adapt-to-constant-change-b1d4366a9eba.

Perrault, R., Shoham, Y., Brynjolfsson, E., Clark, J., Etchemendy, J., Grosz, B., Lyons, T., Manyika, J., Mishra, S. and Niebles, J.C., 2019. *The AI Index 2019 Annual Report*. AI Index Steering Committee, Human-Centered AI Institute, Stanford University, Stanford, CA.

Raghu, M., Blumer, K., Corrado, G., Kleinberg, J., Obermeyer, Z. and Mullainathan, S., 2019. The algorithmic automation problem: Prediction, triage, and human effort. *arXiv preprint arXiv:1903.12220*.

Ruamviboonsuk, P., Krause, J., Chotcomwongse, P., Sayres, R., Raman, R., Widner, K., Campana, B.J., Phene, S., Hemarat, K., Tadarati, M. and Silpa-Archa, S., 2019. Deep learning

versus human graders for classifying diabetic retinopathy severity in a nationwide screening program. *NPJ digital medicine*, 2(1), pp.1-9.

Varian, H., 2018. Automation vs. Procreation.
http://conference.nber.org/conferences/2018/AIf18/f114523.slides.pdf.

Webb, M., 2019. The Impact of Artificial Intelligence on the Labor Market. *Available at SSRN 3482150.*

Weigel, G., Ramaswamy, A., Sobel, L., Salganicoff, A., Cubanski, J. and Freed, M., 2020. Opportunities and barriers for telemedicine in the US during the COVID-19 emergency and beyond. Washington, DC: Kaiser Family Foundation, May, 11, 2020.
https://www.kff.org/womens-health-policy/issue-brief/opportunities-and-barriers-for-telemedicine-in-the-u-s-during-the-covid-19-emergency-and-beyond/